



Self-consistent solution of cosmological radiation-hydrodynamics and chemical ionization

Daniel R. Reynolds^{a,*}, John C. Hayes^b, Pascal Paschos^c, Michael L. Norman^{c,d}

^a Mathematics, Southern Methodist University, Dallas, TX 75275-0156, United States

^b Lawrence Livermore National Lab, P.O. Box 808, L-551, Livermore, CA 94551, United States

^c Ctr. for Astrophysics and Space Sciences, U.C. San Diego, La Jolla, CA 92093, United States

^d Physics Department, U.C. San Diego, La Jolla, CA 92093, United States

ARTICLE INFO

Article history:

Received 8 January 2009

Received in revised form 14 May 2009

Accepted 8 June 2009

Available online 17 June 2009

PACS:

98.80.Bp

02.60.Cb

02.60.Lj

Keywords:

Radiation

Hydrodynamics

Ionization

Cosmology

Numerical methods

Implicit methods

ABSTRACT

We consider a PDE system comprising compressible hydrodynamics, flux-limited diffusion radiation transport and chemical ionization kinetics in a cosmologically-expanding universe. Under an operator-split framework, the cosmological hydrodynamics equations are solved through the piecewise parabolic method, as implemented in the Enzo community hydrodynamics code. The remainder of the model, including radiation transport, chemical ionization kinetics, and gas energy feedback, form a stiff coupled PDE system, which we solve using a fully-implicit inexact Newton approach, and which forms the crux of this paper. The inner linear Newton systems are solved using a Schur complement formulation, and employ a multigrid-preconditioned conjugate gradient solver for the inner Schur systems. We describe this approach and provide results on a suite of test problems, demonstrating its accuracy, robustness, and scalability to very large problems.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

A fundamental physics component in cosmology and astrophysics applications is the transport of ionizing radiation along with the interaction of that radiation with the hydrodynamic motion and ionization state of the surrounding gas. One example currently receiving a great deal of attention is cosmic reionization [1], our motivation in this work. Observations indicate that an early population of UV emitting galaxies photoionized the intergalactic hydrogen and helium gas when the universe was about 800 million years old (redshift ~ 8). A computational challenge is to calculate this process self-consistently, coupling the radiative transfer of ionizing photons, the ionization kinetics and photo-heating of the gas, and the attendant hydrodynamic motions. This problem is challenging because the physics is numerically stiff and cosmic reionization is intrinsically three-dimensional, involving the growth, percolation, and overlap of ionization zones around an irregular and evolving distribution of galaxies with time-dependent luminosities. In addition, the problem inherits the large range

* Corresponding author. Tel.: +1 214 768 4339; fax: +1 214 768 2355.

E-mail addresses: reynolds@smu.edu (D.R. Reynolds), jchayes@llnl.gov (J.C. Hayes), ppaschos@minbari.ucsd.edu (P. Paschos), mlnorman@ucsd.edu (M.L. Norman).

of spatial scales ($\sim 10^5$) intrinsic in galaxy formation simulations, necessitating the use of spatially-adaptive mesh or particle-tree methods and large-scale parallel computing [2,3].

A variety of 3D radiative transfer methods are under development to tackle this problem [4]. These necessarily simplify the description of the radiation field in order to render the problem tractable. These methods include ray tracing using long and short characteristics [5–8], Monte Carlo [9,10], and moment methods [11–13]. However, only some of these codes allow solution of the coupled problem on spatially-adaptive grids, and very few allow distributed-memory parallelism. More importantly, all of these codes handle the interactions between hydrodynamic, radiative, and chemical processes in an explicit, operator-split fashion, which ignores stiff couplings that often arise between these components. When this happens, such codes must dramatically reduce allowable timesteps, or resort to subcycling, to ensure stability and accuracy of the coupled simulations.

Radiation transport and chemical ionization time scales are much faster than typical hydrodynamic time scales. This is particularly evident in dense gas bound to galaxies where recombination and light-crossing times are short compared to the age of the universe (Hubble time). Moreover, these processes are very tightly-coupled since radiation induces ionization that in turn affects opacities. While time-explicit, adaptive subcycling schemes have been developed that are capable of returning accurate solutions in all regimes of interest [7], it is our view that for both computational efficiency and solution accuracy, tightly-coupled implicit methods require investigation. Here we present such a method.

We solve ionizing radiation transport, chemical ionization kinetics, and gas photo-heating using a fully-implicit inexact Newton method. This algorithm is coupled to a cosmological hydrodynamics code through an explicit, operator-split formalism. The inner linear Newton systems are solved using a Schur complement formulation, which neatly decouples the local microphysics from the transport calculation. Radiation transport is modeled in the flux-limited diffusion approximation for simplicity, although our approach can be easily generalized to higher-order moment schemes. The use of Schur complements allows the application of optimal and scalable multigrid methods for the solution of the scalar radiation diffusion equation. We describe our algorithm in detail. We then illustrate our method's accuracy, robustness, and parallel scalability against a suite of verification tests of increasing size and complexity. In its current implementation, we are restricted to calculations on uniform Cartesian grids. An extension of our algorithm on block structured adaptive meshes is under development, and requires only modifications to the inner multigrid linear solver.

The paper is organized as follows. In Section 2 the governing equations for cosmological radiation hydrodynamics are presented under two different assumptions about the radiation-matter coupling: a two temperature model assuming local thermodynamic equilibrium (LTE), and a non-LTE ionization kinetics multispecies model. In Section 3 we describe our solution procedures for splitting off the hydrodynamic calculation, and our coupled implicit radiation-ionization-gas energy system. Results from solution verification tests are presented in Section 4, as well as parallel scalability tests. Conclusions follow in Section 5.

2. Cosmological radiation-hydrodynamics-ionization model

We consider the coupled system of partial differential equations

$$\partial_t \rho_b + \frac{1}{a} \mathbf{v}_b \cdot \nabla \rho_b = -\frac{1}{a} \rho_b \nabla \cdot \mathbf{v}_b, \quad (1)$$

$$\partial_t \mathbf{v}_b + \frac{1}{a} \mathbf{v}_b \cdot \nabla \mathbf{v}_b = -\frac{\dot{a}}{a} \mathbf{v}_b - \frac{1}{a \rho_b} \nabla p - \frac{1}{a} \nabla \phi, \quad (2)$$

$$\partial_t e + \frac{1}{a} \mathbf{v}_b \cdot \nabla e = -\frac{2\dot{a}}{a} e - \frac{1}{a \rho_b} \nabla \cdot (p \mathbf{v}_b) - \frac{1}{a} \mathbf{v}_b \cdot \nabla \phi + G - A, \quad (3)$$

$$\partial_t n_i + \frac{1}{a} \nabla \cdot n_i \mathbf{v}_b = \alpha_{ij} n_e n_j - n_i \Gamma_i^{ph}, \quad i = 1, \dots, N_s, \quad (4)$$

$$\partial_t E + \frac{1}{a} \nabla \cdot E \mathbf{v}_b = \nabla \cdot D \nabla E - m \frac{\dot{a}}{a} E + 4\pi \eta - c \kappa E. \quad (5)$$

Here, the first three Eqs. (1)–(3) correspond to the equations of ideal-gas dynamics in a coordinate system that is comoving with the expanding universe [14]. These equations correspond to mass, momentum and energy conservation, respectively, in which $\mathbf{v}_b \equiv a(t) \dot{\mathbf{x}}$ is the proper peculiar baryonic velocity, p is the proper pressure, and the total gas energy per unit mass is given by e . The modified gravitational potential is given by ϕ , which satisfies the comoving form of Poisson's equation,

$$\nabla^2 \phi = \frac{4\pi g}{a} (\rho_b + \rho_{dm} - \langle \rho \rangle), \quad (6)$$

where g provides the gravitational constant, ρ_b and ρ_{dm} are the baryonic and dark matter densities, respectively, and $\langle \rho \rangle$ is the cosmic mean density. The densities ρ_i are comoving, relating to the proper densities through the relation $\rho_i \equiv \rho_{i, \text{proper}} a(t)^3$. Here $a(t) \equiv (1+z)^{-1}$ denotes the cosmological expansion parameter for a smooth homogeneous background,

where the redshift z is a function of time only. All spatial derivatives are taken with respect to the comoving position $\mathbf{x} \equiv \mathbf{r}/a(t)$. The hydrodynamics equations are closed as usual with the ideal-gas equation of state,

$$e = \frac{p}{\rho_b(\gamma - 1)} + \frac{1}{2}|\mathbf{v}_b|^2, \quad (7)$$

where γ is the ratio of specific heats, taken to be 5/3 throughout this work.

The hydrodynamics equations are coupled with the elemental species rate Eq. (4) and an equation describing the flux-limited diffusion (FLD) approximation of radiation transport in a cosmological medium (5) [12,15]. In these equations i denotes the i th chemical species (out of N_s total), n_i is the comoving number density, n_e is the electron number density, n_j corresponds to other ions that react with the species i , and α_{ij} are the rate coefficients defining these interactions [16,17]; E corresponds to the comoving radiation energy density. The parameter m controls whether E is monochromatic at a specified frequency ν ($m = 0$), or an integrated grey radiation energy density ($m = 1$).

