



Modeling UV Radiation Feedback from Massive Stars. I. Implementation of Adaptive Ray-tracing Method and Tests

Jeong-Gyu Kim^{1,2}, Woong-Tae Kim¹, Eve C. Ostriker², and M. Aaron Skinner^{2,3}

¹ Department of Physics & Astronomy, Seoul National University, Seoul 08826, Korea; jgkim@astro.snu.ac.kr, wkim@astro.snu.ac.kr

² Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA; eco@astro.princeton.edu

³ Lawrence Livermore National Laboratory, 7000 East Ave., Livermore, CA 94550-9234, USA; skinner15@llnl.gov

Received 2017 July 11; revised 2017 November 13; accepted 2017 November 15; published 2017 December 15

Abstract

We present an implementation of an adaptive ray-tracing (ART) module in the *Athena* hydrodynamics code that accurately and efficiently handles the radiative transfer involving multiple point sources on a three-dimensional Cartesian grid. We adopt a recently proposed parallel algorithm that uses nonblocking, asynchronous MPI communications to accelerate transport of rays across the computational domain. We validate our implementation through several standard test problems, including the propagation of radiation in vacuum and the expansions of various types of H II regions. Additionally, scaling tests show that the cost of a full ray trace per source remains comparable to that of the hydrodynamics update on up to $\sim 10^3$ processors. To demonstrate application of our ART implementation, we perform a simulation of star cluster formation in a marginally bound, turbulent cloud, finding that its star formation efficiency is 12% when both radiation pressure forces and photoionization by UV radiation are treated. We directly compare the radiation forces computed from the ART scheme with those from the M_1 closure relation. Although the ART and M_1 schemes yield similar results on large scales, the latter is unable to resolve the radiation field accurately near individual point sources.

Key words: H II regions – methods: numerical – radiation: dynamics – radiative transfer – stars: formation

1. Introduction

Radiative feedback plays a vital role in regulating star formation on various scales (e.g., Krumholz et al. 2014). On small scales, radiation from (accreting) protostars raises the temperature of the surrounding gas, affecting the evolution of individual accretion disks and suppressing fragmentation of the dense cluster-forming gas into very low mass objects (e.g., Krumholz 2006; Whitehouse & Bate 2006; Offner et al. 2009). On intermediate scales, ultraviolet (UV) radiation emitted by massive stars not only reduces cold neutral gas available for star formation via photoionization and photodissociation (Whitworth 1979; Williams & McKee 1997; Matzner 2002; Krumholz et al. 2006) but also produces thermal and radiation pressure that controls the dynamics of H II regions by inducing expansion (Krumholz & Matzner 2009; Murray et al. 2010; Lopez et al. 2011; Kim et al. 2016). Ionizing radiation that escapes from star-forming regions also produces diffuse ionized gas in the Milky Way and other galaxies (e.g., Haffner et al. 2009) and contributes to reionization of the intergalactic medium (IGM) in the universe at high redshift (e.g., Barkana & Loeb 2001). On large scales in galaxies, far-UV radiation from young OB associations is the dominant heating source of the diffuse neutral interstellar medium (ISM) via the photoelectric effect on small dust grains (e.g., Wolfire et al. 2003). This controls the thermal pressure that contributes to supporting the vertical gravitational weight of the ISM in galactic disks and thus represents an important feedback loop that self-regulates star formation (e.g., Ostriker et al. 2010; Kim et al. 2013). Therefore, it is essential to follow effects of radiative feedback to address questions such as star formation efficiencies and the mechanisms of cloud destruction for giant molecular clouds (GMCs), as well as a wide variety of other physical issues in the ISM and IGM.

Since star formation and feedback involve highly nonlinear processes, radiation hydrodynamic (RHD) simulations have become an indispensable tool in understanding the impact of radiation on star cluster formation occurring in turbulent clouds. For example, several recent studies relied on numerical simulations to investigate the effects of photoionization (Dale et al. 2012, 2013; Walch et al. 2012; Howard et al. 2016; Gavagnin et al. 2017) and radiation pressure from dust-reprocessed infrared (Skinner & Ostriker 2015) or stellar UV (Raskutti et al. 2016) radiation on cloud dispersal. Dale et al. (2014) and Dale (2017) studied the combined effects of ionizing radiation with stellar winds. More recently, Geen et al. (2016), Grudić et al. (2016), and Shima et al. (2016) explored how the radiation feedback works together with supernova explosions to destroy parent molecular clouds.

To treat radiation feedback properly, it is important to solve radiative transfer (RT) equations accurately and efficiently. Despite increasing demand for RHD simulations, RT still remains numerically challenging for a number of reasons: high dimensionality, nonlocal and multiscale behavior of interactions between radiation and matter, and difficulty in choosing a suitable frame for evaluating radiation and fluid variables, to name a few (Mihalas & Auer 2001; Castor 2004). A common approach for the numerical solution of RT problems is to take the angular moments of the transfer equation and adopt a closure relation to truncate the hierarchy of moments. As the time-dependent moment equations can be written as hyperbolic conservation laws, high-order Godunov methods are often adopted to solve the time-dependent RT problem.

In astrophysics, the most widely used moment method is the flux-limited diffusion approximation, in which the radiation flux is calculated by taking a local gradient of the radiation energy density (Levermore & Pomraning 1981; Krumholz et al. 2007a; González et al. 2015), applying a limiter if the gradient is very steep to prevent superluminal transport. RT

solvers based on flux-limited diffusion are suitable for describing radiation fields in optically thick fluids. However, they are of limited accuracy in treating an optically thin medium with a complex source distribution, as well as in casting shadows in the transition zones where the optical depth varies significantly (e.g., Hayes & Norman 2003; Davis et al. 2014).

Another approach for RT is to use the M_1 closure relation that assumes that the intensity is invariant under rotation about the direction of radiation flux (e.g., Levermore 1984; González et al. 2007; Rosdahl et al. 2013; Skinner & Ostriker 2013). For a single point source, the M_1 closure model can correctly describe the radiation field in both the optically thin and thick regimes (see Skinner & Ostriker 2013). When there are multiple sources distributed widely, however, M_1 can fail in the optically thin regime. For example, when two beams going in different directions interact with each other, they unphysically merge rather than crossing (e.g., Frank et al. 2012), since the local flux is assumed to be unidirectional. For diffuse radiation in systems with both optically thin and thick regions, a more accurate closure relation can be obtained from the formal solution of the RT equation using the multidirectional method of short characteristics (Davis et al. 2012). For point-like sources that have strong angular variations in emissivity, however, artificial anisotropic structure can arise from inaccuracy in interpolating intensity over neighboring cells in the short-characteristics method (Finlator et al. 2009).

For problems in which the total emission is dominated by a small number of point sources (and in which the reprocessed diffuse radiation is negligible), one may directly integrate the RT equation by calculating the column density (or, equivalently, the optical depth) to every point in the simulation domain starting from the sources. A method utilizing short characteristics calculates the column densities from the sources by performing upwind interpolation over cells (Mellema et al. 2006). Alternatively, the long-characteristics method computes the column densities to each zone by following rays emitted by all sources until they reach the domain boundary (Abel et al. 1999; Lim & Mellema 2003). This allows one to calculate the radiation field more accurately over the entire domain at the expense of excessively resolving the regions close to the sources when the number of ray directions is large.

To alleviate the inefficiency of the traditional long-characteristics method, Abel & Wandelt (2002) developed a novel, adaptive ray-tracing (ART) technique that has since been widely used to describe ionizing radiation originating from massive stars, in various astronomical contexts (e.g., Krumholz et al. 2007b; Bisbas et al. 2009; Wise & Abel 2011; Baczynski et al. 2015). In the ART method, rays are created at point sources and successively split as they are traced outward based on the Hierarchical Equal Area isoLatitude Pixelization (HEALPix) scheme (Górski et al. 2005). The salient feature of this method is that via ray splitting the angular resolution of radiation adapts to the local hydrodynamical resolution, such that the number of rays from each point source intersecting each grid zone remains approximately constant. While this RT method is quite efficient in achieving an accurate solution for the radiation field, existing codes employing the ART method often suffer from poor parallel performance. This is due to excessive overhead for interprocessor communication as rays traverse the interfaces between subdomains (see Section 2.2), making some implementations of the ART method essentially

inapplicable to large-scale simulations with multiple point sources.

Very recently, Rosen et al. (2017) developed a hybrid RHD module for the *Orion* code by combining the ART method with the flux-limited diffusion method. It makes use of the former to describe direct radiation from point sources, while taking advantage of the latter to treat diffuse radiation arising from gas and dust. Their parallel algorithm for the ART uses completely nonblocking, asynchronous communication, which greatly improves the parallel scaling of ray tracing over previous implementations using a synchronous communication algorithm. The scaling tests in Rosen et al. (2017) showed that the cost of the ART module in *Orion* remains comparable to that of hydrodynamics up to $\sim 10^3$ processors.

In this work, we describe the implementation of an ART module for multiple point sources in the grid-based magneto-hydrodynamics (MHD) code *Athena* (Stone et al. 2008) and present test results. Our implementation closely follows the parallelization strategy proposed by Rosen et al. (2017) and includes a few new features that further improve parallel performance. To model hydrogen ionization and recombination processes, we adopt a simple and efficient explicit scheme based on an analytic approximation. We first measure the scalability of our implementation by performing weak and strong scaling tests. We then apply the code to the standard test problems, namely, the expansion of R-type and D-type ionization fronts, as well as expansion of a dusty H II region driven by both thermal and radiation pressures. By comparing the numerical results with analytic or semianalytic solutions, we demonstrate the accuracy of our ART implementation.

A key application for ART is to follow the detailed effects of radiation produced by OB stars on their natal GMCs. To this end, we have combined our ART module with other physics packages implemented in *Athena* and run numerical RHD simulations of star cluster formation in turbulent, self-gravitating, unmagnetized clouds; these models are similar to the RHD simulations of Skinner & Ostriker (2015) and Raskutti et al. (2016) but include both ionizing and nonionizing radiation computed using ART. In this paper, we demonstrate the practical application of our ART module via an example simulation for a fiducial model. For comparison to the radiation field computed using the M_1 closure relation in the moment-based *Hyperion* code (Skinner & Ostriker 2013), we have also run an RT model with nonionizing radiation only, using an identical set of sources and density distribution. We find that the large-scale radiation fields from the two RT methods are quite similar to each other, but there is a non-negligible difference at small scales near the point sources, where resolution is inherently limited for pure moment methods like that in *Hyperion*.

The rest of this paper is organized as follows. In Section 2, we present the basic equations that we solve and briefly review the algorithm of ART together with our parallelization strategy. We also describe the subcycling method to solve the equation for hydrogen ionization and recombination. Section 3 presents the results of the scaling tests and the tests of the expansion of H II regions. In Section 4, we describe the numerical setup for simulations of star cluster formation in turbulent molecular clouds and compare the radiation fields based on the ART method with those from the moment method with the M_1 closure relation. Finally, we summarize and discuss future applications of the code in Section 5.

2. Numerical Method

Stellar UV radiation can alter the chemical, thermal, and dynamical state of the ISM through various processes. In this paper, we consider the two most basic processes: photoionization of hydrogen atoms and direct radiation pressure applied to the gas/dust mixture. The governing equations we adopt read as follows:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \quad (1)$$

$$\frac{\partial}{\partial t}(\rho \mathbf{v}) + \nabla \cdot (\rho \mathbf{v} \mathbf{v} + P \mathbf{I}) = -\rho \nabla \Phi + \mathbf{f}_{\text{rad}}, \quad (2)$$

$$\frac{\partial n_{\text{H}^0}}{\partial t} + \nabla \cdot (n_{\text{H}^0} \mathbf{v}) = \mathcal{R} - \mathcal{I}, \quad (3)$$

$$\nabla^2 \Phi = 4\pi G \rho, \quad (4)$$

where $\rho = 1.4m_{\text{H}}n_{\text{H}}$ is the gas density, with n_{H} representing the total hydrogen number density, \mathbf{v} is the velocity, P is the thermal pressure, Φ is the gravitational potential, and \mathbf{f}_{rad} is the force per unit volume due to absorption of stellar UV radiation. In Equation (3), n_{H^0} refers to the number density of neutral hydrogen, and \mathcal{I} and \mathcal{R} denote the volumetric ionization and recombination rates of hydrogen, respectively, whose functional forms are given in Section 2.3. Numerical solution of Equations (1) and (2) uses the methods described by Stone et al. (2008), with the radiation force source term updated using an operator splitting. The Poisson Equation (4) is solved using fast Fourier transforms (FFTs) with vacuum boundary conditions (Skinner & Ostriker 2015). At any location on the grid, the radiation flux and energy density are required to compute the source terms \mathbf{f}_{rad} and \mathcal{I} , respectively; these radiation terms are computed via our implementation of ART.

In what follows, we describe how the radiation fields are computed in the ART method in Section 2.1, while our method of parallelization is presented in Section 2.2. Our algorithms for updating radiation source terms and advancing Equation (3) are described in Section 2.3.

2.1. Adaptive Ray Tracing

Here we briefly review the RT problem for point sources in the framework of ART. We refer the reader to Wise & Abel (2011), Baczynski et al. (2015), and Rosen et al. (2017) for more detailed descriptions.

We consider only the direct radiation field from multiple point sources and ignore diffuse emission and scattered light. Under the assumption that the light-crossing time is much shorter than both the sound-crossing time and the timescale for the change in opacity, streaming radiation reaches equilibrium with matter effectively instantaneously. The radiation intensity should, therefore, satisfy the time-independent transfer equation

$$\mathbf{n} \cdot \nabla I = -\chi I, \quad (5)$$

where I is the radiation intensity, $\chi = \rho\kappa$ is the (isotropic) extinction coefficient of matter per unit path length for opacity κ , and \mathbf{n} is a unit vector parallel to the propagation direction of radiation. Both I and χ depend on the position \mathbf{x} , time t , and frequency ν , which are suppressed for notational simplicity.

We solve Equation (5) along a set of rays that discretize the directions in solid angle with respect to individual radiation sources. For simplicity, we consider monochromatic radiation emergent from a single point source located at \mathbf{x}_{src} with luminosity L ; an extension to polychromatic radiation from multiple sources is straightforward by taking a summation over a discretized set of frequencies and over rays from different sources. We inject photon packets from \mathbf{x}_{src} and carry them radially outward along the direction of any given ray by calculating the absorption rates of the photon energy and momentum over the grid cells they pass through. The direction of propagation \mathbf{n}_{ray} is determined using the HEALPix scheme of Górski et al. (2005), which divides the unit sphere into $N_{\text{ray}}(\ell) = 12 \times 4^\ell$ equal-area pixels at the level $\ell \geq 0$. Denoting the initial level by ℓ_0 , each injected photon packet on a given ray carries luminosity $L_{\text{ray}}(r=0) = L/N_{\text{ray}}(\ell_0)$, subtends a solid angle $\Omega_{\text{ray}}(\ell_0) = 4\pi/N_{\text{ray}}(\ell_0)$ from the source, and propagates along the ray $\mathbf{x} = \mathbf{x}_{\text{src}} + r\mathbf{n}_{\text{ray}}$, where $r = |\mathbf{x} - \mathbf{x}_{\text{src}}|$ measures the distance from the source.

Integrating Equation (5) over the whole solid angle gives the radiation energy equation

$$\nabla \cdot \mathbf{F} = -c\chi\mathcal{E}, \quad (6)$$

where c is the speed of light, \mathcal{E} is the radiation energy density, and \mathbf{F} is the radiation flux. With $\mathbf{F} = \hat{\mathbf{r}}Le^{-\tau(r,\mathbf{n})}/(4\pi r^2)$ and $\mathcal{E} = |\mathbf{F}|/c$ for streaming radiation, Equation (6) can be written as

$$\frac{\partial L_{\text{ray}}}{\partial r} = -\chi L_{\text{ray}}, \quad (7)$$

where $L_{\text{ray}}(r) = Le^{-\tau(r,\mathbf{n}_{\text{ray}})}/(4\pi)$ is the discretized luminosity at distance r in the region subtended by Ω_{ray} . We compute the length of a line segment Δr between the two consecutive cell interfaces intersected by a ray (see, e.g., Wise & Abel 2011; Baczynski et al. 2015). The corresponding cell optical depth is $\Delta\tau = \chi\Delta r$. The absorption rates of the radiation energy and momentum by the material along the path Δr in a given cell are then $\Delta L_{\text{ray}} = L_{\text{ray,in}}(1 - e^{-\Delta\tau})$ and $\mathbf{n}_{\text{ray}}\Delta L_{\text{ray}}/c$, respectively, where $L_{\text{ray,in}}$ is the luminosity of the ray entering the cell. The luminosity of a photon packet on a given ray is therefore reduced by ΔL_{ray} as it traverses the cell.