The baryonic gas is coupled to collisionless dark matter solely through their self-consistent gravitational field via (6). The dark matter density is evolved using the Particle-Mesh method described in [18–20]. As the N-body method is standard and not the focus of this paper, we do not elaborate on it here.

2.1. Model coupling

In addition to the advective components of the chemistry and radiation equations, coupling between these equations arise through a number of spatially-local reaction terms. The radiation energy density and chemical number densities affect the gas energy through the heating and cooling rates G and \mathcal{A} , respectively. The ionization and recombination rates α_{ij} and emissivity η depend on the gas temperature,

$$T = (\gamma - 1) \frac{\rho \mu m_p}{\rho_b k_b}, \quad (8)$$

where m_p corresponds to the mass of a proton, μ corresponds to the local molecular weight, and k_b is Boltzmann's constant. Finally, the photoionization rate Γ_i^{ph} depends on the radiation energy density, and the opacity κ depends on the state of chemical ionization.

In determining these coupling terms we distinguish between two cases: those in local thermodynamic equilibrium (LTE) and those that are not (nLTE). In the LTE case the chemical species are assumed to be in equilibrium, and hence their Eq. (4) may be omitted from the time-dependent system (1)–(5). For problems in this regime, the coupling terms resemble those typically encountered in radiation-hydrodynamics simulations [12,21,22],

$$\eta_{LTE}(T) = \kappa_p B = \frac{\kappa_p \sigma_{SB}}{\pi} T^4, \quad G_{LTE}(\rho_b, E) = \frac{c\kappa}{\rho_b} E, \quad \mathcal{A}_{LTE}(\rho_b, T) = \frac{4\pi}{\rho_b} \eta_{LTE}(T), \quad (9)$$

where c is the speed of light, σ_{SB} is the Stefan–Boltzmann constant, and κ_p and κ correspond to the problem-dependent Planck mean and total opacities for the gas.

For simulations that may not be approximated as being in local thermodynamic equilibrium, these coupling terms involve the dynamically-changing chemical ionization states. Here, the combined opacity depends on the local ionization states n_i , the emissivity η depends on both T and n_i , the gas heating rate G depends on both E and n_i , and the gas cooling rate \mathcal{A} depends on T and n_i , with the corresponding formulas given in the references [16,23–26].

2.2. Cosmological flux-limited radiation diffusion model details

We derive the cosmological flux-limited radiation diffusion equation (5) from the general multi-frequency version [15],

$$\partial_t E_\nu + \frac{1}{a} \nabla \cdot (E_\nu \mathbf{v}_b) = \nabla \cdot (D \nabla E_\nu) + \nu \frac{\dot{a}}{a} \partial_\nu E_\nu + 4\pi \eta_\nu - c\kappa_\nu E_\nu. \quad (10)$$

Through assumption of a given radiation frequency spectrum, $\chi_E(\nu)$, the frequency-dependent radiation energy density may be written in the form $E_\nu(\mathbf{x}, t, \nu) = \tilde{E}(\mathbf{x}, t) \chi_E(\nu)$. With this, we define the single “grey” radiation energy density used in the model (1)–(5) as

$$E(\mathbf{x}, t) = \int_{\nu_0}^{\infty} E_\nu(\mathbf{x}, t, \nu) d\nu = \tilde{E}(\mathbf{x}, t) \int_{\nu_0}^{\infty} \chi_E(\nu) d\nu. \quad (11)$$

The radiation Eq. (5) may then be derived through integration of the Eq. (10) over frequencies ranging from the ionization threshold of Hydrogen ($h\nu_0 = 13.6$ eV) to infinity; integration of the term $\nu \frac{\dot{a}}{a} \partial_\nu E_\nu$ gives rise to the term $-\frac{\dot{a}}{a}$ in (5). We note that this approximation (11) is valid only if the assumed spectrum $\chi_E(\nu)$ is defined such that the indefinite integral exists, as is the case for quasar and stellar type spectra where it scales with frequency as $E_\nu \propto \nu^{-\beta_q}$ where $\beta_q > 1$. However, through this formulation we may also consider problems involving a monochromatic radiation energy density, since such energy densities may also be expanded as $E_\nu(\mathbf{x}, t, \nu) = \tilde{E}(\mathbf{x}, t) \chi_E(\nu)$, where for radiation at the monochromatic frequency ν_k , the assumed spectrum is given through the Dirac-delta function $\chi_E(\nu) = \delta_{\nu_k}(\nu)$. In such cases, the term $\nu \frac{\dot{a}}{a} \partial_\nu E_\nu$ vanishes, giving rise to the

parameter m in (5). For standard grey radiation approximations, we assume a radiation spectrum of the form of either a power law, $\chi_E(\nu) = \left(\frac{\nu}{\nu_0}\right)^\alpha$, $\alpha < -1$, or as a T_B blackbody spectrum, $\chi_E(\nu) = 8\pi h \left(\frac{\nu}{c}\right)^3 / \left(\exp\left(\frac{h\nu}{k_B T_B}\right) - 1\right)$.

As is standard with FLD approximations to radiation transfer, one must take special care in construction of the diffusion coefficient function D . In its simplest form, D may be written as $D = \frac{c}{3\kappa_T}$, where $\kappa_T = \kappa_A + \kappa_S$ is the *total extinction coefficient*, κ_A corresponds to total absorption (here taken to be the opacity κ) and κ_S corresponds to scattering [21]. Use of this form for the diffusion coefficient, however, results in an infinite signal speed for the radiative flux $-D\nabla E$. To preserve causality, the analytic form of D is modified with a dimensionless *flux-limiter* whose particular form may be tuned to the specific problem of interest, but whose overriding purpose is to guarantee that the radiation transfer Eq. (5) gives the correct numerical behavior in the limiting cases of (nearly) isotropic and free-streaming radiation. Several choices for flux-limited forms of D have been proposed in the literature [27,28]. We consider the diffusion coefficient to be of the form

$$D(E) = \begin{bmatrix} D_1(E) & 0 & 0 \\ 0 & D_2(E) & 0 \\ 0 & 0 & D_3(E) \end{bmatrix}, \quad (12)$$

where

$$D_i(E) = \frac{c(2\kappa_T + R_i)}{6\kappa_T^2 + 3\kappa_T R_i + R_i^2},$$

with $R_i = |\partial_i E|/E$, $i = 1, 2, 3$. We note that this function has been reformulated from its original version [12] to allow increased numerical precision for scattering-free simulations involving extremely small opacities (i.e. $\kappa_T = \kappa_A = \kappa \ll 1$), as is typical in cosmology applications.

3. Solving the coupled system

3.1. Operator-split hydrodynamics with radiative feedback

Since typical astrophysical and cosmological simulations involve the hydrodynamic motion of gases encountering shocks, whereas radiation diffusion and chemical kinetics processes are more of reaction–diffusion type, we choose to solve the coupled system (1)–(5) in an operator-split fashion. In this approach, a time step t^n to t^{n+1} is taken using the general steps

- (i) Deposit the dark matter particles onto the mesh to calculate the ρ_{dm}^n .
- (ii) Solve for the gravitational potential ϕ resulting from the densities ρ_b and ρ_{dm} using equation (6).
- (iii) Evolve the dark matter particles using the Particle-Mesh Method [18–20].
- (iv) Evolve the hydrodynamics Eqs. (1)–(3) with a high-order, explicit-time upwind method. In this step, use the velocity field \mathbf{v}_b to advect both the chemical number densities n_i and radiation energy density E .
- (v) Using a high-order implicit-time method, solve a coupled reaction–diffusion system to obtain the time-evolved number densities n_i , radiation energy density E and gas energy e .

In order to allow us to split the Eqs. (1)–(5) into the two steps (iv) and (v) above, we consider the gas energy as consisting of two components, $e = e_h + e_c$, where e_h is the fluid energy arising from the hydrodynamic evolution of the system, and e_c is the gas energy *correction* arising from the couplings with radiation and chemistry. Under this decomposition, the energy conservation equation (3) may be equivalently written as

$$\partial_t(e_h + e_c) + \frac{1}{a} \mathbf{v}_b \cdot \nabla(e_h + e_c) = -\frac{2\dot{a}}{a}(e_h + e_c) - \frac{1}{a\rho_b} \nabla \cdot (p\mathbf{v}_b) - \frac{1}{a} \mathbf{v}_b \cdot \nabla\phi + G - \Lambda. \quad (13)$$

Under this splitting, the hydrodynamic solver used in step (iv) of the operator-split algorithm solves the system of equations

$$\partial_t \rho_b + \frac{1}{a} \mathbf{v}_b \cdot \nabla \rho_b = -\frac{1}{a} \rho_b \nabla \cdot \mathbf{v}_b, \quad (14)$$

$$\partial_t \mathbf{v}_b + \frac{1}{a} \mathbf{v}_b \cdot \nabla \mathbf{v}_b = -\frac{\dot{a}}{a} \mathbf{v}_b - \frac{1}{a\rho_b} \nabla p - \frac{1}{a} \nabla \phi, \quad (15)$$

$$\partial_t e_h + \frac{1}{a} \mathbf{v}_b \cdot \nabla e_h = -\frac{2\dot{a}}{a} e_h - \frac{1}{a\rho_b} \nabla \cdot (p\mathbf{v}_b) - \frac{1}{a} \mathbf{v}_b \cdot \nabla \phi \quad (16)$$

$$\partial_t n_i + \frac{1}{a} \nabla \cdot n_i \mathbf{v}_b = 0, \quad (17)$$

$$\partial_t E + \frac{1}{a} \nabla \cdot E \mathbf{v}_b = 0, \quad (18)$$

to evolve the solution at t^n , $(\rho_b^n, \mathbf{v}_b^n, e^n, n_i^n, E^n)$, to the time-updated variables at t^{n+1} , $(\rho_b^{n+1}, \mathbf{v}_b^{n+1}, e_h^{n+1})$, and the advected variables (n_i^*, E^*) . For this step, we employ the Piecewise Parabolic Method (PPM) [29], on a regular finite-volume spatial grid, implemented in the community astrophysics code Enzo [2,20,30].