Utilizing the lab-frame equations of RHD (e.g., Mihalas & Auer 2001), these quantities are related to the volume-averaged radiation energy density and flux in a given cell as

$$\mathcal{E} = \frac{1}{\chi\Delta V} \sum_{\text{rays}} \frac{\Delta L_{\text{ray}}}{c}, \quad (8)$$

$$\mathbf{F} = \frac{1}{\chi\Delta V} \sum_{\text{rays}} \Delta L_{\text{ray}} \mathbf{n}_{\text{ray}}, \quad (9)$$

where ΔV is the cell volume and the summation is taken over all rays passing through the cell. Note that these expressions satisfy the flux-limiting condition $|\mathbf{F}| \leq c\mathcal{E}$ (Levermore 1984) and tend to the exact value in the limit of infinite angular resolution. In updating Equation (2) over time step Δt , an amount $\Delta t \mathbf{f}_{\text{rad}} = \Delta t \chi \mathbf{F}/c$ is added to the gas momentum in each cell using Equation (9) for \mathbf{F} (see Equation (22) below; in practice, these operator-split increments to the gas momentum are applied after each radiation subcycle). Since the photon momentum on a set of rays traversing a cell is reduced by

$\sum_{\text{rays}} \Delta L_{\text{ray}} \mathbf{n}_{\text{ray}} / c$ per unit time, this update is manifestly conservative of momentum.

To ensure that each grid cell is sampled by at least m_{ray} rays, we split a parent ray into four child rays at one higher HEALPix level if the solid angle $(\Delta x)^2 / r^2$ (corresponding to the maximum angle subtended by a given face of the current cell as seen from the source) is smaller than $m_{\text{ray}} \Omega_{\text{ray}}$. With a uniform grid spacing of Δx , this corresponds to the maximum distance $r_{\text{max}}(\ell) = [3 / (\pi m_{\text{ray}})]^{1/2} 2^\ell \Delta x$ that rays can travel at level ℓ . The new child rays are cast at positions $\mathbf{x}_{\text{src}} + r_{\text{max}}(\ell) \mathbf{n}_{\text{ray}}$. We follow photon packets as they traverse rays, dividing them equally when a parent ray splits into children. A ray is terminated and its photon packet is destroyed either where a ray exits the computational domain or where the photon packet is completely absorbed. The latter occurs when the total optical depth $\tau(r, \mathbf{n}_{\text{ray}})$ from the source is larger than a specified value τ_{max} , which we set to 7 as a default value. In order to alleviate the geometrical artifacts arising from the use of the crossing length rather than the ray–cell volume intersection, we randomly rotate the injection directions of the rays at every time step (Krumholz et al. 2007b).

2.2. Parallelization

Although conceptually simple and easy to implement, the major drawback of the ART method so far has been its poor parallel performance. There are two chief impediments to achieving scalable performance on a distributed memory platform: load imbalance and communication overhead. First, the amount of work needed per processor for the ART is roughly proportional to the number of the ray–cell crossings. Since this depends not only on the spatial distribution of sources but also on the opacity of the material, processors most likely have unequal workloads depending on the problem geometry. This imposes an inherent limit to scalability of ART with static domain decomposition. One may alleviate this issue by dynamically adjusting the local domain sizes to evenly distribute the workload.

The second problem is more serious and common to long-characteristic methods implemented with domain decomposition, in which the data from different subdomains are local to individual processors. For rays traveling across multiple subdomains, processors need to share information such as the propagation directions of the rays, optical depth, etc., in order to integrate Equation (7) along the characteristics. Because a ray may be terminated by the $\tau > \tau_{\text{max}}$ condition before traversing a subdomain that lies along the \mathbf{n}_{ray} direction, it cannot be determined in advance when, from where, or how many rays will enter a particular subdomain. The size and pattern of the data that need to be communicated among processors are therefore highly irregular. This poses a major challenge to the parallelization of the ART code. Indeed, the communication overhead has been the dominant performance bottleneck in the existing long-characteristic methods (e.g., Rijkhorst et al. 2006; Wise & Abel 2011).

Rosen et al. (2017) presented an efficient parallel algorithm for the ART module implemented in *Orion*, an adaptive mesh refinement (AMR) code. The key idea in their algorithm is to make use of the nonblocking message-passing operations (MPI_Isend, MPI_Iprobe, and MPI_Test) provided in the MPI-3 standard⁴ for communication of ray information

between processors. This allows processors to carry on their work (local transport of photon packets along rays) while communication is pending (see also Baczynski et al. 2015). Thus, the transport of rays within and between subdomains can be executed asynchronously, greatly reducing the idle time spent waiting for other processors to finish their work.

Another hallmark of the Rosen et al. (2017) algorithm is the use of a global “destroy” counter N_{dest} , which keeps track of how many rays have been terminated in the whole computational domain. Whenever a ray at the HEALPix level ℓ is terminated, we increment the local destroy counter by $N_{\text{dest,ray}} = 4^{\ell_{\text{max}} - \ell}$, where ℓ_{max} is the maximum level allowed. The global destroy counter is then a simple sum of the local destroy counters, and it stays synchronized across the processors by nonblocking communications. When all rays are terminated, we have $N_{\text{dest}} = N_{\text{dest,max}} = N_{\text{src}} \times 12 \times 4^{\ell_{\text{max}}}$, where N_{src} is the number of point sources in the entire domain. Updating the destroy counter as a global shared variable allows processors to determine when to exit the work-communication cycle without relying on synchronous, blocking communication.

We have implemented ART in *Athena* following the parallelization approach of Rosen et al. (2017), with some additional modifications to improve parallel performance. The schematic overview of the ray-tracing algorithm is shown in Figure 1. Our ART algorithm runs the following steps:

Initialization (executed only once at the start of the simulation):

Find neighbor grids⁵ that share faces, edges, or corners, and allocate memory for arrays `my_pp_list` and `exit_pp_list`, which will store local and outgoing ray (or photon packet) information, respectively. Each processor has one `my_pp_list` and n_{nbr} `exit_pp_list`, where n_{nbr} is the number of the neighbor grids.

Step 1: If any sources exist within my grid, create an initial set of directions \mathbf{n}_{ray} and corresponding L_{ray} values and store the photon packet information in `my_pp_list`. Compute $N_{\text{dest,max}}$ from N_{src} summed over all grids.

Step 2: Transport photon packets along their respective \mathbf{n}_{ray} directions within the local domain until either the total number of photon packets that need to be passed to neighbor grids exceeds $n_{\text{exit,max}}$ or no local ray is left in `my_pp_list`. Each ray is followed until it (1) needs to split, (2) reaches the grid boundaries, or (3) is terminated with $\tau \geq \tau_{\text{max}}$. Stack information on photon packets for rays leaving the local grid in the particular `exit_pp_list` for each neighbor. Calculate radiation energy density and flux for cells through which rays pass, as described in Section 2.1.

Step 3: Loop over the list of the neighbors and send photon packet information in `exit_pp_list` to the neighbor grids (MPI_Isend), checking whether the previous operation has completed (MPI_Testsome).

Step 4: Check for incoming messages from neighbors (MPI_Iprobe). If there are any, execute a blocking receive for each of them (MPI_Recv) and copy the received data to `my_pp_list`.

⁵ In *Athena* without mesh refinement, the computational domain is divided into a set of rectangular “grids,” each of which is owned by a single processor. Thus, “grid” is equivalent to “subdomain” for *Athena*.

⁴ <http://mpi-forum.org/docs/mpi-3.1>

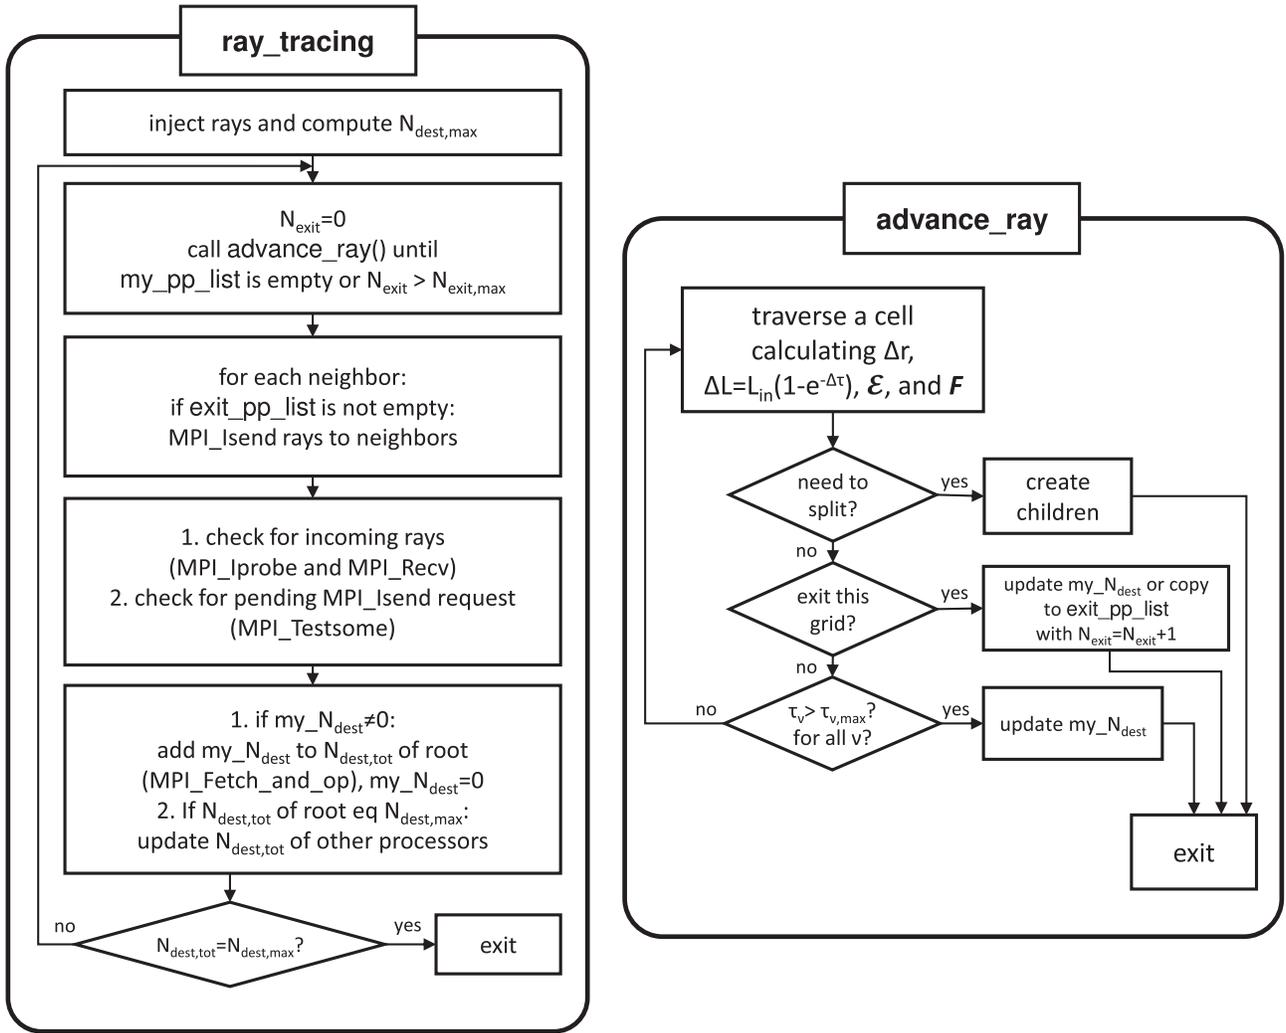


Figure 1. Flowchart of the ART algorithm: `ray_tracing` performs one ray tracing throughout the global computational domain, while `advance_ray` traces one ray within a local subdomain (“grid” in *Athena*).

Step 5: If `my_pp_list` is empty, add N_{dest} to the global destroy counter $N_{\text{dest,tot}}$ in the root processor (`MPI_Fetch_and_op`) and reset N_{dest} to 0.

Step 5-1 (root processor only): If $N_{\text{dest,tot}} = N_{\text{dest,max}}$, update $N_{\text{dest,tot}}$ to $N_{\text{dest,max}}$ in the other processors (`MPI_Fetch_and_op`).

Step 6: Go back to step 2 if $N_{\text{dest,tot}} \neq N_{\text{dest,max}}$. Exit the ray tracing when $N_{\text{dest,tot}} = N_{\text{dest,max}}$.

Our implementation of the ART method has two notable differences compared to that of Rosen et al. (2017). First, in Rosen et al. (2017) photon packets in a given grid are communicated to neighbors only after all rays have been traced to the grid boundaries. Therefore, processors without a source have to wait, repeatedly checking for rays coming from neighbors, until the intervening grids toward the sources successively finish their work. In our implementation, instead of waiting until all local ray tracing is complete, communications between neighbors are initiated as soon as the number of rays traced to the subdomain boundary exceeds a certain prescribed number $n_{\text{exit,max}}$, reducing the idle time spent by downstream processors (see also Baczynski et al. 2015). Although it is ideal to communicate a single photon packet right after the ray tracing reaches the subdomain boundary to

keep the downstream processors busy whenever possible, there is an optimal granularity, i.e., ratio of computation to communication, or an optimal value of $n_{\text{exit,max}}$ that gives the best performance owing to the finite overhead required for function calls and synchronization. The choice of $n_{\text{exit,max}}$ certainly depends on the source distribution, the domain decomposition, and machine specifications. In the case of a single point source at the box center, we find that $n_{\text{exit,max}} = 100$ gives the best performance, and we adopt this value for the rest of this paper.

Second, unlike in Rosen et al. (2017), we perform atomic (that is, uninterruptible) memory updates⁶ of the destroy counter utilizing the `MPI_Fetch_and_op` function, which is one of the remote memory access (RMA) operations in the MPI library. In the RMA operations, also called one-sided communications, processors are allowed to access a remote target memory without involving the explicit intervention of the remote processor. The RMA operations have comparatively low overheads since one processor specifies all communication parameters. We find that the use of the `MPI_Fetch_and_op` function helps to reduce

⁶ Atomic operations are guaranteed to complete without interference from other operations, ensuring correct results even when multiple processors try to access $N_{\text{dest,tot}}$ simultaneously.

substantially the overall cost of ray tracing when a large number of processors ($N_{\text{core}} \geq 256$) are used.

In addition to these differences, we find that the following implementation details can further improve parallel efficiency:

1. We use arrays for storing a collection of local and outgoing/incoming rays, to enable copying data to/from the communication buffer in one chunk. We dynamically allocate memory for arrays, which doubles in size if the old array becomes full.
2. As pointed out by Baczynski et al. (2015), we store only necessary information required for ray tracing, splitting, etc., in the ray data structure to minimize the amount of data transferred between processors. These include the HEALPix level ℓ , HEALPix pixel number, luminosity L_{ray} , source position \mathbf{x}_{src} , ray direction \mathbf{n}_{ray} , distance from the source, and the (integer) indices of the current cell in a grid.
3. We treat `my_pp_list` as a stack and access its elements in a last-in-first-out manner. For example, if a parent ray splits, child rays are stacked immediately on top of one another or on a sibling of the parent. This implies that we recursively follow a ray and its child until the grid boundary is reached (see also Baczynski et al. 2015).
4. In step 1, photon packets are stacked in such a way that every 12 contiguous blocks of elements in `my_pp_list` have ray directions that belong to the distinct 12 pixels at the base HEALPix level. The pixel numbers (in the nested numbering scheme) are arranged in 12 hierarchical tree structures, corresponding to the 12 base-level pixels. With $\ell_0 = 4$, for example, the HEALPix numbers of injected rays can be ordered as $0, 4^4, 2 \times 4^4, \dots, 11 \times 4^4, 1, 1 + 4^4, 1 + 2 \times 4^4, \dots$.

All of the above help to distribute workload to downstream processors as fast as possible and at a similar rate in all directions.

2.3. Update of Radiation Source Terms

We make a number of simplifying assumptions in modeling the photochemistry of hydrogen and its thermodynamic state. First, we adopt the on-the-spot approximation in which every diffuse Lyman-continuum photon resulting from a recombination to the ground state is locally reabsorbed. We neglect the diffuse ionizing radiation that can be important in the evolution of radiatively driven collapse and instabilities of ionization fronts (Haworth & Harries 2012). Second, we do not consider collisional ionizations, which are negligible compared to photoionizations in the temperature range of our interest ($T \lesssim 10^4$ K). Third, we do not solve an energy equation that accounts for radiative heating and cooling. Instead, we simply assign the gas temperature according to

$$T = T_{\text{ion}} - \left(\frac{x_{\text{n}}}{2 - x_{\text{n}}} \right) (T_{\text{ion}} - T_{\text{neu}}), \quad (10)$$

where $x_{\text{n}} \equiv n_{\text{H}^0}/n_{\text{H}}$ is the neutral gas fraction and T_{ion} and T_{neu} are the prescribed temperatures of the fully ionized ($x_{\text{n}} = 0$) and purely neutral ($x_{\text{n}} = 1$) states, respectively (e.g., Henney et al. 2005; Kim et al. 2016).⁷ For our tests presented in this

⁷ In reality, the gas temperature profile exhibits a peak immediately behind an ionization front due to spectral hardening. However, Lefloch & Lazareff (1994) found that the detailed functional form of $T(x_{\text{n}})$ does not significantly affect the dynamics of H II regions.

work, we adopt the constant temperatures $T_{\text{ion}} = 8000$ K and $T_{\text{neu}} = 20$ K, the latter of which falls within the temperature range inferred from the line ratios of low- J rotational transitions of CO molecules (e.g., Yoda et al. 2010). The gas thermal pressure is then set by $P = [1.1 + (1 - x_{\text{n}})]n_{\text{H}}k_{\text{B}}T$, accounting for 10% of the helium content. The two-temperature isothermal equation of state has been adopted by numerous numerical studies of H II regions and is valid as long as the timescale for approaching thermal equilibrium is short compared to the dynamical timescale (e.g., Lefloch & Lazareff 1994; Williams 2002; Gritschneider et al. 2009; Mackey et al. 2014; Steggle et al. 2017). Although idealized, Equation (10) is a simple and practical approach to following the pressure-driven dynamical expansion and internal structure of H II regions.

Equation (3) describes temporal changes of the neutral hydrogen fraction, with the recombination and ionization rates given by

$$\mathcal{R} = \alpha_{\text{B}}n_{\text{e}}n_{\text{H}^+} \quad (11)$$

and

$$\mathcal{I} = n_{\text{H}^0} \int_{\nu_{\text{L}}}^{\infty} \frac{c\mathcal{E}_{\nu}}{h\nu} \sigma_{\text{H}}(\nu) d\nu, \quad (12)$$

respectively, where \mathcal{E}_{ν} is the radiation energy density per unit frequency, $\alpha_{\text{B}} = 3.03 \times 10^{-13} (T/8000 \text{ K})^{-0.7} \text{ cm}^3 \text{ s}^{-1}$ is the case B recombination coefficient (Osterbrock 1989; Krumholz et al. 2007b), $n_{\text{H}^+} = n_{\text{e}} = n_{\text{H}}(1 - x_{\text{n}})$ is the number density of protons and electrons, ν_{L} is the Lyman limit corresponding to $h\nu_{\text{L}} = 13.6$ eV, and $\sigma_{\text{H}}(\nu)$ is the photoionization cross section. In practice, Equation (12) is evaluated as a discrete summation over a finite number of frequency bins. In this work, we use one frequency bin for ionizing radiation, denoted by the subscript ‘‘i.’’ Equation (12) then becomes

$$\mathcal{I} = n_{\text{H}^0} \frac{c\mathcal{E}_{\text{i}}}{h\nu_{\text{i}}} \langle \sigma_{\text{H}} \rangle \quad (13)$$

$$= \frac{1}{\Delta V} \frac{n_{\text{H}^0} \langle \sigma_{\text{H}} \rangle}{\chi_{\text{i}}} \sum_{\text{rays}} \frac{\Delta L_{\text{ray},\text{i}}}{h\nu_{\text{i}}}, \quad (14)$$

where $\mathcal{E}_{\text{i}} = \int_{\nu_{\text{L}}}^{\infty} \mathcal{E}_{\nu} d\nu$ is the energy density of ionizing radiation, as evaluated using Equation (8), $h\nu_{\text{i}}$ is the mean energy of the ionizing photons, and

$$\langle \sigma_{\text{H}} \rangle = \int_{\nu_{\text{L}}}^{\infty} \frac{\mathcal{E}_{\nu}/h\nu}{\mathcal{E}_{\text{i}}/h\nu_{\text{i}}} \sigma_{\text{H}}(\nu) d\nu \quad (15)$$

is the frequency-averaged effective cross section. The actual value of $\langle \sigma_{\text{H}} \rangle$ depends on the spectral shape of the incident radiation \mathcal{E}_{ν} , which in turn depends on both the optical depth from the source and the spectral types of ionizing stars. For simplicity, we take the constant values $h\nu_{\text{i}} = 18$ eV and $\langle \sigma_{\text{H}} \rangle = 6.3 \times 10^{-18} \text{ cm}^2$ in the present work.

We solve Equation (3) in two steps: hydrodynamic update and source update by treating \mathcal{R} and \mathcal{I} as source terms. The source update requires solving

$$\frac{dx_{\text{n}}}{dt} = \frac{\mathcal{R} - \mathcal{I}}{n_{\text{H}}} = \alpha_{\text{B}}n_{\text{e}}(1 - x_{\text{n}}) - x_{\text{n}}\Gamma, \quad (16)$$

where $\Gamma \equiv \mathcal{I}/n_{\text{H}}^0$. Since the timescale for changes of x_{n} is almost always smaller than the hydrodynamic time step, we update x_{n} explicitly using subcycling, as explained below.

Assuming that α_{B} and Γ are constant during a substep and using $n_{\text{e}} = n_{\text{H}}(1 - x_{\text{n}})$, the right-hand side of Equation (16) is quadratic in x_{n} . Altay & Theuns (2013) showed that this has an analytic solution⁸

$$x_{\text{n}}(t) = x_{\text{eq}} + \frac{(x_{+} - x_{\text{eq}})(x_0 - x_{\text{eq}})K}{(x_{+} - x_0) + (x_0 - x_{\text{eq}})K}, \quad (17)$$

where $x_0 = x_{\text{n}}(t_0)$, $K = \exp[-(x_{+} - x_{\text{eq}})(t - t_0)\alpha_{\text{B}}n_{\text{H}}]$, $x_{+} = x_{\text{eq}}^{-1}$, and

$$x_{\text{eq}} = \frac{2\alpha_{\text{B}}n_{\text{H}}}{\Gamma + 2\alpha_{\text{B}}n_{\text{H}} + \sqrt{\Gamma^2 + 4\alpha_{\text{B}}n_{\text{H}}\Gamma}} \quad (18)$$

is the equilibrium neutral fraction. When $\Gamma = 0$, Equation (16) has a solution

$$x_{\text{n}}(t) = \frac{x_0 + (1 - x_0)\alpha_{\text{B}}n_{\text{H}}t}{1 + (1 - x_0)\alpha_{\text{B}}n_{\text{H}}t}, \quad (19)$$

which should be applied to cells completely shielded from ionizing radiation. While Equation (17) is exact, it is not computationally robust when the denominator is close to zero, possibly resulting in inaccurate x_{n} due to amplified roundoff errors.

Alternatively, if we (incorrectly) treat n_{e} as being constant, Equation (16) yields a solution

$$x_{\text{n}}(t) = x_{\text{eq},0} + (x_0 - x_{\text{eq},0})e^{-(t-t_0)/t_{1-r}}, \quad (20)$$

where $t_{1-r} = (\Gamma + \alpha_{\text{B}}n_{\text{e}})^{-1}$ is the ionization/recombination time and $x_{\text{eq},0} = \alpha_{\text{B}}n_{\text{e}}/(\Gamma + \alpha_{\text{B}}n_{\text{e}})$ is the equilibrium neutral fraction (e.g., Schmidt-Voigt & Koeppen 1987; Mellema et al. 2006). Mellema et al. (2006) adopted Equation (20) to implicitly update the time-averaged ionization fraction in their C^2 -ray method. In the Appendix, we present the test results of Equation (20) on the temporal changes of x_{n} in a single cell exposed to a fixed radiation field. It turns out that although Equation (20) is based on the incorrect assumption of constant n_{e} , in practice it gives almost identical results to those with Equation (16). In addition, Equation (20) is robust and guarantees that x_{n} always lies between x_0 and $x_{\text{eq},0}$. We also find (see Appendix) that it is more accurate than methods based on the backward-difference formula often used in the literature (e.g., Anninos et al. 1997). We thus use Equation (20) to calculate x_{n} in our implementation of the subcycle.

We determine the size of substeps as

$$\Delta t_{\text{ss}} = C \times \min\left(\frac{n_{\text{H}}}{|\mathcal{I} - \mathcal{R}|}\right), \quad (21)$$

with a constant coefficient C . Taking $C = 0.1$ restricts the change of x_{n} per substep to below 0.1 (see also Baczynski et al. 2015). The minimum value of Δt_{ss} is usually from the cells in transition layers where $0.1 \lesssim x_{\text{n}} \lesssim 1$ and $x_{\text{eq}} \approx 0$.

Our overall computation procedure is as follows. We first evolve the hyperbolic terms in Equations (1)–(3) for a full

hydrodynamic time step Δt using the existing Godunov-type scheme in the *Athena* code. The total gas density n_{H} and the neutral gas density n_{H}^0 are then available as inputs to the ART module for radiation (and ionization/recombination) subcycles. Next, we perform the ART for each radiation subcycle to calculate \mathcal{E} and Γ and determine Δt_{ss} . We then update the neutral fraction by using Equation (20); given an updated x_{n} , the electron density is updated to $n_{\text{e}} = n_{\text{H}}(1 - x_{\text{n}})$. At each subcycle over Δt_{ss} , we explicitly calculate the radiation force

$$\mathbf{f}_{\text{rad}} = \frac{1}{c} \sum_{\nu_j} \chi_{\nu_j} \mathbf{F}_{\nu_j}, \quad (22)$$

using Equation (9) for each frequency bin, and add $\mathbf{f}_{\text{rad}} \Delta t_{\text{ss}}$ to the gas momentum density if radiation pressure is switched on. The total opacity in the ionizing and nonionizing bins is calculated as

$$\chi_{\text{i}} = n_{\text{H}}^0 \langle \sigma_{\text{H}} \rangle + n_{\text{H}} \sigma_{\text{d}} \quad (23)$$

and

$$\chi_{\text{n}} = n_{\text{H}} \sigma_{\text{d}}, \quad (24)$$

respectively, where we use the constant-attenuation cross section per hydrogen atom $\sigma_{\text{d}} = 1.17 \times 10^{-21} \text{ cm}^2 \text{ H}^{-1}$ (Draine 2011).

3. Results of Code Tests

We now present the results of various tests intended to verify the performance and accuracy of our implementation of the ART, including its ability to simulate the dynamics of H II regions. Unless otherwise noted, we adopt $m_{\text{ray}} = 4$ and $\ell_0 = 4$ as fiducial values for the angular resolution and the initial HEALPix level, respectively (Baczynski et al. 2015; Rosen et al. 2017).

3.1. Scaling Tests

To measure the parallel performance of the ART and its cost relative to that of the hydrodynamic solver, we conduct weak and strong scaling tests similar to the ones presented in Rosen et al. (2017). The tests are run on 16-core Intel Sandy Bridge nodes in the Tiger cluster at Princeton University.

In the weak scaling test, the whole computational domain is subdivided into N_{core} identical grids, each having 32^3 cells over a 1 pc^3 volume with a source at each grid center. Each grid is assigned to a processor, and we time the execution of the ART and the hydrodynamics solver, varying the number of processors N_{core} . Figure 2(a) shows the wall-clock time (t_{wall}), averaged over 10 cycles, taken for a single ART (circles) and hydrodynamics (triangles) update. We consider problem sizes between $N_{\text{core}} = 2^3$ and 2^{10} . The open circles correspond to the case in which rays are terminated at $d_{\text{max}} = 0.6 \text{ pc}$, such that communication is only with immediately adjacent neighbors. Overall, the wall-clock time is nearly flat with varying N_{core} , indicating excellent parallel efficiency (a horizontal line would represent perfect weak scaling for this case). Quantitatively, the run time to perform the ray tracing is 0.050 s when $N_{\text{core}} = 8$, slightly shorter than the hydrodynamics update, and increases to 0.086 s when $N_{\text{core}} = 2^{10}$, with parallel efficiency of 58%. For the weak scaling test, we also consider the case where rays are followed until they exit the simulation domain, which is plotted as filled circles. This is in

⁸ These expressions can be generalized to include the collisional ionizations and the contribution to n_{e} from heavy elements. See Appendix C3 in Altay & Theuns (2013).

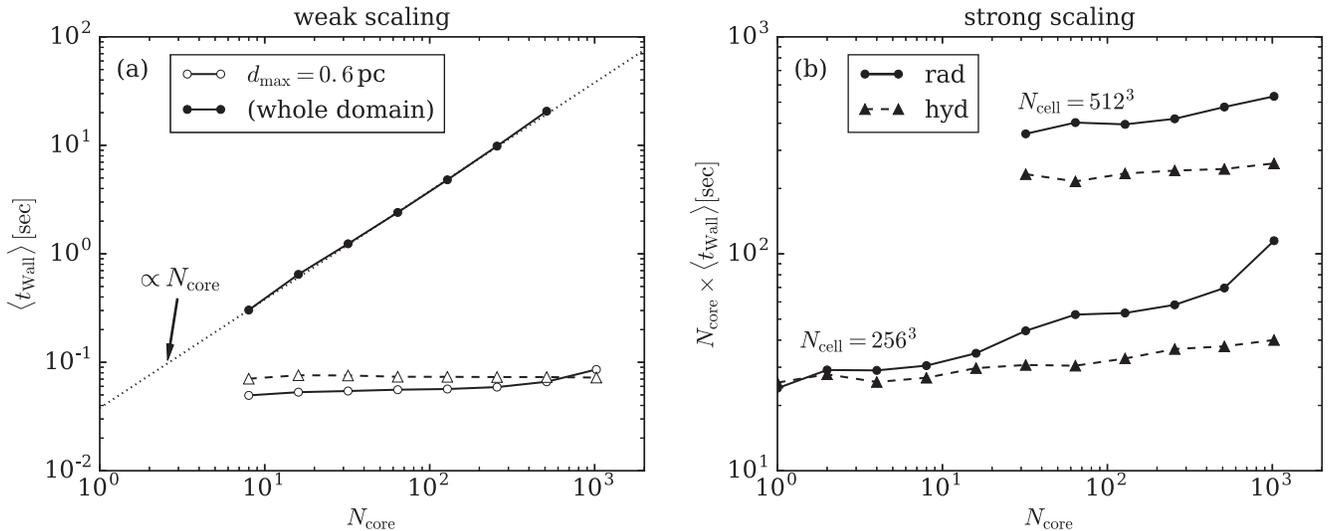


Figure 2. (a) Wall-clock time vs. N_{core} in the weak scaling test for the ART (circles) and hydrodynamics solver (triangles) in which every 32^3 grid belonging to a given processor has a single point source at the center. The open circles correspond to the case in which each ray is terminated 0.6 pc from its originating source, while filled circles denote the case in which all rays are extended until they hit the simulation domain boundary. (b) Wall-clock time multiplied by N_{core} vs. N_{core} in the strong scaling test. A single point source is placed at the center of the box with 256 or 512 cells per side.

good agreement with the prediction of the perfect scaling for a processor workload that increases with the number of sources and hence rays per cell $\propto N_{\text{source}} \propto N_{\text{core}}$, with $\langle t_{\text{Wall}} \rangle \propto N_{\text{core}}$ shown as a dotted line, indicating that the processors are busy most of the time.

In the strong scaling test, the total problem size remains fixed, with the domain decomposed into a varying number of grids. We consider a domain with 256^3 or 512^3 cells and add a single point source at the domain center. We do not restrict the distance that rays can travel from the source: rays extend to the domain boundaries. Figure 2(b) plots the wall-clock time to complete one ray tracing (solid) and hydrodynamics update (dashed) multiplied by N_{core} . A horizontal line would represent perfect strong scaling for this case. For the domain with 256^3 cells, the run time increases by a factor of 4.74 as N_{core} varies from 1 to 1024, corresponding to a scaling efficiency of 21%. For the domain with 512^3 cells, we start from $N_{\text{core}} = 32$, since the problem becomes memory bound, and obtain a relative parallel efficiency of 67% on 1024 cores. It is remarkable that the cost of ray tracing in our implementation is comparable to that of hydrodynamic updates even though rays pass through multiple subdomains to reach the domain boundaries. This indicates that our communication approach is successfully distributing work throughout the domain and, in particular, that “downstream” processors do not suffer from being idle.