The remainder of the coupled system,

$$\partial_t e_c = -\frac{2\dot{a}}{a} e_c + G - A, \quad (19)$$

$$\partial_t n_i = \alpha_{ij} n_e n_j - n_i \Gamma_i^{ph}, \quad (20)$$

$$\partial_t E = \nabla \cdot D \nabla E - m \frac{\dot{a}}{a} E + 4\pi\eta - ckE, \quad (21)$$

is then solved using a fully-implicit nonlinear solution approach to evolve the advected variables $(0, n_i^*, E^*)$ to the time-evolved quantities $(e_c^{n+1}, n_i^{n+1}, E^{n+1})$. Here we may assume that $e_c^n = 0$ since the hydrodynamic solver uses the full energy at t^n in its evolution, i.e. $e_h^n = e^n$. Once this step is finished, we compute the time-evolved total energy as the sum of the hydrodynamic portion e_h and the adjustments due to radiative feedback e_c , i.e. $e^{n+1} = e_h^{n+1} + e_c^{n+1}$. The treatment of the implicit radiation, chemical ionization, and gas energy feedback system (19)–(21) serves as our focus for the remainder of this section.

3.2. Solving the radiation, ionization and energy feedback system

Under a method-of-lines approach, we consider a two level, up-to-second-order accurate theta-scheme for implicit integration of our system (19)–(21),

$$e_c^{n+1} + \Delta t \theta \mathcal{L}_e^{n+1} = e_c^n + \Delta t (\theta - 1) \mathcal{L}_e^n, \quad (22)$$

$$n_i^{n+1} + \Delta t \theta \mathcal{L}_{n_i}^{n+1} = n_i^n + \Delta t (\theta - 1) \mathcal{L}_{n_i}^n, \quad (23)$$

$$E^{n+1} + \Delta t \theta [\mathcal{D}_E^{n+1} + \mathcal{L}_E^{n+1}] = E^n + \Delta t (\theta - 1) [\mathcal{D}_E^n + \mathcal{L}_E^n]. \quad (24)$$

Here, the parameter θ defines the implicit integration method: $\theta = 1$ corresponds to a first-order implicit Euler method, $\theta = 0.5$ gives a second-order time-centered approach (i.e. Crank–Nicolson). We note that in the ensuing computational results from Section 4, we have typically taken $\theta = 0.51$ to provide a nearly-second-order time integration while avoiding the “ringing” traditionally associated with fully time-centered approaches [31,32]. For the above equations, we have defined the diffusive operator

$$\mathcal{D}_E = \mathcal{D}_E(E, n_i) \equiv -\nabla \cdot D \nabla E, \quad (25)$$

and we have defined the local “reaction” operators as

$$\mathcal{L}_e = \mathcal{L}_e(e_c, E, n_i) \equiv \frac{2\dot{a}}{a} e_c - G + A \quad (26)$$

$$\mathcal{L}_{n_i} = \mathcal{L}_{n_i}(n_i, e_c, E) \equiv n_i \Gamma_i^{ph} - \alpha_{ij} n_e n_j \quad (27)$$

$$\mathcal{L}_E = \mathcal{L}_E(E, e_c, n_i) \equiv m \frac{\dot{a}}{a} E - 4\pi\eta + ckE. \quad (28)$$

The Eqs. (22)–(24) form a coupled nonlinear system of reaction–diffusion equations for evolution of the fluid energy correction e_c , the elemental number densities n_i , and the radiation energy density E . Denoting the vector of unknowns $U = (e_c, n_i, E)^T$, we first define the nonlinear residual function for the time step $t^n \rightarrow t^{n+1}$, as

$$f(U) \equiv U + \Delta t \theta \begin{pmatrix} \mathcal{L}_e(U) \\ \mathcal{L}_{n_i}(U) \\ \mathcal{D}_E(U) + \mathcal{L}_E(U) \end{pmatrix} - \begin{pmatrix} g_{e_c}^n \\ g_{n_i}^n \\ g_E^n \end{pmatrix}, \quad (29)$$

where the vectors g_*^n are formed using the previous time-level information from (22)–(24). In order to evolve the coupled implicit system, we solve the nonlinear problem $f(U) = 0$ for the updated vector of unknowns U^{n+1} . For this nonlinear solve, we use a *globalized Inexact Newton's Method* [33,34], in which we apply an iterative process for convergence toward the solution U^{n+1} in the following manner.

Given an initial guess $U_0 \approx U(t^{n+1})$, we iterate toward a solution U^{n+1} satisfying $\|f(U^{n+1})\| < \varepsilon \ll 1$ (we typically choose $\varepsilon = 10^{-7}$):

1. Approximately solve the linearized Newton system, $\|J(U_k)S_k + f(U_k)\| < \delta_k$, to tolerance δ_k for the correction vector S_k . Here, $J(U_k) \equiv \frac{\partial}{\partial U} f(U_k)$, and we typically choose the tolerance as $\delta_k = 10^{-6} \|f(U_k)\|$.
2. Update the vector of unknowns as $U_{k+1} = U_k + \lambda_k S_k$, where $\lambda_k \in (\lambda_{min}, 1]$ is the *line-search parameter* [35,36].

We measure convergence of the Newton iteration with the RMS norm

$$\|v\| = \left(\frac{\|v\|_2^2}{N(N_s + 2)} \right)^{1/2}, \tag{30}$$

where $N(N_s + 2)$ is the number of unknowns in v (N spatial cells, $N_s + 2$ variables), since such a norm does not grow artificially larger with mesh refinement. The key to efficiency of the inexact Newton algorithm lies in a fast and robust solver for the linear systems $JS = -f$. Once such a solver has been provided, the algorithm exhibits very fast convergence – superlinear for this choice of δ_k [33,37]. Moreover, for diffusive PDE systems similar to the one solved here, the Newton convergence rate has been shown to be *independent* of spatial resolution [38], suggesting that this entire implicit algorithm should allow scalability to the limits of the inner linear solver.

3.2.1. Linear solver

In solving the system (29), we make one approximation within the Newton system matrices $J(U)$, wherein we lag the n_i dependence of \mathcal{D}_E in (25) to the previous Newton iterate. Mathematically, this results in a full Newton step for all but the limiter's dependence on the chemical opacities, which are instead converged through a fixed-point iteration. The resulting solution retains the accuracy and stability of the full Newton iteration, albeit with theoretically slower convergence. However, in practice we have not noticed any increase in nonlinear iterations due to this approximation, and most importantly it results in inexact Newton matrices with the form

$$J(U) = I + \Delta t \theta \begin{bmatrix} J_{e,e} & J_{e,n} & J_{e,E} \\ J_{n,e} & J_{n,n} & J_{n,E} \\ J_{E,e} & J_{E,n} & J_{E,E} \end{bmatrix}, \tag{31}$$

where nearly all of the blocks are given by the spatially-local components,

$$\begin{aligned} J_{e,e} &\equiv [\partial_e \mathcal{L}_e] & J_{e,n} &\equiv [\partial_{n_1} \mathcal{L}_e \quad \partial_{n_2} \mathcal{L}_e \quad \dots] & J_{e,E} &\equiv [\partial_E \mathcal{L}_e], \\ J_{n,e} &\equiv \begin{bmatrix} \partial_e \mathcal{L}_{n_1} \\ \partial_e \mathcal{L}_{n_2} \\ \vdots \end{bmatrix} & J_{n,n} &\equiv \begin{bmatrix} \partial_{n_1} \mathcal{L}_{n_1} & \partial_{n_2} \mathcal{L}_{n_1} & \dots \\ \partial_{n_1} \mathcal{L}_{n_2} & \partial_{n_2} \mathcal{L}_{n_2} & \dots \\ \vdots & \vdots & \ddots \end{bmatrix} & J_{n,E} &\equiv \begin{bmatrix} \partial_E \mathcal{L}_{n_1} \\ \partial_E \mathcal{L}_{n_2} \\ \vdots \end{bmatrix}, \\ J_{E,e} &\equiv [\partial_e \mathcal{L}_E] & J_{E,n} &\equiv [\partial_{n_1} \mathcal{L}_E \quad \partial_{n_2} \mathcal{L}_E \quad \dots], \end{aligned} \tag{32}$$

and the only block containing spatial couplings is $J_{E,E} \equiv [\partial_E(\mathcal{D}_E + \mathcal{L}_E)]$. Thus, although the Jacobian matrix contains couplings both within and between variables, it has a very desirable structure: all inter-variable couplings occur locally in space, and the only nonlocal couplings are within the block $J_{E,E}$, consisting of a scalar-valued reaction–diffusion operator.