Compared to the strong scaling presented by Rosen et al. (2017) (their Figure 8), our result for the 256^3 box shows a performance improvement by a factor of more than 10. Differences in machine specifications and hierarchical grid structures may contribute to these differences in scaling results. For instance, the *Orion* AMR code used by Rosen et al. (2017) employs a patch-based AMR method in which the domain is decomposed into a set of grids of uniform cell spacing, and multiple grids may be assigned to a single core. The same patch-based method is applied even without mesh refinement. Finding neighbors on the patch-based mesh is nontrivial and requires looping over local grids until the next grid is found. This is in contrast to the domain decomposition method adopted by *Athena*, in which a single core covers only a single

grid. Therefore, the patch-based scheme requires the additional cost of finding neighbor grids when a ray needs to be passed between grids even when those grids reside on a single core. This most likely accounts for some of the differences in the scaling results. Further optimizations described in Section 2.2 may also contribute to the performance improvement.

3.2. Radiation in a Vacuum

We now assess the accuracy of our implementation of the ART method for recovering the inverse-square law of radiation energy density around a point source in vacuum. With a limited angular resolution, the ART method is unable to achieve perfect spherical symmetry under Cartesian geometry. We study the effects of the angular resolution parameter m_{ray} and the ray rotation on the accuracy of the calculated radiation energy density.

As a computational domain, we consider a cubic box with a side length of 2 pc and place a point source with luminosity $L = 1L_{\odot}$ at the center. The box is discretized into 128^3 cells. Figure 3(a) shows the distributions of the volume-averaged radiation energy density \mathcal{E} for $(m_{\text{ray}}, \ell_0) = (4, 4)$ (blue), $(m_{\text{ray}}, \ell_0) = (10, 4)$ (red), and $(m_{\text{ray}}, \ell_0) = (10^3, 10)$ (yellow) against the normalized distance from the source. Each distribution is shifted by a constant factor along the ordinate for clear comparison. The deviations of \mathcal{E} relative to the cell-centered value $L/(4\pi r^2 c)$, plotted as the black dashed line, are mostly on the order of a few percent for $m_{\text{ray}} = 4$. We also run 10 different instances of the ray trace with $m_{\text{ray}} = 4$, randomly varying the ray orientation. The mean values of the resulting \mathcal{E} are plotted as green symbols, demonstrating that the errors introduced by geometrical artifacts can be reduced by rotating the directions of ray injection (Krumholz et al. 2007b).

Figure 3(b) plots as solid lines the 10th, 50th (median), and 90th percentiles of the relative errors $|\Delta\mathcal{E}|/\mathcal{E}$ within spherical shells centered at the source, after taking the case with $(m_{\text{ray}}, \ell_0) = (10^3, 10)$ as the reference solution. The median value of the relative error is 1%–4% for $(m_{\text{ray}}, \ell_0) = (4, 4)$. The case with ray rotation achieves a median accuracy similar to that with $m_{\text{ray}} = 10$. The sawtooth patterns in the relative

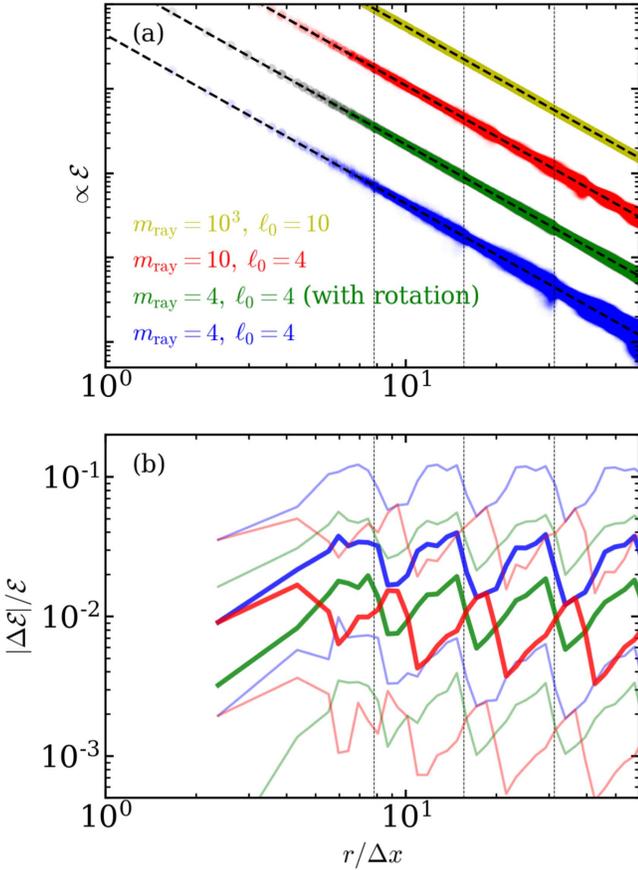


Figure 3. Test of the ART on the radiation field around a single point source in vacuum. (a) Radiation energy density vs. the distance from the source (normalized by the grid spacing Δx) with $(m_{\text{ray}}, \ell_0) = (4, 4)$ (blue and green), $(m_{\text{ray}}, \ell_0) = (10, 4)$ (red), and $(m_{\text{ray}}, \ell_0) = (10^3, 10)$ (yellow). The data for each case are shifted along the ordinate for clear comparison. The green dots show the averages of 10 ARTs, each with different ray orientation. The black dashed lines denote the analytic solution $\mathcal{E} \propto r^{-2}$. (b) 10th, 50th (median; thick line), and 90th percentile values of the relative error in each radial bin. The vertical dashed lines indicate the distance at which rays split for $(m_{\text{ray}}, \ell_0) = (4, 4)$.

errors reflect the radial variation of the angular resolution: the resolution becomes gradually worse with increasing r at a given HEALPix level and suddenly increases at the ray-splitting radii $r_{\text{max}}(\ell)$, plotted as the vertical dashed lines for $(m_{\text{ray}}, \ell_0) = (4, 4)$. Varying m_{ray} from 2 to 256, we find that the median value of the relative errors can be fitted as $2.0(m_{\text{ray}}/4)^{-1.32}\%$, which is steeper than $\propto m_{\text{ray}}^{-0.6}$ obtained by Wise & Abel (2011).

3.3. Expansion of H II Regions

Next, we perform three tests of the expansion of an H II region embedded in a uniform medium. These are classical problems in astrophysics with well-known analytic solutions and have thus been the standard tests for RHD codes incorporating the effects of ionization and recombination (e.g., Mellema et al. 2006; Krumholz et al. 2007b; Wise & Abel 2011; Baczynski et al. 2015; Bisbas et al. 2015). In this section, we adopt $T_{\text{neu}} = 100$ K and $T_{\text{ion}} = 8000$ K.

3.3.1. R-type Ionization Front

We consider an R-type ionization front created by a central source in a uniform static medium with hydrogen number

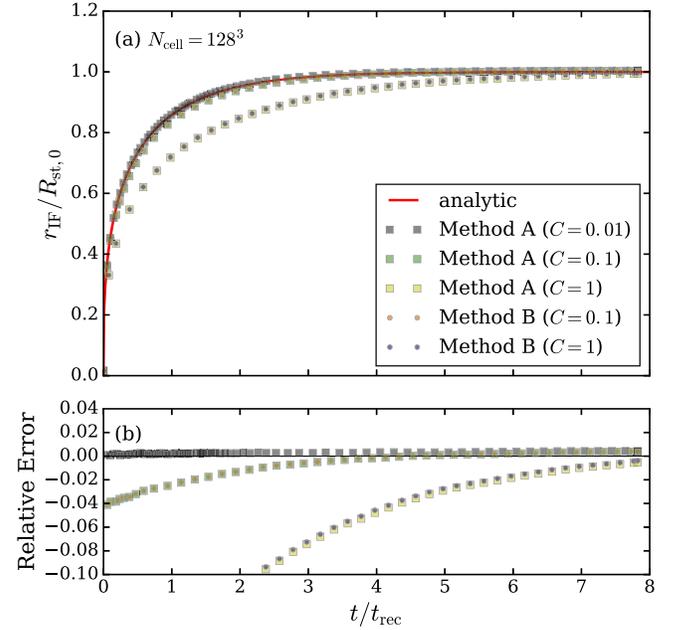


Figure 4. Test of the expansion of an R-type ionization front in a uniform medium. Methods A (squares) and B (circles) refer to using Equations (17) and (20), respectively. (a) Ionization front radius vs. time with different time-stepping coefficient C in Equation (21). The analytic solution is shown as the red solid line. (b) Fractional errors relative to the analytic solution.

density n_{H} . At $t = 0$, the source starts to emit ionizing photons at a constant rate of Q_i . The expansion of the ionization front is very rapid at early time, without inducing significant gas motions. If there is no dust extinction and the recombination coefficient is taken as constant, the ionization front expands as

$$r_{\text{IF}} = R_{\text{St},0}(1 - e^{-t/t_{\text{rec}}})^{1/3}, \quad (25)$$

where $t_{\text{rec}} = 1/(\alpha_{\text{B}}n_{\text{H}})$ is the recombination time and $R_{\text{St},0} = [3Q_i/(4\pi\alpha_{\text{B}}n_{\text{H}}^2)]^{1/3}$ is the dustless Strömgen radius (Spitzer 1978). For our test, we take $n_{\text{H}} = 10^2 \text{ cm}^{-3}$, $Q_i = 10^{49} \text{ s}^{-1}$, $\sigma_{\text{H}} = 6.3 \times 10^{-18} \text{ cm}^2$, and $\alpha_{\text{B}} = 3.02 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$. The simulation domain is a cubic box with sides of $2.4R_{\text{St},0} = 7.2 \text{ pc}$, which is resolved by a grid of $N_{\text{cell}} = 128^3$ cells. The optical depth over one neutral cell is $\Delta\tau = n_{\text{H}}\sigma_{\text{H}}\Delta x = 109(128/N_{\text{cell}}^{1/3})(Q_i/10^{49} \text{ s}^{-1})^{1/3}(n_{\text{H}}/10^2 \text{ cm}^{-3})^{1/3}$, so that the background medium is highly optically thick initially. The test simulations are run until $t = 8t_{\text{rec}}$, with the hydrodynamic updates turned off.

We first examine the effect of differing time step in the subcycling. Figure 4 plots the temporal changes of r_{IF} , defined as the radius at which $x_{\text{n}} = 0.5$, and the relative errors compared to Equation (25). The results based on Equation (17) (Method A) with $C = 0.01, 0.1$, and 1 in Equation (21) are shown as squares. For $C = 0.1$, the simulation results agree with the analytic solution within $\lesssim 5\%$, similar to the results of Baczynski et al. (2015). The results using Equation (20) (Method B) with $C = 0.1$ and 1 are shown as circles, which are almost identical to those from Equation (17) with the same C . We also explore the effect of varying spatial resolution by taking Method B with $C = 0.1$. Figure 5 plots the resulting r_{IF} for $N_{\text{cell}} = 64^3, 128^3$, and 256^3 , all of which agree within 2%. The relative errors are $\sim 4\%$ – 5% at early time and decrease to less than 1% at $t/t_{\text{rec}} = 8$. The above results suggest that Method B with $C = 0.1$ reproduces the evolution of R-type

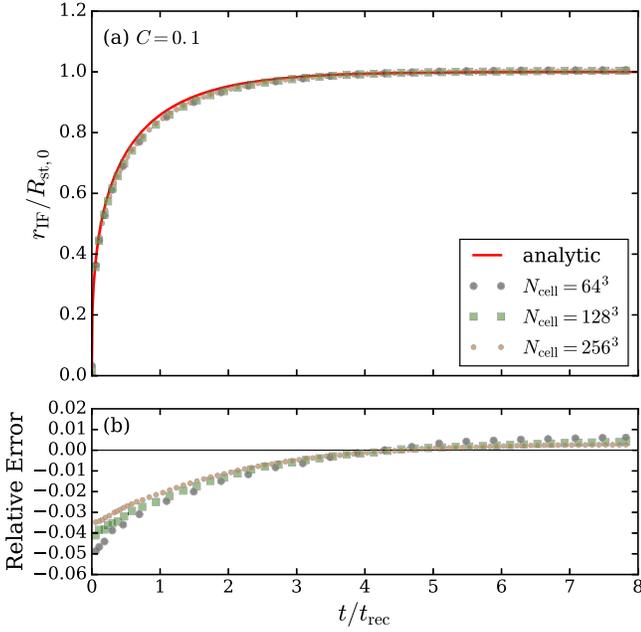


Figure 5. Test of the expansion of an R-type ionization front in a uniform medium. Same as Figure 4, using Method B with $C = 0.1$, but with a range of spatial resolution.

ionization fronts quite well. They can be followed more accurately, albeit at a higher cost, if one restricts Δt_{ss} more strictly, for example, by limiting the relative changes in x_n to less than 10% per update (e.g., Krumholz et al. 2007b; Mackey 2012).

3.3.2. D-type Ionization Front

As an R-type ionization front approaches $R_{St,0}$, its expansion speed falls below twice the sound speed c_{ion} in the ionized region. This develops an isothermal shock in front of the ionization front, which in turn undergoes a transition to a (weak) D-type front (Shu 1992). If the shell of the swept-up material between the ionization and shock fronts is geometrically thin, the radius of the shell r_{sh} ought to satisfy the following momentum equation:

$$\frac{d}{dt} \left(M_{sh} \frac{dr_{sh}}{dt} \right) = 4\pi r_{sh}^2 c_{ion}^2 \rho_{ion}, \quad (26)$$

where $M_{sh} = (4\pi/3)\rho_0 r_{sh}^3$ is the shell mass for ρ_0 , which is the density of the background medium, and $\rho_{ion} = \rho_0 (r_{sh}/R_{St,0})^{-3/2}$ is the density in the ionized region (assuming instantaneous ionization equilibrium). Hosokawa & Inutsuka (2006) found that Equation (26) has a self-similar solution:

$$r_{sh} = R_{St,0} \left(1 + \frac{7}{2\sqrt{3}} \frac{c_{ion} t}{R_{St,0}} \right)^{4/7}. \quad (27)$$

Spitzer (1978) also solved for a self-similar solution by employing the requirement of $\rho_0 (dr_{sh}/dt)^2 \approx c_{ion}^2 \rho_{ion}$ and derived a similar expression, with $2\sqrt{3}$ in the denominator of the coefficient in Equation (27) replaced by 4 (see also Bisbas et al. 2015).

For the test of D-type fronts, we set up a cubic domain with 128^3 cells whose side is $20R_{St,0}$ long. The domain is filled with

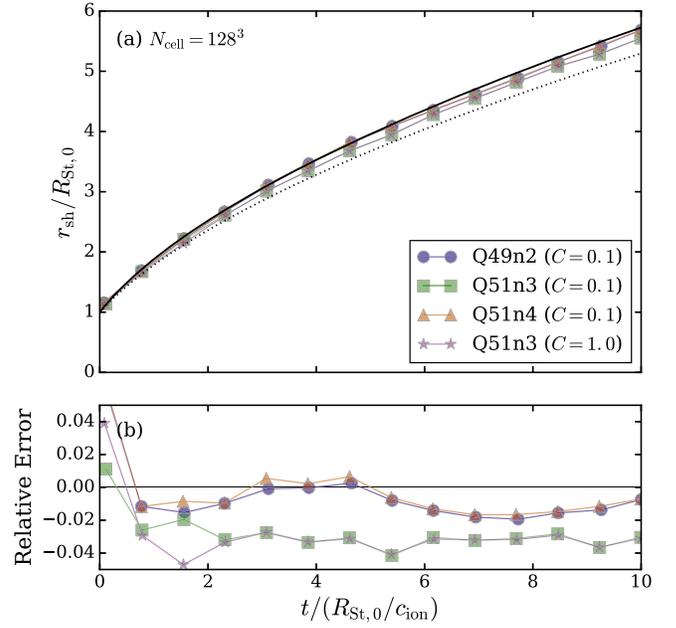


Figure 6. Test of the expansion of a D-type ionization front in a uniform medium. (a) Shell radius r_{sh} vs. time. The black solid and dotted lines represent the analytic solution of Hosokawa & Inutsuka (2006) and Spitzer (1978), respectively. (b) Errors relative to the analytic solution of Hosokawa & Inutsuka (2006).

neutral gas with $T_{neu} = 100$ K, and a source placed at the center starts to emit Q_i ionizing photons per unit time from $t = 0$. Unlike in the R-type front tests, we turn on the hydrodynamics updates so that the gas responds self-consistently to the pressure of gas at T_{ion} produced by the ionizing radiation. The simulations are run up to $t = 10R_{St,0}/c_{ion}$ using Method B with $C = 0.1$. At each time, we determine the shell radius r_{sh} as the position where the gas density is maximal. Figure 6 plots the resulting r_{sh} and the relative errors as functions of time for $(Q_i/s^{-1}, n_H/cm^{-3}) = (10^{49}, 10^2)$, $(10^{51}, 10^3)$, and $(10^{51}, 10^4)$. Shown also as black solid and dotted lines are the solutions of Hosokawa & Inutsuka (2006) and Spitzer (1978), respectively. Our results agree with the former better, with typical relative errors less than 3%, which is consistent with the results of Bisbas et al. (2015).⁹ Using a less stringent time step size with $C = 1.0$, the errors become slightly larger, but the computational cost of the radiation module is reduced by a factor of 2.5.