In keeping with a block-structured view of the Jacobian (31), we rewrite the Newton system $JS = -f$ in the form

$$\begin{bmatrix} M & U \\ L & D \end{bmatrix} \begin{pmatrix} s_M \\ s_E \end{pmatrix} = - \begin{pmatrix} f_M \\ f_E \end{pmatrix}, \tag{33}$$

where

$$\begin{aligned} M &= I + \Delta t \theta \begin{bmatrix} J_{e,e} & J_{e,n} \\ J_{n,e} & J_{n,n} \end{bmatrix}, & U &= \Delta t \theta \begin{bmatrix} J_{e,E} \\ J_{n,E} \end{bmatrix}, \\ L &= \Delta t \theta [J_{E,e} \quad J_{E,n}], & D &= I + \Delta t \theta [J_{E,E}], \end{aligned} \tag{34}$$

$s_M = [s_e, s_n]^T$ and $f_M = [f_e, f_n]^T$. We note that the only matrix containing spatial dependencies is D , so under an appropriate variable ordering the other sub-matrices are block diagonal. Hence, we may efficiently invert M to obtain s_M as a function of s_E :

$$Ms_M + Us_E = -f_M \Rightarrow s_M = -M^{-1}(f_M + Us_E).$$

Inserting this into the second row, we have the single equation for s_E ,

$$(D - LM^{-1}U)s_E = -f_E + LM^{-1}f_M.$$

Therefore, this *Schur complement* formulation [39] for solution of the linear Newton system (33) proceeds with the following steps:

- (i) Set $\tilde{f}_M = M^{-1}f_M$.
- (ii) Solve the system $(D - LM^{-1}U)s_E = -f_E + L\tilde{f}_M$ for s_E .
- (iii) Recover the remaining solution pieces, $s_M = -\tilde{f}_M - M^{-1}Us_E$.

We examine each of these steps below.

The step (i) corresponds to solving the linear system

$$\begin{bmatrix} I + \Delta t \theta J_{e,e} & \Delta t \theta J_{e,n} \\ \Delta t \theta J_{n,e} & I + \Delta t \theta J_{n,n} \end{bmatrix} \begin{pmatrix} \tilde{f}_e \\ \tilde{f}_n \end{pmatrix} = \begin{pmatrix} f_e \\ f_n \end{pmatrix}. \quad (35)$$

Due to the spatial locality of each component in M , we order the equations and unknowns in this system so that application of M , and more notably M^{-1} , may be performed independently in every spatial cell. Such solves consist of dense matrix algebra on $(1 + N_s) \times (1 + N_s)$ linear systems (for N_s chemical densities). In addition, at this step we compute the matrix $M^{-1}U$ through one additional solve with M , which will be used in the following steps.

The step (ii) corresponds to solving the system $(D - LM^{-1}U)s_E = L\tilde{f}_M - f_E$. This is denoted the *Schur complement* system, with the matrix $S = D - LM^{-1}U$. We note that due to the spatially-local nature of the matrices L and $M^{-1}U$, we may form S by constructing the diffusive sub-matrix D , followed by updates to the diagonal entries corresponding to the entries of $LM^{-1}U$. Similarly, construction of the right-hand side $L\tilde{f}_M - f_E$ may occur independently at each spatial location. Once this system has been computed, we use a multigrid-preconditioned conjugate gradient parallel linear solver from the HYPRE library [40,41] to perform the scalar-valued solve, $s_E = S^{-1}(L\tilde{f}_M - f_E)$. We note that this is the only step in the solution of the Jacobian systems (33) that requires communication between processors. Moreover, we point out that in recent tests the HYPRE library has demonstrated ideal weak scaling up to over 100,000 processors for diffusion problems similar to the one encountered in this work [42]. As this solver comprises the majority of the non-local components within our nonlinear solver, we therefore expect similar scalability for the overall implicit solution approach described here.

The final step (iii) in solution of the system (33) is to recover the solution components $s_M = (s_e, s_n)^T$ via the system $s_M = -\tilde{f}_M - M^{-1}Us_E$. Again, since we have already computed the spatially-local matrix $M^{-1}U$ and the vector $\tilde{f}_M = M^{-1}f_M$ in step (i), we may trivially obtain the remaining solution components through cell-local matrix-vector products and vector operations, $s_M = -\tilde{f}_M - (M^{-1}U)s_E$.

3.2.2. Multiphysics/cosmology units

As with any multi-physics system, special care must be taken when solving such systems computationally due to disparate scales between variables and equations. This problem is especially evident in cosmology applications, where in CGS units one may typically encounter specific gas energies on the order of 10^{12} , number densities on the order of 10^{-27} , and radiation energy densities on the order of 10^{-15} , with all proper density values decreasing in time due to cosmological expansion. To this end, we define the scaled variables

$$\tilde{e}_c = e_c/u_e, \quad \tilde{E} = E/u_E, \quad \tilde{n}_i = n_i/u_n, \quad \tilde{x} = x/u_x, \quad \tilde{t} = t/u_t, \quad (36)$$

where the constants u_e, u_E, u_n, u_x and u_t correspond to the typical magnitudes of gas energy, radiation energy density, chemical number density, length and time at the start of the simulation. We note that due to our use of *comoving* values for E, n_i and x , these constants are all redshift-independent, with the *proper* values of these quantities given by

$$E_{\text{proper}} = E/a^3(t) = \tilde{E} \frac{u_E}{a^3(t)}, \quad n_{i,\text{proper}} = n_i/a^3(t) = \tilde{n}_i \frac{u_n}{a^3(t)}, \quad x_{\text{proper}} = xa(t) = \tilde{x}u_x a(t). \quad (37)$$

The constants are supplied on a problem-dependent basis, to allow for adaptable, on-the-fly rescaling of simulations ranging from normalized test problems to cosmological reionization. With these rescaled variables, we rewrite our Eqs. (22)–(24) as the normalized system

$$\tilde{e}_c^{n+1} + \Delta \tilde{t} \theta \tilde{\mathcal{L}}_e^{n+1} = \tilde{e}_c^n + \Delta \tilde{t} (\theta - 1) \tilde{\mathcal{L}}_e^n, \quad (38)$$

$$\tilde{n}_i^{n+1} + \Delta \tilde{t} \theta \tilde{\mathcal{L}}_n^{n+1} = \tilde{n}_i^n + \Delta \tilde{t} (\theta - 1) \tilde{\mathcal{L}}_n^n, \quad (39)$$

$$\tilde{E}^{n+1} + \Delta \tilde{t} \theta [\tilde{\mathcal{D}}_E^{n+1} + \tilde{\mathcal{L}}_E^{n+1}] = \tilde{E}^n + \Delta \tilde{t} (\theta - 1) [\tilde{\mathcal{D}}_E^n + \tilde{\mathcal{L}}_E^n]. \quad (40)$$

Here, the operators $\tilde{\mathcal{L}}_e, \tilde{\mathcal{L}}_n, \tilde{\mathcal{L}}_E$ and $\tilde{\mathcal{D}}_E$ have correspondingly absorbed the renormalization constants u_i . These equations, along with the normalized solution vector $\tilde{U} = (\tilde{e}_c, \tilde{n}_i, \tilde{E})^T$ are then used within the solution strategy described in Section 3.2.

3.2.3. Adaptive time step selection

A strong appeal of using implicit methods is their stability with respect to time step size; however such freedom gives rise to the question of what time step should be used. At one extreme, we may choose a large step to achieve overall efficiency of the simulation, with little to no knowledge of the resulting temporal accuracy. At the other extreme, we may choose a very small time step for an accurate solution, resulting in inefficient simulations due to the increased cost of solving the nonlinear systems at each step. As the approach described here is operator-split, in which the hydrodynamics is solved using an explicit approach, we are therefore bound by the hydrodynamic CFL stability limit; however for most problems involving radiation and chemical ionization, the dynamic time scales of interest remain significantly faster than the hydrodynamic time scale. Thus the question of how to adaptively choose the time step size remains.