3.3.3. Dusty H II Region with Radiation Pressure

In the preceding tests, we have ignored the presence of dust grains and radiation pressure on them. They are efficient in absorbing (both ionizing and nonionizing) UV photons and transfer momentum to the gas through collisional coupling. The momentum deposition from radiation pressure may dominate in driving the expansion of H II regions in dense, massive star-forming environments (Krumholz & Matzner 2009; Fall et al. 2010; Murray et al. 2010; Kim et al. 2016). Radiation pressure also causes a nonuniform-density distribution inside an H II region in static equilibrium (Draine 2011). Time-dependent, spherical models developed by Kim et al. (2016) confirmed that

⁹ We note that the benchmark tests in Bisbas et al. (2015) considered hydrogen-only gas and adopted $T_{ion} = 10^4$ K and $\alpha_B = 2.7 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$, corresponding to a $\sim 30\%$ and $\sim 10\%$ difference in c_{ion} and α_B from ours, respectively. However, our result is unlikely to be affected by the specific choice of T_{ion} provided that $T_{ion} \gg T_{neu}$.

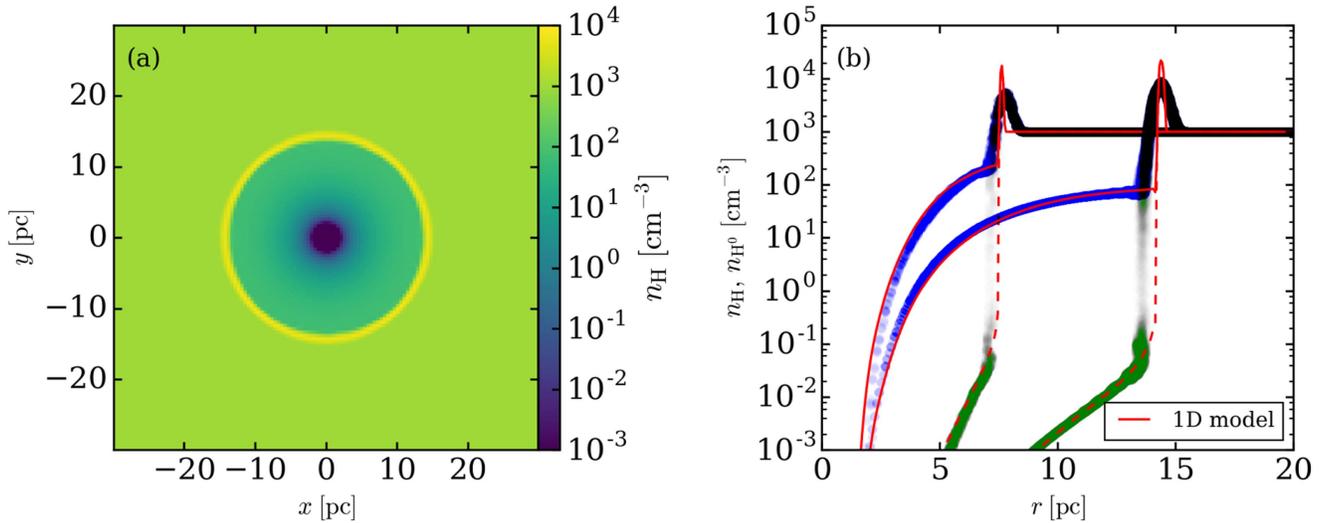


Figure 7. Test of the expansion of a dusty H II region with $Q_i = 10^{51} \text{ s}^{-1}$ and $n_{\text{H}} = 10^3 \text{ cm}^{-3}$. (a) Slice of n_{H} through the $z = 0$ plane at $t = 2.26$ Myr. (b) Radial profiles of n_{H} (blue) and n_{H}^0 (green) at $t = 0.67$ and $t = 2.26$ Myr. The red lines represent the results of a one-dimensional simulation in spherical coordinates.

when $Q_i n_{\text{H}}$ is large, the expansion of dusty H II regions is dominated by radiation pressure, with nonuniform internal structure.

Let L_n denote the photon luminosity below energy $h\nu_L = 13.6 \text{ eV}$ from a source, and let τ_{edge} and ρ_{edge} be the dust optical depth and gas density at the edge of the ionized region, respectively. The equation of motion of the shell is then modified to

$$\frac{d}{dt} \left(M_{\text{sh}} \frac{dr_{\text{sh}}}{dt} \right) = \frac{L_n}{c} e^{-\tau_{\text{edge}}} + 4\pi r_{\text{sh}}^2 c_{\text{ion}}^2 \rho_{\text{edge}} \quad (28)$$

(Kim et al. 2016). Given τ_{edge} and ρ_{edge} as a function of r_{sh} from Draine (2011), Equation (28) can readily be integrated to yield r_{sh} as a function of time, to which our test results are compared.

We consider an initially uniform medium with $n_{\text{H}} = 10^3 \text{ cm}^{-3}$ and a central source with $Q_i = 10^{51} \text{ s}^{-1}$ and $L_n = 1.5Q_i h\nu_i$, where $h\nu_i = 18 \text{ eV}$ is the mean energy of hydrogen-ionizing photons. We take a constant dust absorption cross section $\sigma_d = 10^{-21} \text{ cm}^2 \text{ H}^{-1}$ and ionized gas temperature $T_{\text{ion}} = 8000 \text{ K}$, corresponding to the dust opacity parameter of $\gamma \equiv (2ck_B T_{\text{ion}} \sigma_d) / (\alpha_B h\nu_i) = 7.58$ (Equation (7) of Draine 2011).¹⁰ We take a computational domain with side $20R_{\text{St},0} = 60 \text{ pc}$, resolved by 128^3 cells. For comparison, we also run a one-dimensional simulation in spherical coordinates with the same set of the physical parameters, but resolving the radial domain $0.1 \text{ pc} < r < 30.1 \text{ pc}$ using 512 cells (for details of the simulation setup, see Kim et al. 2016).

Figure 7 plots (a) the density distribution in the $z = 0$ plane at $t = 2.26$ Myr and (b) the radial profiles of n_{H} (blue) and n_{H}^0 (green) at $t = 0.67$ and 2.26 Myr. The corresponding one-dimensional results are compared as the red lines. As expected, radiation pressure on dust creates a central cavity devoid of gas and dust, with the density decaying toward the center approximately as $n_{\text{H}} \propto \exp(-r_0/r)$, where $r_0 = \sigma_d(Q_i h\nu_i + L_n) / (8\pi ck_B T_{\text{ion}})$ (Rodríguez-Ramírez et al. 2016). Although the shocked shell in the three-dimensional model has a larger thickness than the

one-dimensional counterpart owing to poor spatial resolution, the shell radius agrees quite well. Figure 8 compares $r_{\text{sh}}(t)$ in the simulations with that from the solution of Equation (28), again showing good agreement (within 5% for $t > 0.5$ Myr) between the numerical and semianalytic results. The difference between r_{IF} and r_{sh} in the three-dimensional model is primarily due to the limited spatial resolution; the true physical difference between r_{IF} and r_{sh} is smaller.

4. Application to Star Cluster Formation in Turbulent Clouds

So far our tests have been limited to idealized problems, in which the background is a uniform medium. By combining our implementation of ART with other physics modules already existing in the *Athena* code, we now demonstrate application of our ART module to an important practical astronomical problem, namely, the formation of a star cluster in a molecular cloud. We follow the collapse of a turbulent, self-gravitating cloud and subsequent star formation until the associated UV radiation feedback from multiple sources halts further star formation and disperses the remaining gas. In this section, we first describe the numerical setup and report strong scaling for the fiducial model. We then compare the nonionizing UV radiation field computed by the ART with that based on the M_1 closure scheme of Skinner & Ostriker (2013), as implemented for point sources in the single-scattering approximation.

4.1. Numerical Setup

We solve Equations (1)–(4) to study evolution of self-gravitating gas interacting with UV radiation from multiple point sources. We adopt *Athena*'s HLLC Riemann solver, the van Leer integrator (Stone & Gardiner 2009), piecewise linear spatial reconstruction, and strict outflow boundary conditions.

We use the sink particle method of Gong & Ostriker (2013) to treat cluster formation and gas accretion processes. For this, we create a sink particle when a cell meets the following three conditions simultaneously: (1) its density exceeds the threshold corresponding to the Larson–Penston solution of an isothermal collapse, (2) it is a local minimum of the gravitational potential, and (3) it has a converging velocity field along all three

¹⁰ For the numerical value of γ , we took $\alpha_B \simeq 2.59 \times 10^{-13} (T/10^4 \text{ K})^{-0.7} \text{ cm}^3 \text{ s}^{-1}$ from Krumholz et al. (2007b), which is slightly different from $\alpha_B \simeq 2.56 \times 10^{-13} (T/10^4 \text{ K})^{-0.83} \text{ cm}^3 \text{ s}^{-1}$ adopted by Draine (2011).

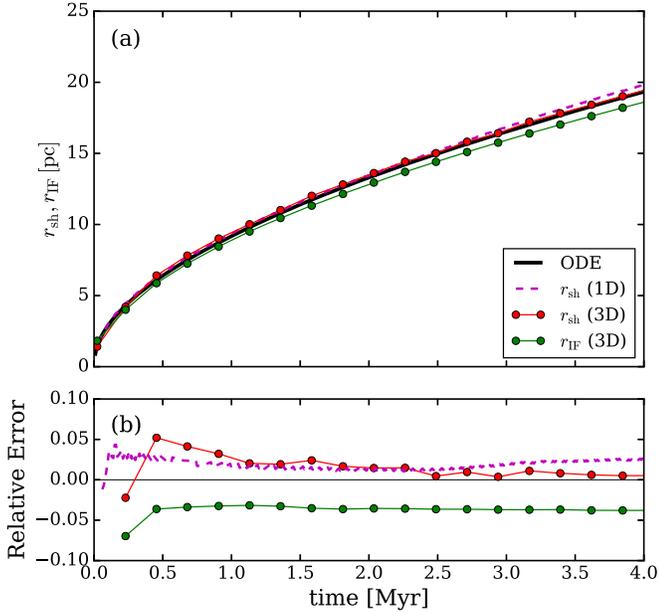


Figure 8. (a) Shell radius (red) and ionization front radius (green) vs. time for an expanding dusty H II region with $Q_i = 10^{51} \text{ s}^{-1}$ and $n_{\text{H}} = 10^3 \text{ cm}^{-3}$. The black solid line draws the solution of Equation (28), while the dashed line gives the shell radius obtained from the one-dimensional simulation. (b) Relative errors with respect to the solution of Equation (28).

principal axes. The mass accretion rate onto the sink particle is determined by the flux returned from the Riemann solver at the boundary faces of the 3^3 ghost cells surrounding the sink particle. The ghost cells in this sink particle control volume are reset after the hydrodynamics integration by extrapolation from the nearest active cells. To prevent spurious mass and momentum flows from the ghost cells to the active cells, we use the “diode”-like boundary condition.¹¹ When two sink particles come close enough together to make their control volumes overlap, we simply merge them in such a way that conserves the total mass and momentum. We have also tested the sink particle method of Bleuler & Teyssier (2014) and found similar results.

The gravitational potential from gas and star particles is calculated using an FFT Poisson solver with open boundary conditions (Skinner & Ostriker 2015) after mapping star particles’ mass onto the mesh via the triangular-shaped-cloud scheme (Gong & Ostriker 2013). For the ART, we adopt $m_{\text{ray}} = 4$ and $C = 0.1$ in the subcycling and rotate ray orientation randomly every hydrodynamic cycle. Since the control volume encompassing a star particle is regarded as a ghost zone, we do not allow radiation to interact with the gas in the control volume.

Due to limited mass resolution, star particles in our simulation represent subclusters rather than individual stars. Star particles emit radiation in two frequency bins: hydrogen ionizing and nonionizing photons with luminosity denoted by $L_i = Q_i h\nu_i$ and $L_n = L - L_i$, respectively. To assign the luminosity, we first calculate the total sink mass M_* in the whole domain at a given time. We then use the stellar population synthesis code SLUG based on the Chabrier initial mass function (Krumholz et al. 2015) to calculate the light-to-mass ratio $\Psi(M_*)$ and the ionizing photon

¹¹ This boundary condition is the same as the outflow boundary condition, except that the normal velocity component at the boundaries is set equal to zero if gas flows from the ghost to active cells.

Table 1
Global Properties of the Radiation Fields from the ART and M_1 Schemes

f_* (%) (1)	$t/t_{\text{ff},0}$ (2)	$f_{\text{esc,ART}}$ (3)	f_{esc,M_1} (4)	$\mathcal{F}_{\text{out,ART}}$ (5)	$\mathcal{F}_{\text{out},M_1}$ (6)
10	0.7	0.29	0.28	0.079	0.073
50	1.0	0.29	0.30	0.20	0.15
90	1.75	0.63	0.63	0.64	0.59

Note. Column (1): fraction of the total stellar mass formed. Column (2): time in units of the initial free-fall time. Columns (3) and (4): escape fraction of the photons through an enclosing sphere computed by the ART and M_1 methods, respectively. Columns (5) and (6): normalized, volume-integrated radial force from the ART and M_1 methods, respectively. The measurements of f_{esc} and \mathcal{F}_{out} are made for spheres with radii 22, 25, and 29 pc centered at r_{CM} at $t_{10\%}$, $t_{50\%}$, and $t_{90\%}$, respectively.

rate per unit mass $\Xi(M_*)$ for a cluster with mass M_* at birth (Kim et al. 2016).¹² Finally, a sink particle with mass m_* is assigned to emit $L = \Psi(M_*)m_*$ and $Q_i = \Xi(M_*)m_*$. Using these conversion factors, for example, a stellar cluster of mass $(10^2, 10^3, 10^4) M_\odot$ has a total bolometric luminosity and ionizing photon production rate of $(1.1 \times 10^4, 7.3 \times 10^5, 8.9 \times 10^6) L_\odot$ and $(1.5 \times 10^{46}, 3.5 \times 10^{49}, 4.8 \times 10^{50}) \text{ s}^{-1}$, respectively. For simplicity, we do not allow for age variation of Ψ and Ξ and take constant values of $h\nu_i = 18 \text{ eV}$ for the mean energy of ionizing photons and $\sigma_{\text{H}} = 6.3 \times 10^{-18} \text{ cm}^2$ for the photoionization cross section.

Our problem initialization is largely similar to the approach of Skinner & Ostriker (2015) and Raskutti et al. (2016). We consider an isolated, uniform-density sphere of mass M_{cl} and radius R_{cl} placed at the center of the cubic box with side $L = 4R_{\text{cl}}$. The rest of the box is filled with a rarefied medium with density 10^3 times smaller than that of the cloud. The total gas mass inside the box is thus $1.014M_{\text{cl}}$. We impose a turbulent velocity field realized by a Gaussian random distribution with power spectrum $|v^2| \propto k^{-4}$ over the wave-number range $k \in [2, 64] \times 2\pi/L$ (Stone et al. 1998). The amplitude of the velocity field is adjusted such that the total kinetic energy E_{kin} is equal to the absolute value of the gravitational potential energy $E_G = -\frac{3}{5}GM_{\text{cl}}^2/R_{\text{cl}}$, making the initial cloud marginally gravitationally bound. The corresponding virial parameter is $\alpha_{\text{vir}} \equiv 2E_{\text{K}}/|E_G| = 2$ at $t = 0$. For the particular model presented here, we set $M_{\text{cl}} = 5 \times 10^4 M_\odot$ and radius $R_{\text{cl}} = 15 \text{ pc}$, and we resolve the simulation box using 256^3 cells. The initial cloud conditions are therefore the same as listed in Table 1 for the fiducial model of Raskutti et al. (2016), although here with $T_{\text{neu}} = 20 \text{ K}$, the sound speed in the neutral gas is $c_s = 0.27 \text{ km s}^{-1}$, and the opacity for nonionizing radiation is $\kappa = 500 \text{ cm}^2 \text{ g}^{-1}$ (rather than $c_s = 0.2 \text{ km s}^{-1}$ and $\kappa = 1000 \text{ cm}^2 \text{ g}^{-1}$ in Raskutti et al. 2016), and also here Ψ varies with total stellar mass rather than being set to a fixed value $\Psi = 2000 \text{ erg s}^{-1} \text{ g}^{-1}$.