To that end, we adaptively choose the time steps as the largest possible that additionally satisfy a prescribed accuracy requirement. We estimate this accuracy through comparison of the updated solution U^{n+1} with an explicit predictor for that solution U^{pred} . Defining the weighting vector in a spatial cell i for the variable v as

$$\omega_{i,v} = \sqrt{|U_{i,v}^{n+1} U_{i,v}^{pred}|} + 1, \quad (41)$$

(which assumes normalized values of U_v), we estimate the local accuracy of the current time step as

$$\varepsilon_{loc} = \left(\frac{1}{N(N_s + 2)} \left\| \frac{U^{n+1} - U^{pred}}{\omega} \right\|_p^p \right)^{1/p}, \quad (42)$$

where we have used the standard p -norm (including $p = \infty$ as the ‘max’ norm, in which case we do not divide by $N(N_s + 2)$), and where the quotient inside the norm is taken pointwise. With this estimate, we set the new time step to

$$\Delta t^{n+1} = \frac{\tau_{tol} \Delta t^n}{\varepsilon_{loc}}, \quad (43)$$

which should provide the maximal value that still satisfies the desired integration accuracy tolerance, τ_{tol} , assuming that U^{pred} approximates the time-evolved solution U^{n+1} to $O(\Delta t)$.

Here, the vector ω is designed so that ε_{loc} estimates the average relative change in each solution component, and includes the harmonic mean of the predicted and new states to allow increased robustness in the case of cosmology-type problems where variables change by orders of magnitude across cells and time steps. The value of p is typically taken to be ∞ in the ensuing test problems; however such a choice may limit parallel scaling since such a measure is sensitive to pointwise changes, of which there are many more as dynamics propagate throughout an increasingly refined domain. Lastly, we use the explicit predictor as the initial guess for the Newton method, U_0 , which we describe in the following Section 3.2.4.

3.2.4. Explicit predictor

A well known property of Newton’s method is that its robustness and efficiency benefit greatly from an accurate initial guess. To this end, we provide the predicted initial Newton iterate

$$U_0 = \begin{pmatrix} e_{c,0} \\ n_{i,0} \\ E_0 \end{pmatrix} = \begin{pmatrix} e_c^n \\ n_i^n \\ E^n \end{pmatrix} + \Delta t \begin{pmatrix} \mathcal{L}_e^n \\ \mathcal{L}_{n_i}^n \\ \mathcal{D}_E^n + \mathcal{L}_E^n \end{pmatrix}, \quad (44)$$

i.e. we use an initial guess given by the $O(\Delta t)$ -accurate explicit Euler update to the coupled system (19)–(21). As this provides only an initial guess to the solution, its instability at larger Δt will not affect the temporal stability of the overall method, since the solution to each step must satisfy the implicit system (22)–(24). However, as we use an adaptive time-stepping strategy, for very fast dynamics (that give rise to very small Δt), such an initial guess may already satisfy the nonlinear tolerance $\|f(U_0)\| < \varepsilon$ and the solver will not require any Newton iterations, effectively allowing an adaptive explicit/implicit simulation of the coupled system (19)–(21).

3.2.5. Adaptive computation with supplied radiation spectrum $\chi_E(v)$

The final detail that we describe in this solution approach relates to the choice of assumed radiation spectrum $\chi_E(v)$. As noted in Section 2.2, we may choose either a monochromatic or an integrated ‘‘grey’’ radiation equation, based on the choice of this assumed spectrum. This choice affects all terms involving the radiation energy density E_v in the general radiation energy equation (10) and in the coupling terms G and Γ_i^{ph} . As each of these terms involve a product of the form $f(v)E_v$, integration over v converts these to

$$\int_{v_0}^{\infty} f(v)E_v(\mathbf{x}, t, v) dv = \tilde{E}(\mathbf{x}, t) \int_{v_0}^{\infty} f(v)\chi_E(v) dv. \quad (45)$$

We therefore allow a user-defined functional form for $\chi_E(v)$, which we then numerically integrate to high accuracy upon initialization of the simulation, providing the relevant constants necessary to convert the v -dependent Eq. (10) to the monochromatic or grey integrated Eq. (5).

4. Numerical results

We present test problems designed to verify the accuracy of the radiation diffusion and chemical ionization modules in conjunction with hydrodynamical fluid motions. Since a number of distinct processes may compete for importance in a full simulation, we begin with simple tests that isolate various single components, and subsequently build upon those results with more sophisticated problems that couple additional physics. We begin with a radiation diffusion problem (Section 4.1) that exercises the diffusion term of the radiation Eq. (5) in the absence of energy, chemical, or hydrodynamic coupling; this test is followed by an examination of the matter-radiation coupling terms (Section 4.2) in an infinite uniform medium, a diffusion wave with material coupling (Section 4.3), and a non-equilibrium radiating shock problem (Section 4.4), all of which assume chemical equilibrium. Our attention then turns to problems including ionization, beginning with an HII ionization front propagating through a static, isothermal medium (Section 4.5), followed by a cosmological I-front propagation problem that exercises the cosmology terms and units (Section 4.6). We then consider a fully-coupled radiation-hydrodynamics-ionization calculation (Section 4.7). This section concludes with additional calculations (Section 4.8) demonstrating the parallel scalability of the radiation diffusion module, which of all components places the highest communication demands upon a domain-decomposed parallel calculation.

We note that for all test problems except (Sections 4.6 and 4.8), we use a non-cosmological problem (i.e. $z = 0$ and $a = 1$). In problems (Sections 4.6 and 4.8), the cosmological parameters are described therein.

4.1. Free-streaming radiation

Because the standard diffusion equation is parabolic, the associated signal speed of the diffusion variable is formally infinite. However, in reality radiation fronts propagate at speeds bounded by the speed of light in vacuum, so we modify the diffusion coefficient in our radiation energy Eq. (5) with a flux-limiter, as discussed in Section 2.2. Our first test problem verifies the correct action of this limiter by examining the propagation of a planar radiation front through a transparent medium. Radiation is assumed to propagate along the x -axis of our computational mesh; a Dirichlet boundary condition is imposed on the left boundary specifying an incident radiation energy density of 1.0 erg cm^{-3} . Physically, the expectation is that with a sufficiently small (but nonzero, due to numerical constraints) opacity, a sharp radiation front will move through the domain at the speed of light. The Planck and Rosseland mean opacities are assigned a constant value of 10^{-6} cm^{-1} , ensuring an essentially transparent medium. The spatially uniform initial value of the radiation energy density is $10^{-4} \text{ erg cm}^{-3}$.

The computational mesh has a domain length of 1.0 cm along the propagation direction of the light wave. We have run the problem for 8.3391 picoseconds, which is one quarter of the light-crossing time for this length. Fig. 1 shows a series of curves resulting from calculations at mesh sizes of 128, 256, 512, 1024, and 2048 zones along the x -axis. The dashed line indicates the expected location of the radiation front, ct , where c is the speed of light and t the evolution time. In the absence of the flux-limiter, the numerical curves would give the formal $t \rightarrow \infty$ solution for the diffusion equation, which for our problem parameters would be a nearly horizontal profile for E throughout the domain (given the nearly zero opacity). That the

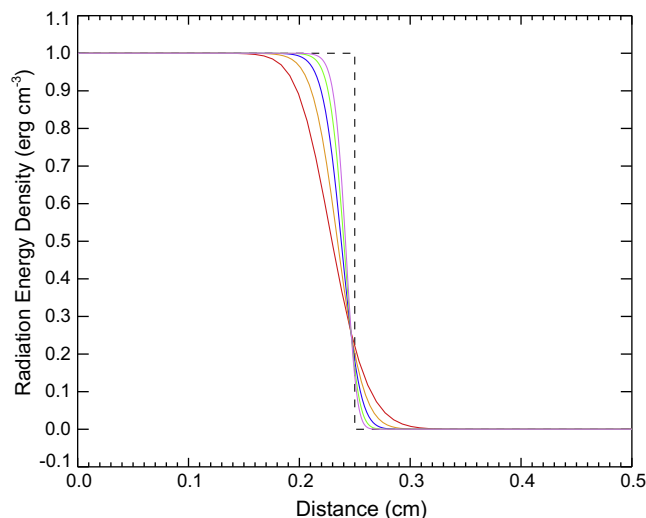


Fig. 1. Curves of E vs. x for mesh sizes of 128 (red), 256 (orange), 512 (blue), 1024 (green), and 2048 (violet) zones. The analytical solution (black dashed line) is a step function centered at $x = 0.25 \text{ m}$. (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

curves capture the correct location of the radiation front is due entirely to the action of the limiter. The sharpening of this front with increased resolution is evident.

Also apparent is a slight lag (about 0.01 cm) between the analytical location of the radiation front and the numerical location, taken as the common intersection point of the numerical curves. The size of this lag depends upon the choice of the adaptive time step. In the language of (Section 3.2.3), we compute Δt using $p = \infty$ and we vary τ_{tol} , which here corresponds to the maximum allowed fractional change in the radiation energy density per timestep. Fig. 2 illustrates this effect by showing 3 curve pairs, each of which has been computed with a different choice of τ_{tol} : 0.01 (green), 0.05 (blue), and 0.25 (red). In each colored pair, the solid curve was obtained using 512 zones, and the dashed curve shows the 128-zone result. Two curves at different mesh resolutions are provided to identify the “consensus” value of the light front location via their point of intersection. This location is seen to converge as $\tau_{\text{tol}} \rightarrow 0$.

4.2. Matter-radiation equilibration in a homogeneous medium

We now consider a problem in which e and E are spatially uniform but are initialized to values far away from equilibrium. This case thus isolates the matter-radiation coupling terms in the gas and radiation energy equations. The parameters for this test were published by Turner and Stone [43], who assumed an isotropic medium characterized by a single opacity of $4 \times 10^{-8} \text{ cm}^{-1}$, a gas density of $10^{-7} \text{ g cm}^{-3}$, and an average particle mass of $0.6 m_H$, where m_H is the mass of a hydrogen atom. Coupled to this medium is a radiation field with a uniform value of $10^{12} \text{ erg cm}^{-3}$. From this value we compute a “radiation temperature”, $T_r \equiv (E/a_r)^{1/4}$, of about $3.4 \times 10^6 \text{ K}$. Here we have defined a_r as the radiation constant, $7.56 \times 10^{-15} \text{ erg cm}^{-3} \text{ K}^{-1}$. Two cases are considered: one in which the initial gas energy density is $e = 10^{10} \text{ erg cm}^{-3}$, and one in which the initial value is $e = 10^2 \text{ erg cm}^{-3}$. For the stated parameters, these energies correspond to gas temperatures of roughly 4.8×10^8 and 4.8 K , respectively, which therefore bracket the radiation temperature. In both cases, however, the initial radiation temperature is sufficiently high that the radiation energy density should remain constant to good approximation as the gas evolves to thermal equilibrium. To see this clearly, consider the effective heat capacity of a unit volume of the radiation “gas” as compared to that for the material. For radiation, this number is simply $1.0 \text{ cm} \times a_r T_r^3$, which evaluates to roughly $3 \times 10^5 \text{ erg K}^{-1}$. In contrast, the gamma value and mean particle mass translate to a *specific heat* ($\text{erg g}^{-1} \text{ K}^{-1}$) of roughly 2.0×10^8 , which yields, for our assumed density, a material heat capacity of 20 erg K^{-1} for a unit control volume. The physical result is that the radiation field has an effectively infinite thermal reservoir when compared to the material.

If the radiation energy density is formally assumed to be constant, the gas energy equation may be written as a simple ODE:

$$\dot{e} = c\kappa E - 4\pi\kappa B(e), \quad (46)$$

where B is the temperature-dependent Planck function

$$B(T) = \frac{ca_r}{4\pi} T^4. \quad (47)$$

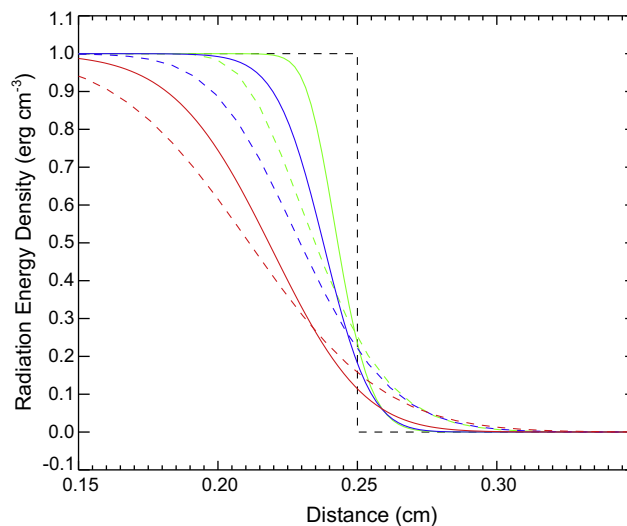


Fig. 2. Radiation energy profiles at 128 (dashed curves) and 512 (solid curves) zones for τ_{tol} of 0.01 (green), 0.05 (blue), and 0.25 (red). Intersections of associated 128 and 512-zone curves indicate the “consensus” location of the radiation front for a given choice of τ_{tol} . Convergence to the correct value is readily observed. (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

Using the ideal-gas law (7), we write B as a function of e and solve the simplified gas energy equation for the equilibrium value of e such that $\dot{e} \equiv 0$,

$$e_{\text{eq}} = \frac{3}{2} \left(\frac{\rho k_B}{0.6 m_H} \right) \left(\frac{E}{a_r} \right)^{1/4}. \tag{48}$$

Notice that this expression is nothing more than the ideal-gas formula for e evaluated at the fixed radiation temperature, the expected result.

The results of our two test calculations are shown in Fig. 3. Both tests were run in a small box domain (4^3 zones) with triply-periodic boundaries. The case of $T(0) > T_r$ is indicated by the solid curve; the low- T case is shown by the dashed line. The horizontal dotted line has been placed at $e = e_{\text{eq}}$. Both energy curves converge to the correct result. Note that in this test the opacity serves only to control the timescale to reach thermal equilibrium; neither the value of the equilibrium energy nor the validity of our assumption of constant E are dependent on the value of κ . While Fig. 3 demonstrates convergence to the correct asymptotic value of the gas energy, it provides no information as to whether the rate at which it approaches this value is correct. A quantitative assessment of this latter metric is provided by our next test problem.

We further use this test to examine the conservation properties of the coupled radiation and gas energy solver. In Table 1 we show the value of

$$\frac{\int |E_{\text{total}}(t) - E_{\text{total}}(0)| dx}{\int E_{\text{total}}(0) dx} \tag{49}$$

for both tests at the final time $t = 2.5e-7$ s, run using a variety of nonlinear and linear solver tolerances, ϵ and δ , respectively. We note that in all cases, the total energy is conserved to more than 10 digits of accuracy. Moreover, while the conservation is weakly dependent on the nonlinear solver tolerance, it is entirely independent of the linear solver tolerance. This behavior is most likely due to use of the Schur complement formulation (Section 3.2.1), that exactly solves for coupling between variables to floating point roundoff, leaving the iterative linear solver to handle only the radiation equation. We further note that this is an ideal problem to test conservation of the coupled solver, since it is the only test considered that uses a closed system. We further comment that since the PPM finite-volume method is constructed to satisfy conservation, and the implicit subsolver achieves conservation to high accuracy, overall conservation of the coupled solver follows. However, we note that for problems utilizing non-periodic boundary conditions, chemical ionization cooling, gravitational heating, or cosmological expansion, the model no longer represents a closed system and therefore will not conserve energy.

4.3. Non-equilibrium Marshak waves

This test exercises both radiation diffusion and the physics of matter-radiation coupling. Non-equilibrium Marshak waves characterize the evolution of the radiation field in an initially cold, uniform halfspace on which a radiation source is imposed. The particular form of the Marshak problem described here is originally due to Pomraning [44]. The problem was re-examined by Su and Olson [45], who derived semi-analytic exact solutions for the radiation and gas energy densities and tabulated select values of them on a grid of space and time values. The problem considers a formally 1D semi-infinite domain

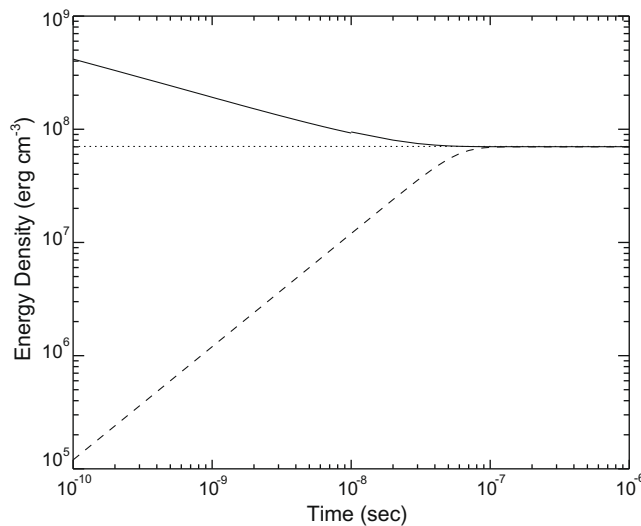


Fig. 3. Evolution to thermal equilibrium of a medium with an initially high (solid line) and low (dashed line) gas energy density. The analytical equilibrium value is shown by the horizontal dotted line.

Table 1

Conservation of total energy for the matter-radiation equilibration test. Relative error in energy conservation (49) for both low-to-high and high-to-low temperature equilibration, for various nonlinear and linear solver tolerances ε and δ , respectively.