4.2. Overall Evolution and Scaling Test

The initial supersonic turbulence creates filamentary structure via shock compression. Subsequently, the filaments become gravitationally unstable and fragment into clumps that undergo runaway collapse to form stars. The first star particle is

¹² Equation (34) in Kim et al. (2016) for the conversion factor $\Xi(M_*)$ contains typographical errors. The correct equation should read $\log_{10}(\Xi/1 \text{ s}^{-1} M_\odot) = 46.7 \mathcal{X}^7 / (7.28 + \mathcal{X}^7)$, where $\mathcal{X} = \log_{10}(M_*/M_\odot)$.

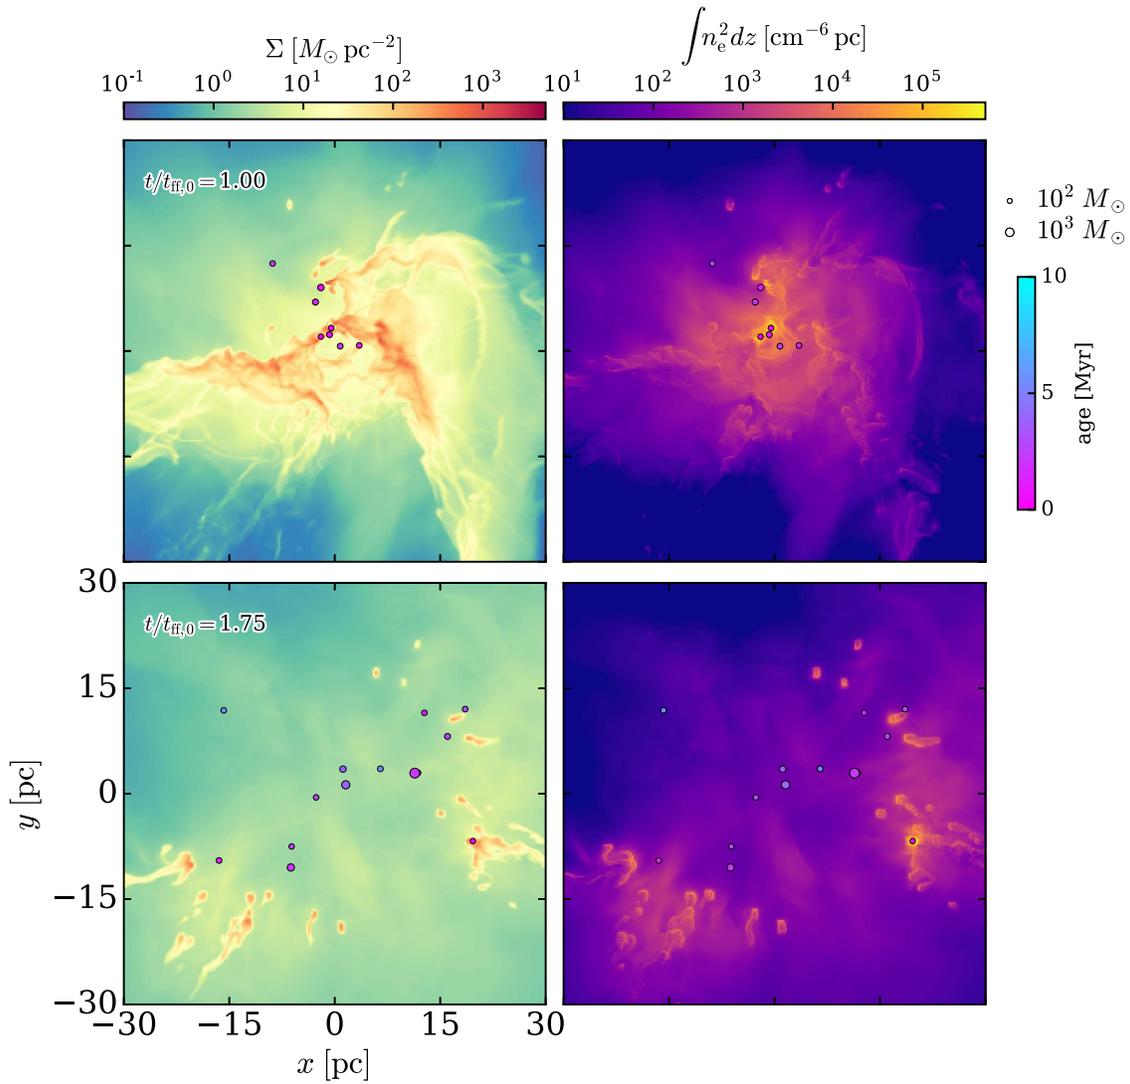


Figure 9. Snapshots of gas surface density (left) and emission measure of ionized gas (right) integrated along the z -axis at $t/t_{\text{ff},0} = 1.0$ (top) and 1.75 (bottom) for the model with $M_{\text{cl}} = 5 \times 10^4 M_{\odot}$ and $R_{\text{cl}} = 15$ pc. The circles mark the projected positions of star particles, with color and size representing their age and mass, respectively.

spawned at $t/t_{\text{ff},0} = 0.44$, where $t_{\text{ff},0} = 4.3$ Myr is the initial free-fall time of the cloud. Star formation continues until $\sim 1.7t_{\text{ff},0}$, creating a total of 24 sink particles. Figure 9 plots the snapshots of gas surface density (left) and emission measure of the ionized gas (right) on the x - y plane at $t/t_{\text{ff},0} = 1.0$ and 1.75, when 32% and 99% of the final stellar mass has assembled, respectively. The star particle positions are marked as circles with their size proportional to the mass. At $t/t_{\text{ff},0} = 2$, the net star formation efficiency is only 12%, with most of the gas ($\gtrsim 70\%$) pushed out of the simulation box owing to photoionization and radiation pressure. We defer to a forthcoming paper a detailed presentation of simulation results on the star formation efficiency, cloud lifetime, the role of photoevaporation versus radiation pressure in cloud disruption, etc., and their dependence on the cloud parameters.

We perform a strong scaling test for this realistic star formation model, varying the number of cores from $N_{\text{core}} = 64$ to 1024. While the time for the hydrodynamic and self-gravity updates remains roughly constant throughout the simulation, the cost of the radiation update scales with the number

of sources and the number of ART substeps taken per hydrodynamic update (typically ~ 10 – 20). Figure 10 plots the average zone cycles per second (i.e., the total number of cells divided by the CPU time) for different physics modules during the time interval of $1.16 \leq t/t_{\text{ff},0} \leq 1.19$, when $\sim 97\%$ of the computational domain is filled with ionized gas and the number of point sources is 13. The zone cycles per second for hydrodynamics and gravity are weakly decreasing functions of N_{core} because of the increasing computation-to-communication ratio. Although the radiation update is the most expensive part owing to multiple sources and subcycling, the cost of the ART normalized by the number of sources and the number of substeps remains roughly constant and is in fact cheaper than the hydrodynamics update. The relative parallel efficiency from 64 to 1024 processors is 72%, 22%, and 67% for the radiation, gravity, and hydrodynamic updates, respectively.¹³

¹³ Inefficient scaling of the FFT gravity module for the largest N_{core} is likely caused by communication overhead associated with global transpose of data among processors.

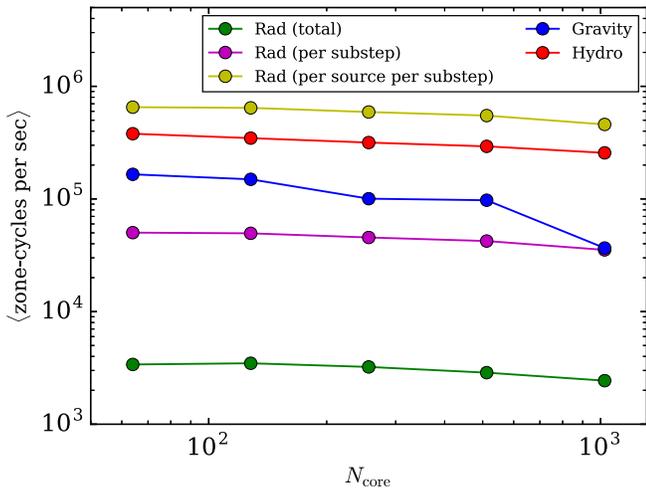


Figure 10. Strong scaling test of a star cluster simulation with $N_{\text{cell}} = 256^3$. The average zone cycles per second during $1.16 \leq t/t_{\text{ff},0} \leq 1.19$ are shown against the number of cores for different physics modules. During this time interval, there are 13 radiation sources, and the typical number of substeps taken per hydrodynamic update is 14. The parallel efficiency of the radiation module from 64 to 1024 cores is 72%.

4.3. Comparison of Radiation Field Computed from M_1 Closure and ART

Raskutti et al. (2016) performed RHD simulations of star formation in turbulent molecular clouds regulated by nonionizing radiation pressure on dust alone. They used the *Hyperion* code (Skinner & Ostriker 2013) employing the M_1 closure relation to evolve the two-moment RT equations. The *Hyperion* radiation solver uses explicit integration in time with a reduced speed of light, employing an HLL-type Riemann solver to compute radiation fluxes between cells, the piecewise linear spatial reconstruction, and the first-order backward Euler method for radiation energy and flux absorption source term updates. For their fiducial model with $M_{\text{cl}} = 5 \times 10^4 M_{\odot}$, $R_{\text{cl}} = 15$ pc, and $\alpha_{\text{vir}} = 2.0$, Raskutti et al. (2016) obtained a final star formation efficiency of 42%, larger than the 12% found in our ART simulation that considers both ionizing and nonionizing radiation, as described above.

To study whether the difference in the star formation efficiency is due to the effect of ionizing radiation and/or to the method of solving the RT equations, we have run an additional ART simulation with the same set of cloud parameters, but by turning off ionizing radiation. To match the parameters of Raskutti et al. (2016), for this comparison run we adopt a constant mass-to-light ratio $\Psi = 2000 \text{ erg s}^{-1} \text{ g}^{-1}$, dust absorption cross section $\sigma_{\text{d}} = 2.34 \text{ cm}^2 \text{ g}^{-1}$, and an isothermal equation of state with isothermal sound speed $c_s = 0.2 \text{ km s}^{-1}$. While overall evolution of this model is very similar to that of Raskutti et al. (2016), the resulting final star formation efficiency is $\varepsilon = 0.25$. This discrepancy in ε between the two results can be largely attributed to the difference in the radiation solvers.

Although the M_1 closure scheme has been benchmarked for a variety of idealized test problems (e.g., González et al. 2007; Rosdahl et al. 2013; Skinner & Ostriker 2013; Rosdahl & Teyssier 2015), it is valuable to check how reliable its radiation field is compared to that based on the ART method for a problem with complex distributions of sources and matter. For this purpose, we take the data at two different epochs from the ART simulation with only nonionizing radiation and use them

as inputs to the *Hyperion* radiation solver. With *Hyperion* we read in the gas density and the positions \mathbf{r}_* and masses m_* of all sink particles, and we assign the source function j_* of a Gaussian form to each sink as

$$j_*(\mathbf{r}) = \frac{m_* \Psi}{(2\pi\sigma_*^2)^{3/2}} \exp\left(-\frac{|\mathbf{r} - \mathbf{r}_*|^2}{2\sigma_*^2}\right), \quad (29)$$

where the width is taken to be $\sigma_* = (2 \log 2)^{-1/2}$ pc; the corresponding FWHM is 1 pc (Raskutti et al. 2016). We then perform radiation updates, while fixing hydrodynamic variables and the star particles. We integrate radiation variables over 10^3 yr, sufficiently long for the radiation field to reach a steady state.

To explore the differences between the two radiation solutions, we choose two epochs, $t_{10\%} = 0.7t_{\text{ff},0}$ and $t_{90\%} = 1.75t_{\text{ff},0}$, when 10% and 90% of the final stellar mass has assembled, respectively. The former corresponds to an early stage of cluster formation when five sink particles have been created, which are mostly embedded in dense nodes of filaments. At $t = t_{90\%}$, 39 sources are distributed more evenly around the center of the domain with half-mass radius 8.2 pc; radiation has cleared out most of the gas in the immediate surroundings. At the latter time, the system is globally super-Eddington, with the remaining gas actively expelled from the domain by the radiation force. The center of mass, \mathbf{r}_{CM} , of the sink particles is $(-1.9, 5.9, -5.6)$ pc and $(1.1, 0.3, 0.6)$ pc at $t_{10\%}$ and $t_{90\%}$, respectively.

Figure 11 shows the gas structure for the (a) $t_{10\%}$ and (b) $t_{90\%}$ epochs, in slices centered at the star particle located at $\mathbf{r}_{\text{sink}} = (-0.8, -0.4, -20.3)$ pc and the center of mass of all star particles $\mathbf{r}_{\text{CM}} = (1.1, 0.3, 0.6)$ pc, respectively. The white circles mark the positions of star particles within 3 pc from the slice plane. The arrows indicate the directions of the radiation fluxes computed by the ART scheme (white) and M_1 scheme (red). At $t_{10\%}$, there are some discrepancies in the direction of the radiation fluxes. In the immediate vicinity of the star particle, the radiation fluxes computed by the ART scheme are directed radially outward from the sink particle, while those from the M_1 scheme have a nonradial component. In the upper middle region, the radiation flux computed by the ART scheme is directed downward (resulting from a star particle at $z \sim -15$ pc outside the panel), while the M_1 radiation flux is directed upward. We note that this region is heavily shielded by an intervening dense clump located at $(x, y) \sim (-1, -20)$ pc and the radiation field is very weak compared to that in the immediate vicinity of the sink or in other regions at the same distance from the sink. Therefore, while the M_1 flux is inaccurate in this region, it is much smaller (typically by a factor of 10–30) than the mean value at similar distances. At $t_{90\%}$, the radiation fields from the two methods are in good overall agreement with each other, although the flux directions are somewhat different in the lower middle region, where the gas density is quite low and multiple beams going in different directions cross.