ε	δ	$L \rightarrow H$ error	$H \rightarrow L$ error
10^{-7}	10^{-9}	7.04×10^{-12}	1.37×10^{-12}
10^{-10}	10^{-9}	3.07×10^{-13}	3.45×10^{-14}
10^{-4}	10^{-9}	3.55×10^{-11}	1.83×10^{-11}
10^{-7}	10^{-12}	7.04×10^{-12}	1.37×10^{-12}
10^{-7}	10^{-3}	7.04×10^{-12}	1.37×10^{-12}

in which (z, t) denote dimensional space and time coordinates. Ignoring hydrodynamic motions, Su and Olson write simplified forms for the radiation and material-energy equations as

$$\partial_t E(z, t) - \partial_z \left[\frac{c}{3\kappa} \partial_z E(z, t) \right] = c\kappa [a_r T^4(z, t) - E(z, t)], \quad (50)$$

$$c_v(T) \partial_t T(z, t) = c\kappa [E(z, t) - a_r T^4(z, t)], \quad (51)$$

in which κ is the constant opacity, a_r the radiation constant as defined previously, and c_v the specific heat of the material. Note that the flux-divergence term in (50) assumes pure diffusion with no flux-limiter. The matter temperature, T , is assumed to be related to the gas energy, e , via

$$e = \rho c_v T. \quad (52)$$

As written, (50) and (51) are coupled nonlinear PDEs in the dependent variables E and T . Su and Olson linearized the equations by choosing the following form for the specific heat:

$$c_v = \alpha T^3, \quad (53)$$

where α is an arbitrary constant. The T^3 dependence of c_v on temperature has two effects: it allows equations (50) and (51) to be written as linear ODEs in E and T^4 , and it gives the heat capacity of the material the same temperature dependence as the effective heat capacity of the radiation field. With an appropriate choice of α , a problem can therefore be designed in which the material and radiation will both evolve significantly in space and time.

The description of the problem is completed with a specification of the boundary conditions. A Marshak boundary condition is applied to E at $z = 0$:

$$E(0, t) - \frac{2}{3\kappa} \partial_z E(0, t) = \frac{4}{c} F_{\text{inc}}, \quad (54)$$

where F_{inc} is the incident flux at $z = 0$. The boundary condition at $z = \infty$, and initial conditions at $t = 0$ are

$$E(\infty, t) = 0, \quad E(z, 0) = T(z, 0) = 0. \quad (55)$$

Su and Olson construct the linearized equations by defining dimensionless independent and dependent variables, (X, τ) and (u, v) such that

$$X \equiv z\kappa\sqrt{3}, \quad u(X, \tau) \equiv \left(\frac{c}{4F_{\text{inc}}} \right) E(z, t), \quad \tau \equiv \left(\frac{4a_r c\kappa}{\alpha} \right) t, \quad v(X, \tau) \equiv \left(\frac{ca_r}{4F_{\text{inc}}} \right) T^4(z, t). \quad (56)$$

With these definitions, and letting $\epsilon = 4a_r/\alpha$, Eqs. (50), (51), and (54), (55) become

$$\epsilon \partial_\tau u(X, \tau) - \partial_X^2 u(X, \tau) = v(X, \tau) - u(X, \tau), \quad (57)$$

$$\partial_\tau v(X, \tau) = u(X, \tau) - v(X, \tau), \quad (58)$$

$$u(0, \tau) - \frac{2}{\sqrt{3}} \partial_X u(0, \tau) = 1, \quad (59)$$

$$u(\infty, \tau) = 0, \quad (60)$$

$$u(X, 0) = v(X, 0) = 0. \quad (61)$$

The Marshak boundary condition represented by (58) enforces the constraint of constant flux on the left boundary. This is an example of a “mixed” or Robin boundary condition, and as such requires special treatment in Enzo. For the purposes of this verification test, we implement this boundary condition by imposing a Dirichlet condition with a time-varying value of u computed from [45]’s Eq. (36), evaluated at $X = 0$. Because the integrands in their equation are highly oscillatory for $\tau \ll 1$, we substitute the asymptotic expression given by their Eq. (51) when $\tau < 10^{-5}$.

Fig. 4 shows results from a high-resolution simulation with 2048 zones along the X coordinate. The exact solution values tabulated by [45] span the range $0 \leq X \leq 10$. Since the right boundary condition is specified at $X = \infty$, we choose our domain $X \in [0, L]$ such that L is sufficiently large (about 35) for the evolution time of interest that the Dirichlet condition $X(L) = 0$ may be reasonably applied. We choose opacity and coupling parameters $\kappa = \epsilon = 1.0 \text{ cm}^{-1}$. The curves indicate profiles of u (dashed lines) and v (solid lines) for τ values of 1, 3, 10, 30, 100. The squares and circles indicate exact values of u and v , respectively. We have indicated these values on corresponding curves at evolution times sufficiently early that the material and radiation have not yet had time to equilibrate. Fig. 5 shows a resolution study for the curves computed at $\tau = 1$. Curves at mesh sizes of 128 (red), 256 (orange), 512 (blue), 1024 (green), and 2048 (black) are shown. Each calculation is performed with the timestep restriction $\tau_{\text{tol}} = 0.05$. Because this treatment allows for adaptive timesteps, the evident first-order rate of convergence measures the combined effect of time and space discretization methods.

4.4. Subcritical radiating shock waves

We now add hydrodynamic motions to our mix of physics by examining the propagation of shock waves for which the radiation energy plays a significant role in the shock structure and evolution. Radiating shock waves represent a broad class

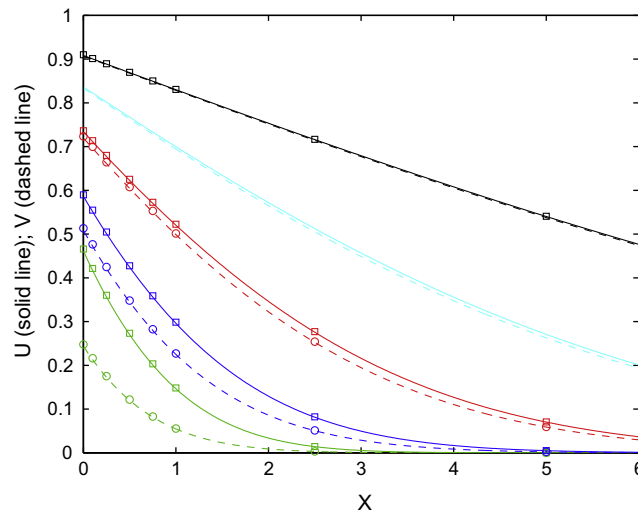


Fig. 4. Curves of e (dashed line) and E (solid line) vs. dimensionless distance, at dimensionless times $\tau = 1, 3, 10, 30$, and 100 (curves shift upward as τ increases). Analytic values for e (squares) and E (circles) are shown for $\tau = 1, 3$, and 10 ; values for E are also provided at $\tau = 100$, by which time e and E are nearly equilibrated.

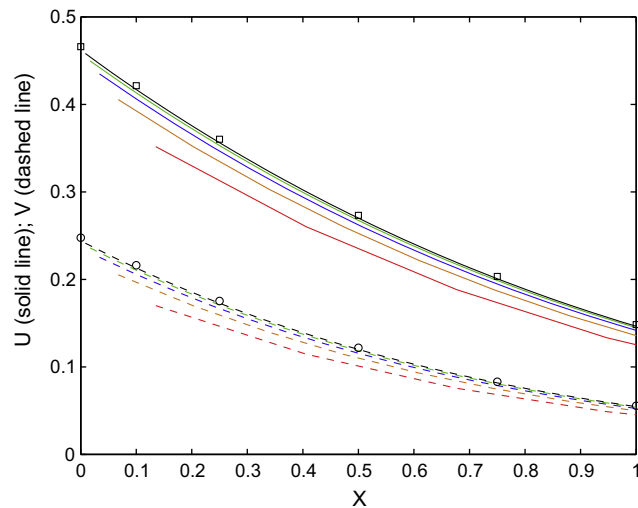


Fig. 5. Marshak problem convergence: curves of u (solid lines) and v (dashed lines) at mesh resolutions of 128 (red), 256 (orange), 512 (blue), 1024 (green) and 2048 (black) zones. Reference solution values are indicated by open circles and squares; the halving of relative error with each doubling of mesh size is readily apparent. (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

of phenomena figuring prominently in both astrophysical and terrestrial applications. The particular formulation of the problem we present is due to Lowrie and Edwards [46], who considered the propagation of planar, steady shock waves in the grey nonequilibrium diffusion limit. Under the assumption of steady flow, [46] transform the coupled gas and radiation energy equations into a set of nonlinear ODEs in dimensionless gas and radiation temperature variables, which must be integrated numerically to achieve semi-analytic solutions. Nonetheless, their radiation diffusion model corresponds identically to that implemented in Enzo in the grey LTE limit, and the unique structure of the post-shock material temperature profile for a given Mach number makes this problem an excellent verification test for computer codes.

We have run the Mach-2 test case described in [46]. The computational domain has a length of 0.1 cm. The material has a uniform initial density of 1.0 g cm^{-3} , a constant specific heat of $2.218056 \times 10^{12} \text{ erg g}^{-1} \text{ eV}^{-1}$, and a uniform initial velocity of $1.9475 \times 10^5 \text{ cm s}^{-1}$. The material and radiation are assumed to be in thermal equilibrium at $t = 0$ at a temperature of 121.6 eV. Outflow and reflecting boundary conditions are imposed upon the left and right boundaries, respectively, resulting in a shock wave that forms near the right boundary and propagates to the left. The total evolution time is 1.73325 nanoseconds.