One quantitative measure of the global dynamical impact of radiation fields is the radial component, $f_{\text{rad},r}$, of the radiation force (per unit volume) relative to \mathbf{r}_{CM} . At each point \mathbf{r} in the domain, this is given by

$$f_{\text{rad},r} = \frac{\chi \mathbf{F}}{c} \cdot \frac{\mathbf{r} - \mathbf{r}_0}{|\mathbf{r} - \mathbf{r}_0|}, \quad (30)$$

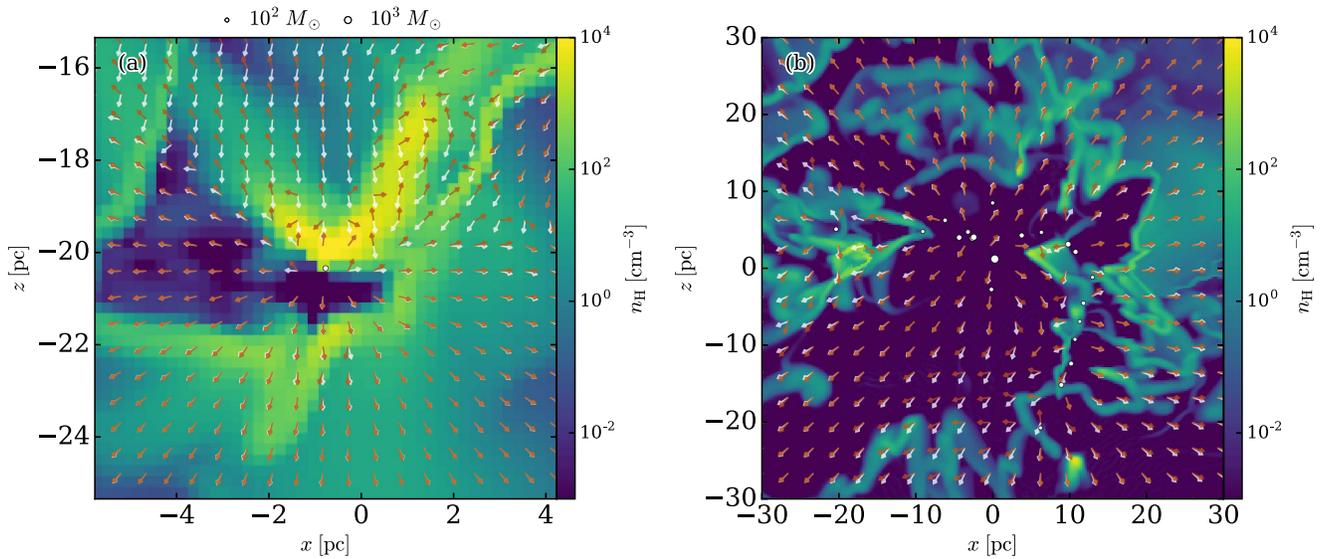


Figure 11. Snapshots of the control model with radiation pressure alone in the plane (a) with $y = -0.4$ pc at $t_{10\%} = 0.7t_{\text{ff},0}$ and (b) with $y = 0.3$ pc at $t_{90\%} = 1.75t_{\text{ff},0}$. All star particles lying within ± 3 pc of the slice plane are shown as white circles. The arrows in white and red represent the direction of the radiation flux \mathbf{F} calculated from the ART and M_1 method, respectively.

where $\mathbf{r}_0 = \mathbf{r}_{\text{CM}}$. Figure 12 compares the probability distribution functions (pdf's) of $|f_{\text{rad},r}|$ computed from ART and from the M_1 method at $t = t_{10\%}$ (left) and $t_{90\%}$ (right); pdf's are shown weighted by the cell mass (top) and volume (bottom). Overall, the M_1 method reproduces the ART radiation field reasonably well, with $|f_{\text{rad},r}(M_1)|/|f_{\text{rad},r}(\text{ART})| \approx 1.0$ at $t_{10\%}$ and 0.89 at $t_{90\%}$, when averaged over the entire range of $|f_{\text{rad},r}|$. At $t = t_{10\%}$ there is a significant scatter relative to the one-to-one relationship (dashed lines), which is likely to be caused by the inability of the M_1 scheme to describe the superposition of streaming radiation from multiple sources going in different directions. The scatter is small at $t = t_{90\%}$ because the radiation sources are more centrally concentrated relative to the gas.

A careful examination of Figure 12(a) shows that $|f_{\text{rad},r}(M_1)|/|f_{\text{rad},r}(\text{ART})| \approx 0.55$ for the largest values of $|f_{\text{rad},r}(\text{ART})| \gtrsim 10^{-29} \text{ dyne cm}^{-3}$ at $t = t_{10\%}$. The corresponding physical locations are in immediate regions surrounding sink particles where both density and radiation flux are high. This discrepancy is suggestive of a possible reason why the total star formation efficiency in the *Hyperion* simulation is larger: with a lower radiation force in the immediate vicinity of sink particles, accretion of material onto sinks may not be limited as strongly as it is in the ART simulation.

To investigate this issue in more detail, we examine the radial component of the radiation force centered on individual star particles. This is computed as in Equation (30), except now \mathbf{r}_0 is the location \mathbf{r}_{sink} of an individual star particle. For example, Figure 13(a) compares the angle-averaged radiation force $\langle f_{\text{rad},r} \rangle = \int f_{\text{rad},r} d\Omega / \int d\Omega$ relative to a sink particle located at $\mathbf{r}_{\text{sink}} = (-0.8, -0.4, -20.3)$ pc, isolated from other sink particles. Although the M_1 radiation force is in good agreement with the ART radiation force for $|\mathbf{r} - \mathbf{r}_{\text{sink}}| > 1.5$ pc, the former significantly underestimates the radiation force within $\lesssim 1.5$ pc of the source. Since the source function in the M_1 scheme (Equation (29)) is smoothed over the finite width of ~ 1 pc, it takes ~ 4 – 5 cells to build up the flux expected for a point-source radiation. In contrast, the ART scheme with the adopted parameters has been proven successful in reproducing

a point-source radiation field with errors of only a few percent (see Section 3.2). Interestingly, the magnitude of the gravitational force (open circles) at $|\mathbf{r} - \mathbf{r}_{\text{sink}}| \lesssim 1$ pc is larger than the radiation force from the M_1 scheme but smaller than that from the ART scheme. This suggests that the direct radiation pressure feedback in the former is less effective in halting accretion than in the latter. The difficulty in resolving flux $\propto r^{-2}$ near point sources in the M_1 scheme may explain the difference in the star formation efficiency between the two simulations.

Although the radiation fields from the two methods are quite different within a few zones of the point sources, they agree both at larger distances from individual sources and on global scales, particularly at late times. Figure 13(b) plots the angle-averaged radial force as computed with ART and with M_1 as functions of the distance from the center of mass of all sinks at $t = t_{90\%}$. Despite a complex source geometry, the results from the two methods agree very well. The fraction of radiation escaping the surface of a sphere of radius $r_{\text{max}} = 2R_{\text{cl}} - |\mathbf{r}_{\text{CM}}| = 29$ pc centered at \mathbf{r}_{CM} is equal at $f_{\text{esc}} = 0.63$ between the two methods.

We also calculate the normalized volume-integrated radial force

$$\mathcal{F}_{\text{out}} = \frac{4\pi c}{(1 - f_{\text{esc}})L_{*,\text{tot}}} \int_0^{r_{\text{max}}} \langle f_{\text{rad},r} \rangle r^2 dr, \quad (31)$$

which would have been equal to unity had all sources with total luminosity $L_{*,\text{tot}}$ been located at \mathbf{r}_{CM} . For multiple sources distributed in space, \mathcal{F}_{out} is reduced owing both to flux cancellation and to misalignment of the radiation force vectors from the radial direction toward \mathbf{r}_{CM} . The values of f_{esc} and \mathcal{F}_{out} from the ART and M_1 schemes are quite comparable to each other. Table 1 summarizes these global properties of the radiation fields for the two radiation solvers at $t = t_{10\%}$, $t_{50\%}$, and $t_{90\%}$.

5. Summary

Radiation feedback from young massive stars has a profound influence on the evolution of their natal clouds and subsequent

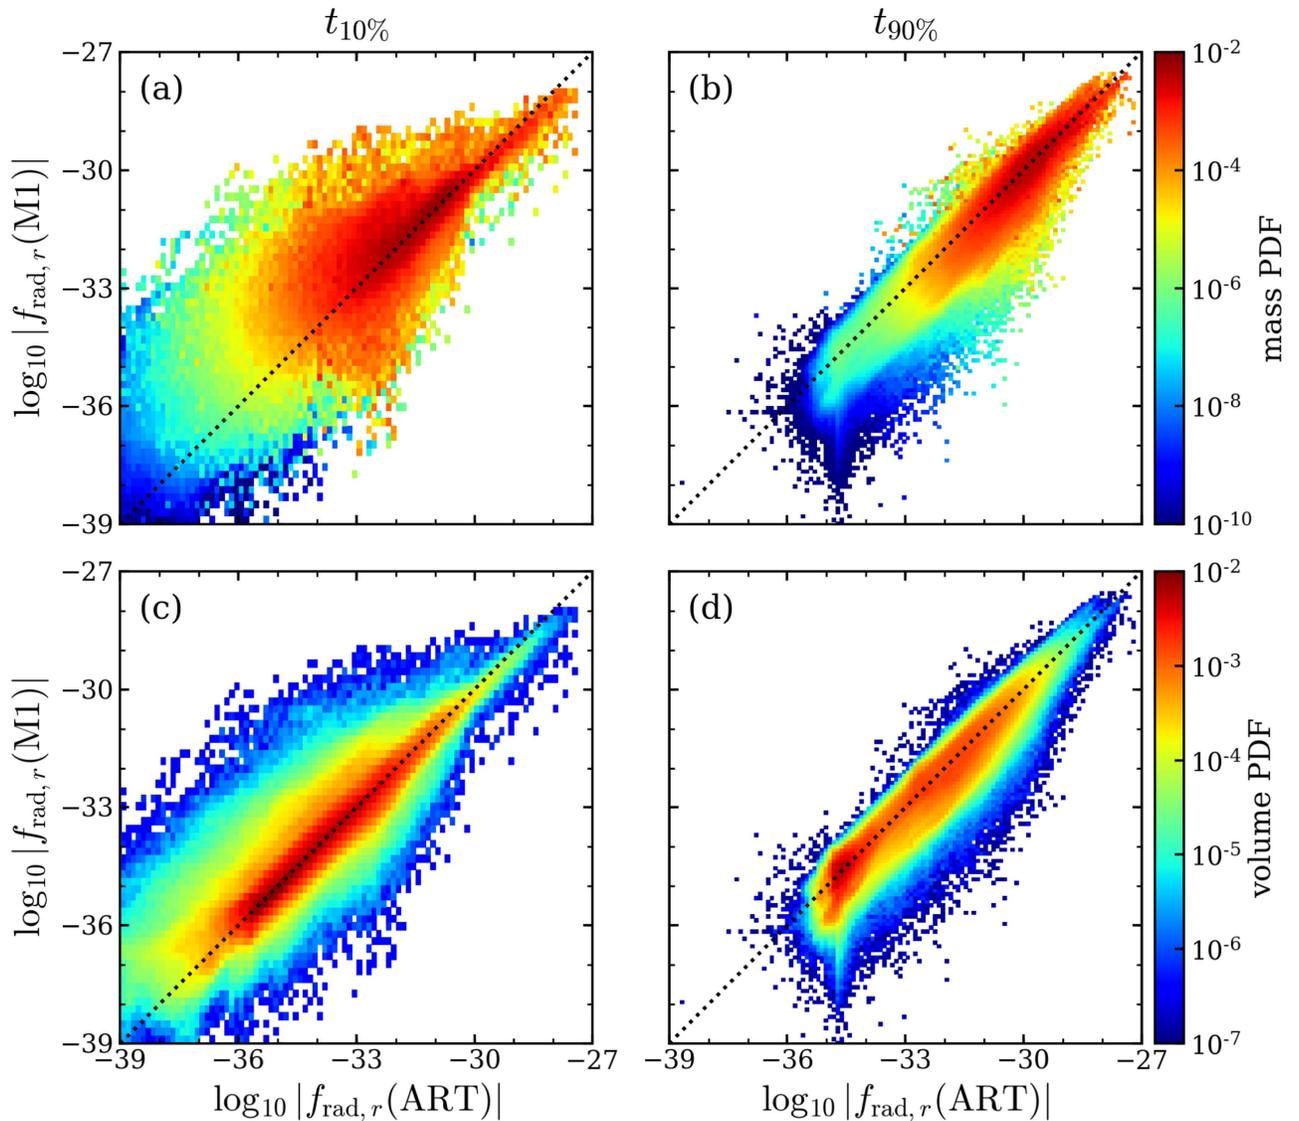


Figure 12. Comparison of the radial component (relative to the CM) of the radiation force (in units of dynes cm^{-3}) computed by the ART and M_1 schemes at times $t_{10\%}$ (left) and $t_{90\%}$ (right). The upper and lower panels show mass- and volume-weighted probability density distributions, respectively. For comparison, the dashed lines show the relation that would be obtained if the radiation force were identical for the two methods.

star formation. Stellar radiation that escapes from dense molecular clouds is also essential in controlling the thermal and ionization states of the diffuse ISM in galaxies and contributes to ionization of the IGM at even further distances. Because of the long-range, multiscale nature of the interactions between matter and radiation in three dimensions and the highly inhomogeneous spatial structure of the gas and radiation source distribution, it is nontrivial to handle RT properly in hydrodynamic simulations. The ART algorithm developed by Abel & Wandelt (2002) provides an accurate means of treating RT from point sources, which retains spatial resolution at a moderate cost by splitting rays as the distance from a given source increases. Rosen et al. (2017) recently improved the parallel performance of the ART method by implementing completely nonblocking, asynchronous MPI communication.

In this paper, we describe our implementation of ART in the Eulerian grid-based code *Athena* and present results of performance tests. We adopt the nonblocking, asynchronous parallelization algorithm suggested by Rosen et al. (2017) for exchanges of information along rays between processors. We

further improve the parallel performance by (1) passing photon packets to neighbor processors whenever a certain number of local ray traces have been completed and (2) making use of one-sided communications to update the “destroy count” of terminated rays. The radiation source terms, hydrogen photoionization, radiation pressure, etc., are all substepped (relative to the hydrodynamic time step) and updated in an operator-split manner after each ART sweep, with the substepping time interval set by the ionization/recombination time via Equation (21). To update the hydrogen ionization state, we use an approximate solution (Equation (20)) of the rate equation, which we find is more numerically robust and gives results that are essentially the same as the full analytic solution (Equation (17)).

We have verified the performance and accuracy of our implementation of the ART scheme on a wide variety of test problems. The results of weak and strong scaling tests (Figure 2) show that the cost of the ART (per source) remains comparable to that of the hydrodynamic update on up to 10^3 processors. The vacuum radiation test shows that if the number

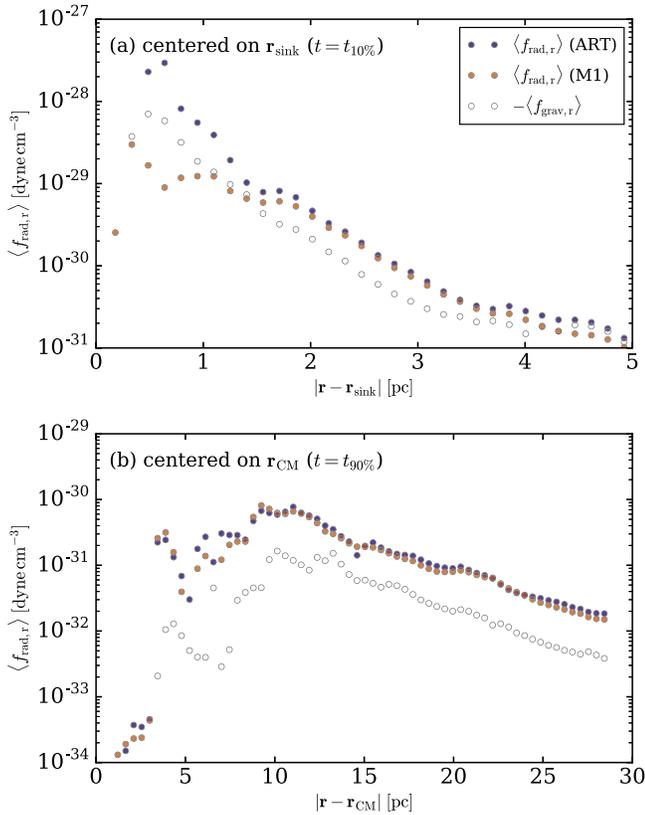


Figure 13. Angle-averaged radial radiation force computed by ART (blue) and M_1 (red) centered (a) at a sink particle at $t = t_{10\%}$ and (b) at the center of mass of all sinks at $t = t_{90\%}$. The absolute values of the angle-averaged gravitational forces are shown as open circles.

of rays passing through a cell is $m_{\text{ray}} = 4$, the median value of the errors in the calculated radiation energy density is only $\sim 1\%$ – 4% , and the accuracy can be further improved by rotating the ray directions randomly (Krumholz et al. 2007b) at each step (Figure 3). Through standard test problems for the expansion of H II regions, we demonstrate that our radiation solver reproduces quite well the expected solutions of expanding R-type and D-type ionization fronts, as well as the expansion of a dusty H II region with radiation pressure (Figures 4–8).

As a practical application demonstrating the use of our code, we conduct a simulation of star cluster formation in a turbulent, self-gravitating molecular cloud with $M_{\text{cl}} = 5 \times 10^4 M_{\odot}$, $R_{\text{cl}} = 15$ pc, and initial virial parameter of $\alpha_{\text{vir}} = 2$. We find that the net star formation efficiency is $\varepsilon = 0.12$, with most of the gas expelled from the simulation box when both photoionization and radiation pressure from UV radiation are present. The strong scaling test for this problem (Figure 10) shows that the parallel efficiency of the ART module is as good as that of the hydrodynamics module. The total cost is, however, dominated by the radiation update, which involves multiple sources and subcycling.

We have also run an analogous model with radiation pressure alone (i.e., without photoionizing radiation), in order to directly compare with results obtained from the *Hyperion* code, which solves the two-moment radiation equations with an M_1 closure relation. Considering an identical radiation source and density distribution, the radiation field computed using the M_1 scheme agrees with that from ART on large scales, even for distributed sources. Since, however, point sources are

smoothed over a finite volume in *Hyperion*, it is not able to reproduce the radiation field accurately near individual sources. As a consequence of the reduced radiation flux near point sources, radiation feedback is less able to limit accretion of nearby material, which likely accounts for the increased net star formation efficiency found with *Hyperion* ($\varepsilon = 0.42$) compared to that with our ART implementation ($\varepsilon = 0.25$). We conclude that one should be cautious when modeling point sources using the M_1 scheme if radiation feedback is important to limiting accretion; one approach might be to expand the control volume of each sink particle such that the radiation flux is well resolved at its boundary.

The test results presented in this paper confirm that our implementation of ART in the *Athena* code is accurate and efficient. In a subsequent work, we shall present results from application of the code to cluster formation and radiation feedback in turbulent molecular clouds. With an accurate and efficient method for treating the effects of radiation, we are able to survey a range of parameters, studying the dependence on the cloud mass and surface density of the star formation efficiency and cloud lifetime, as well as the relative roles of photoevaporation and radiation pressure in shaping and disrupting GMCs.

The authors thank Anna Rosen for making a copy of her code and manuscript on adaptive ray tracing available to us before publication. J.-G.K. wishes to thank Shane Davis, Chang-Goo Kim, Jim Stone, Kengo Tomida, and Gábor Tóth for their helpful discussions and advice. J.-G.K. acknowledges support from the National Research Foundation of Korea (NRF) through the grant NRF-2014-Fostering Core Leaders of the Future Basic Science Program. The work of W.-T.K. was supported by the grant (2017R1A4A1015178) of National Research Foundation of Korea. The work of ECO on this project was supported by the U.S. National Science Foundation under grant AST-1312006 and by NASA under grant NNX14AB49G. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. This article has been assigned an LLNL document release number LLNL-JRNL-737937. The computation of this work was supported by the Supercomputing Center/Korea Institute of Science and Technology Information with supercomputing resources including technical support (KSC-2015-C3-049) and by the PICSciE TIGRESS High Performance Computing Center at Princeton University.

Software: Athena (Stone et al. 2008), SLUG (Krumholz et al. 2015), yt (Turk et al. 2011), numpy (van der Walt et al. 2011), matplotlib (Hunter 2007), IPython (Pérez & Granger 2007), pandas (McKinney 2010).

Appendix

One-cell Test of Photoionization Update

Here we present test results for the hydrogen photoionization update. We consider a single cell with width $\Delta x = 2$ pc initially filled with completely neutral hydrogen of density n_{H} . The cell is exposed to a fixed ionizing radiation field $Q_i = 10^{49} \text{ s}^{-1}$ from $t = 0$. The temporal evolution of the neutral hydrogen fraction in the cell is described by Equation (16), with the photoionization rate given as

$$\Gamma = \frac{1}{n_{\text{H}} x_{\text{n}}} \frac{Q_i}{(\Delta x)^3} (1 - e^{-n_{\text{H}} x_{\text{n}} \sigma_{\text{H}} \Delta x}). \quad (32)$$

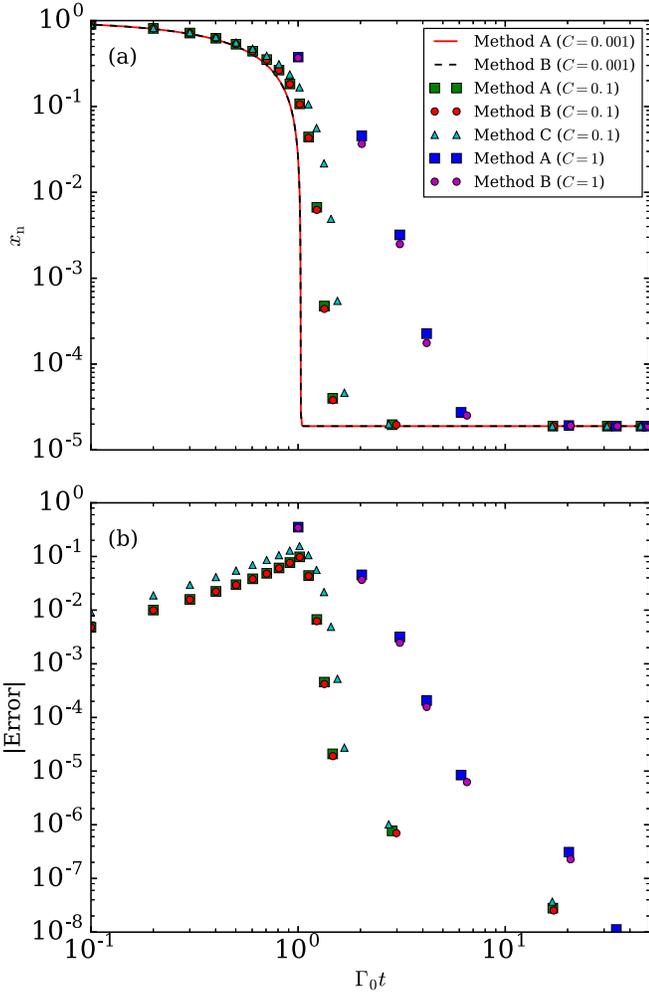


Figure 14. (a) Neutral fraction x_n vs. time obtained using the different update schemes for x_n with various C . Methods A and B use Equations (16) and (17), respectively, while Method C employs Equation (34). (b) Difference of x_n relative to the results with $C = 0.001$, shown as the red solid line in panel (a).

The cell has initial optical depth $\tau_0 = n_H \sigma_H \Delta x = 3880 \gg 1$ and initial ionization rate $\Gamma_0 = Q_i / (n_H \Delta x^3) = 1.7 \times 10^{-6} \text{ s}^{-1} \approx 14 \alpha_B n_H$. The cell is expected to be ionized on the timescale Γ_0^{-1} , eventually settling into a balanced state with an equilibrium neutral fraction of $x_{\text{eq}} \approx \alpha_B n_H / (\Gamma_0 \tau_0) = 1.8 \times 10^{-5}$.

To evolve x_n according to Equation (16), we try three different schemes. First, we directly use the exact expression of Equation (17) with varying values of C in Equation (21), which we call Method A. The second method (Method B) uses Equation (20), again with varying C . We also use a semi-implicit difference method (Method C) to discretize Equation (16) as

$$\frac{x_n^{n+1} - x_n^n}{\Delta t_{\text{ss}}} = \alpha_B n_H (1 - x_n^n)(1 - x_n^{n+1}) - x_n^{n+1} \Gamma, \quad (33)$$

which gives a recurrence relation

$$x_n^{n+1} = \frac{x_n^n + (1 - x_n^n)^2 \alpha_B n_H \Delta t_{\text{ss}}}{1 + \Gamma \Delta t_{\text{ss}}}, \quad (34)$$

where the superscripts “ n ” and “ $n + 1$ ” denote the values at $t = t_0$ and $t = t_0 + \Delta t_{\text{ss}}$, respectively.

Figure 14 plots the resulting temporal changes of x_n from the various methods with different C . The results of Methods A and B with $C = 0.001$, which are almost identical to each other, can be regarded as the true solution. The neutral fraction settles to an equilibrium value at $t \Gamma_0 \gtrsim 1$. Panel (b) plots the errors relative to the $C = 0.001$ results. It is remarkable that the difference between the results of Methods A and B is very small even for C as large as unity and that they are better than those from Method C. The error is largest when $x_n \sim 0.1$ – 0.2 , for which $|dx_n/dt|$ is quite large.

ORCID iDs

Jeong-Gyu Kim <https://orcid.org/0000-0001-6228-8634>
 Woong-Tae Kim <https://orcid.org/0000-0003-4625-229X>
 Eve C. Ostriker <https://orcid.org/0000-0002-0509-9113>
 M. Aaron Skinner <https://orcid.org/0000-0001-8353-8305>

References

- Abel, T., Norman, M. L., & Madau, P. 1999, *ApJ*, 523, 66
 Abel, T., & Wandelt, B. D. 2002, *MNRAS*, 330, L53
 Altay, G., & Theuns, T. 2013, *MNRAS*, 434, 748
 Anninos, P., Zhang, Y., Abel, T., & Norman, M. L. 1997, *NewA*, 2, 209
 Baczynski, C., Glover, S. C. O., & Klessen, R. S. 2015, *MNRAS*, 454, 380
 Barkana, R., & Loeb, A. 2001, *PhR*, 349, 125
 Bisbas, T. G., Haworth, T. J., Williams, R. J. R., et al. 2015, *MNRAS*, 453, 1324
 Bisbas, T. G., Wunsch, R., Whitworth, A. P., & Hubber, D. A. 2009, *A&A*, 497, 649
 Bleuler, A., & Teyssier, R. 2014, *MNRAS*, 445, 4015
 Castor, J. I. 2004, *Radiation Hydrodynamics* (Cambridge: Cambridge Univ. Press), 368
 Dale, J. E. 2017, *MNRAS*, 467, 1067
 Dale, J. E., Ercolano, B., & Bonnell, I. A. 2012, *MNRAS*, 424, 377
 Dale, J. E., Ercolano, B., & Bonnell, I. A. 2013, *MNRAS*, 430, 234
 Dale, J. E., Ngoumou, J., Ercolano, B., & Bonnell, I. A. 2014, *MNRAS*, 442, 694
 Davis, S. W., Jiang, Y.-F., Stone, J. M., & Murray, N. 2014, *ApJ*, 796, 107
 Davis, S. W., Stone, J. M., & Jiang, Y.-F. 2012, *ApJS*, 199, 9
 Draine, B. T. 2011, *ApJ*, 732, 100
 Fall, S. M., Krumholz, M. R., & Matzner, C. D. 2010, *ApJL*, 710, L142
 Finlator, K., Özel, F., & Davé, R. 2009, *MNRAS*, 393, 1090
 Frank, M., Hauck, C. D., & Olbrant, E. 2012, arXiv:1208.0772
 Gavagnin, E., Bleuler, A., Rosdahl, J., & Teyssier, R. 2017, *MNRAS*, 472, A155
 Geen, S., Hennebelle, P., Tremblin, P., & Rosdahl, J. 2016, *MNRAS*, 463, 3129
 Gong, H., & Ostriker, E. C. 2013, *ApJS*, 204, 8
 González, M., Audit, E., & Huynh, P. 2007, *A&A*, 464, 429
 González, M., Vaytet, N., Commerçon, B., & Masson, J. 2015, *A&A*, 578, A12
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
 Gritschneider, M., Naab, T., Walch, S., Burkert, A., & Heitsch, F. 2009, *ApJL*, 694, L26
 Grudić, M. Y., Hopkins, P. F., Faucher-Giguère, C.-A., et al. 2016, arXiv:1612.05635
 Haffner, L. M., Dettmar, R.-J., Beckman, J. E., et al. 2009, *RvMP*, 81, 969
 Haworth, T. J., & Harries, T. J. 2012, *MNRAS*, 420, 562
 Hayes, J. C., & Norman, M. L. 2003, *ApJS*, 147, 197
 Henney, W. J., Arthur, S. J., Williams, R. J. R., & Ferland, G. J. 2005, *ApJ*, 621, 328
 Hosokawa, T., & Inutsuka, S.-i. 2006, *ApJ*, 646, 240
 Howard, C. S., Pudritz, R. E., & Harris, W. E. 2016, *MNRAS*, 461, 2953
 Hunter, J. D. 2007, *CSE*, 9, 90
 Kim, C.-G., Ostriker, E. C., & Kim, W.-T. 2013, *ApJ*, 776, 1
 Kim, J.-G., Kim, W.-T., & Ostriker, E. C. 2016, *ApJ*, 819, 137
 Krumholz, M. R. 2006, *ApJL*, 641, L45
 Krumholz, M. R., Fumagalli, M., da Silva, R. L., Rendahl, T., & Parra, J. 2015, *MNRAS*, 452, 1447
 Krumholz, M. R., Klein, R. I., McKee, C. F., & Bolstad, J. 2007a, *ApJ*, 667, 626
 Krumholz, M. R., & Matzner, C. D. 2009, *ApJ*, 703, 1352

- Krumholz, M. R., Matzner, C. D., & McKee, C. F. 2006, *ApJ*, 653, 361
- Krumholz, M. R., Stone, J. M., & Gardiner, T. A. 2007b, *ApJ*, 671, 518
- Krumholz, M. R., et al. 2014, in *Protostars and Planets VI*, ed. H. Beuther et al. (Tucson, AZ: Univ. Arizona Press), 243
- Lefloch, B., & Lazareff, B. 1994, *A&A*, 289, 559
- Levermore, C. D. 1984, *JQSRT*, 31, 149
- Levermore, C. D., & Pomraning, G. C. 1981, *ApJ*, 248, 321
- Lim, A. J., & Mellema, G. 2003, *A&A*, 405, 189
- Lopez, L. A., Krumholz, M. R., Bolatto, A. D., Prochaska, J. X., & Ramirez-Ruiz, E. 2011, *ApJ*, 731, 91
- Mackey, J. 2012, *A&A*, 539, A147
- Mackey, J., Mohamed, S., Gvaramadze, V. V., et al. 2014, *Natur*, 512, 282
- Matzner, C. D. 2002, *ApJ*, 566, 302
- McKinney, W. 2010, in *Proc. 9th Python in Science Conf., Data Structures for Statistical Computing in Python*, ed. S. van der Walt & J. Millman (Austin, TX: SciPy), 51
- Mellema, G., Iliiev, I. T., Alvarez, M. A., & Shapiro, P. R. 2006, *NewA*, 11, 374
- Mihalas, D., & Auer, L. H. 2001, *JQSRT*, 71, 61
- Murray, N., Quataert, E., & Thompson, T. A. 2010, *ApJ*, 709, 191
- Offner, S. S. R., Klein, R. I., McKee, C. F., & Krumholz, M. R. 2009, *ApJ*, 703, 131
- Osterbrock, D. E. 1989, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei* (Mill Valley, CA: Univ. Science Books)
- Ostriker, E. C., McKee, C. F., & Leroy, A. K. 2010, *ApJ*, 721, 975
- Pérez, F., & Granger, B. E. 2007, *CSE*, 9, 21
- Raskutti, S., Ostriker, E. C., & Skinner, M. A. 2016, *ApJ*, 829, 130
- Rijkhorst, E.-J., Plewa, T., Dubey, A., & Mellema, G. 2006, *A&A*, 452, 907
- Rodríguez-Ramírez, J. C., Raga, A. C., Lora, V., & Cantó, J. 2016, *ApJ*, 833, 256
- Rosdahl, J., Blaizot, J., Aubert, D., Stranex, T., & Teyssier, R. 2013, *MNRAS*, 436, 2188
- Rosdahl, J., & Teyssier, R. 2015, *MNRAS*, 449, 4380
- Rosen, A. L., Krumholz, M. R., Oishi, J. S., Lee, A. T., & Klein, R. I. 2017, *JCoPh*, 330, 924
- Schmidt-Voigt, M., & Koeppen, J. 1987, *A&A*, 174, 211
- Shima, K., Tasker, E. J., & Habe, A. 2016, arXiv:1612.06381
- Shu, F. H. (ed.) 1992, *The Physics of Astrophysics. Volume II: Gas Dynamics* (Mill Valley, CA: Univ. Science Books), 1992
- Skinner, M. A., & Ostriker, E. C. 2013, *ApJS*, 206, 21
- Skinner, M. A., & Ostriker, E. C. 2015, *ApJ*, 809, 187
- Spitzer, L. 1978, *Physical Processes in the Interstellar Medium* (New York: Academic Press)
- Steggles, H. G., Hoare, M. G., & Pittard, J. M. 2017, *MNRAS*, 466, 4573
- Stone, J. M., & Gardiner, T. 2009, *NewA*, 14, 139
- Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, *ApJS*, 178, 137
- Stone, J. M., Ostriker, E. C., & Gammie, C. F. 1998, *ApJL*, 508, L99
- Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *ApJS*, 192, 9
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, 13, 22
- Walch, S. K., Whitworth, A. P., Bisbas, T., Wünsch, R., & Hubber, D. 2012, *MNRAS*, 427, 625
- Whitehouse, S. C., & Bate, M. R. 2006, *MNRAS*, 367, 32
- Whitworth, A. 1979, *MNRAS*, 186, 59
- Williams, J. P., & McKee, C. F. 1997, *ApJ*, 476, 166
- Williams, R. J. R. 2002, *MNRAS*, 331, 693
- Wise, J. H., & Abel, T. 2011, *MNRAS*, 414, 3458
- Wolfire, M. G., McKee, C. F., Hollenbach, D., & Tielens, A. G. G. M. 2003, *ApJ*, 587, 278
- Yoda, T., Handa, T., Kohno, K., et al. 2010, *PASJ*, 62, 1277