Fig. 6 shows the result of a high-resolution simulation (4096 zones along the propagation axis). The curves represent the dimensionless gas (solid curve) and radiation (dashed curve) temperatures, T and T_r . The circles and squares are taken from exact solution data kindly provided by Lowrie for this parameter set. Both the gas and radiation have dimensionless far-field temperatures of 1.0 in the pre-shock state. Examining the gas temperature curve, there are three features of particular interest: the *precursor*, in which the material is preheated ahead of the shock front by the radiation wave which travels ahead of the shock; the *Zel'dovich spike*, shown by the overshoot in temperature at the shock front, and the *radiation relaxation region*, delineated by the decline in the material temperature to its eventual far-field postshock value. Letting T_p denote the maximum preshock value of the gas temperature in the precursor, and T_1 the asymptotic postshock value, we note that the property $T_p < T_1$ identifies this calculation as an example of a *subcritical* radiating shock. In the limit of high Mach number, T_p can become equal to (but never exceed) T_1 , such a shock wave is referred to as *supercritical*.

As vividly demonstrated by [46], the strength of the precursor, the height of the Zel'dovich spike, and the precise temperature structure in the relaxation region are extremely sensitive to the Mach number. While the case we have shown is subcritical, it lies near the limit for which a multidimensional code can reasonably capture this structure without resorting to adaptive mesh refinement. The degree to which we resolve this structure as a function of resolution is shown in Fig. 7, in which we magnify the region near the shock and show gas temperature curves for mesh sizes of 128, 256, 512, 1024, 2048 and 4096 zones. As shown in [46], raising the Mach number results in a dramatic increase in the height of the spike and narrowing of the relaxation region; a proper representation of the postshock structure in a supercritical shock with Enzo must await the implementation of adaptive mesh refinement in our radiation module.

Since this problem considers coupled radiation and hydrodynamics, we also examine how the adaptive time step selection strategy from Section 3.2.3 compares with the hydrodynamic CFL-limited time step. For this problem, the average radiation time step ranged from $1.1\text{e}-5$ down to $9.4\text{e}-6$ for the coarsest (128-cell) to finest (4096-cell) grids, exhibiting a near-constant time step selection that tracks evolution of the radiation field. For these same problems, the hydrodynamic CFL limits on the time steps were $3.3\text{e}-4$, $1.7\text{e}-4$, $4.2\text{e}-5$, $2.1\text{e}-5$, $1.0\text{e}-5$ and $5.0\text{e}-6$. Hence, for most problems the stiff radiation

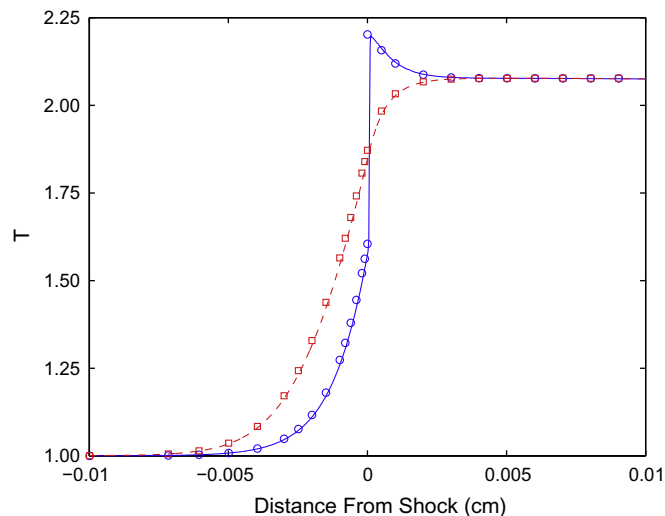


Fig. 6. Subcritical radiating shock test using 4096 spatial zones. Gas and radiation temperatures are plotted in units of their preshock values: T is the solid curve; T_r is the dashed curve. Semi-analytic values for T and T_r are indicated by circles and squares, respectively.

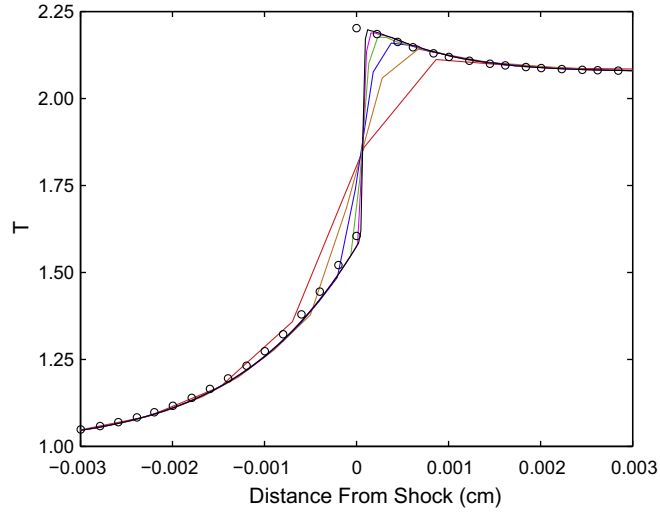


Fig. 7. Subcritical radiating shock convergence: gas temperature vs. distance from the shock for mesh resolutions of 128 (red), 256 (orange), 512 (blue), 1024 (green), 2048 (cyan) and 4096 zones (black). Semi-analytic reference values are indicated by open circles. (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

time scale limits the overall time step size, until very fine grids where the mesh-dependent hydrodynamic CFL stability condition becomes more restrictive.

4.5. Isothermal ionization of a static neutral hydrogen region

Our first test problem incorporating ionization chemistry is due to Iliev et al. [4]. This problem combines radiative transfer and hydrogen ionization in a static astrophysical region. The physical situation of interest is the expansion of an ionized hydrogen (HII) region in a uniform gas around a single monochromatic ionizing source emitting $\dot{N}_\gamma = 5 \times 10^{48}$ photons per second at the ionization frequency of hydrogen ($h\nu = 13.6$ eV). We enforce a fixed gas temperature of $T = 10^4$ K, and a static hydrodynamic state (i.e. $\dot{\rho}_b = \dot{\mathbf{v}}_b = \dot{e} = 0$). In such a problem, the radiation source should rapidly ionize the surrounding hydrogen, and then should develop a spherically-propagating ionization front (I-front) that propagates quickly at first, slows, and then eventually stagnates at an equilibrium position referred to as the Strömgren radius, where ionization ($\text{HI} \rightarrow \text{HII}$) and recombinations ($\text{HII} \rightarrow \text{HI}$) balance. For this scenario, the analytically-provided I-front radius is given by

$$r_I = r_S [1 - \exp(-t/t_{\text{rec}})]^{1/3}, \quad \text{where} \quad (62)$$

$$r_S = \left[\frac{3\dot{N}_\gamma}{4\pi\alpha_B n_H^2} \right]^{1/3}, \quad (63)$$

and the recombination time is given by $t_{\text{rec}} = (\alpha_B n_H)^{-1}$. Here, $\alpha_B = 2.59 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$ is the case B hydrogen recombination coefficient.

We have the following problem parameters: the domain size is $L = 6.6$ kpc in each direction; the initial gas number density is $n_H = 10^{-3} \text{ cm}^{-3}$; the initial radiation energy density is $E = 10^{-20} \text{ erg cm}^{-3}$; the initial ionization fraction (HII/H) is 0.0012; the ionization source is located in the lower corner of the box (the (1, 1, 1) cell); we use reflecting boundary conditions at the x -, y - and z -left boundaries, and outflow conditions at the corresponding right boundaries. For these parameters the Strömgren radius $r_S = 5.4$ kpc, the recombination time $t_{\text{rec}} \approx 3.86 \times 10^{15} \text{ s}$ (≈ 122.4 Myr), and the total simulation time is 500 Myr ($\approx 4t_{\text{rec}}$). The implicit solver parameters used were a convergence norm of $p = 2$, desired solution tolerance of $\tau_{\text{tol}} = 0.01$, time-step parameter of $\theta = 0.51$, and nonlinear solver tolerance of $\varepsilon = 10^{-7}$.

In Fig. 8 we plot the spherically-averaged I-front position and radius with respect to time, for various spatial mesh sizes. The I-front position is computed from our results as the distance at which cells transition from below 50% to above 50% HII fractional density. Assuming a spherical HII region, we compute this radius as $r_S = (8 \frac{3V}{4\pi})^{1/3}$, where V is the volume comprised of all ionized cells (i.e. where $n_{\text{HII}}/n_H \geq 0.5$), and the additional factor of 8 arises due to the fact that our source is in the corner, so we must mirror V into the other 7 octants. We also plot the error in the computed I-front radius and velocity for varying mesh sizes. As can be clearly seen, the computed I-front position is highly accurate, even for coarse spatial grids, with the corresponding accuracy increasing as the mesh is resolved.

In Fig. 9 we show cross-sections of the radiation energy density through the ionization source for 16^3 and 128^3 grids. We note that although the spherical front is jagged for coarse grids, as the mesh is refined we approach the physically-accurate spherical profile. Moreover, this demonstrates that although the flux-limiter (12) is based on one-dimensional derivatives, it does not result in anisotropic propagation biased along axial directions.

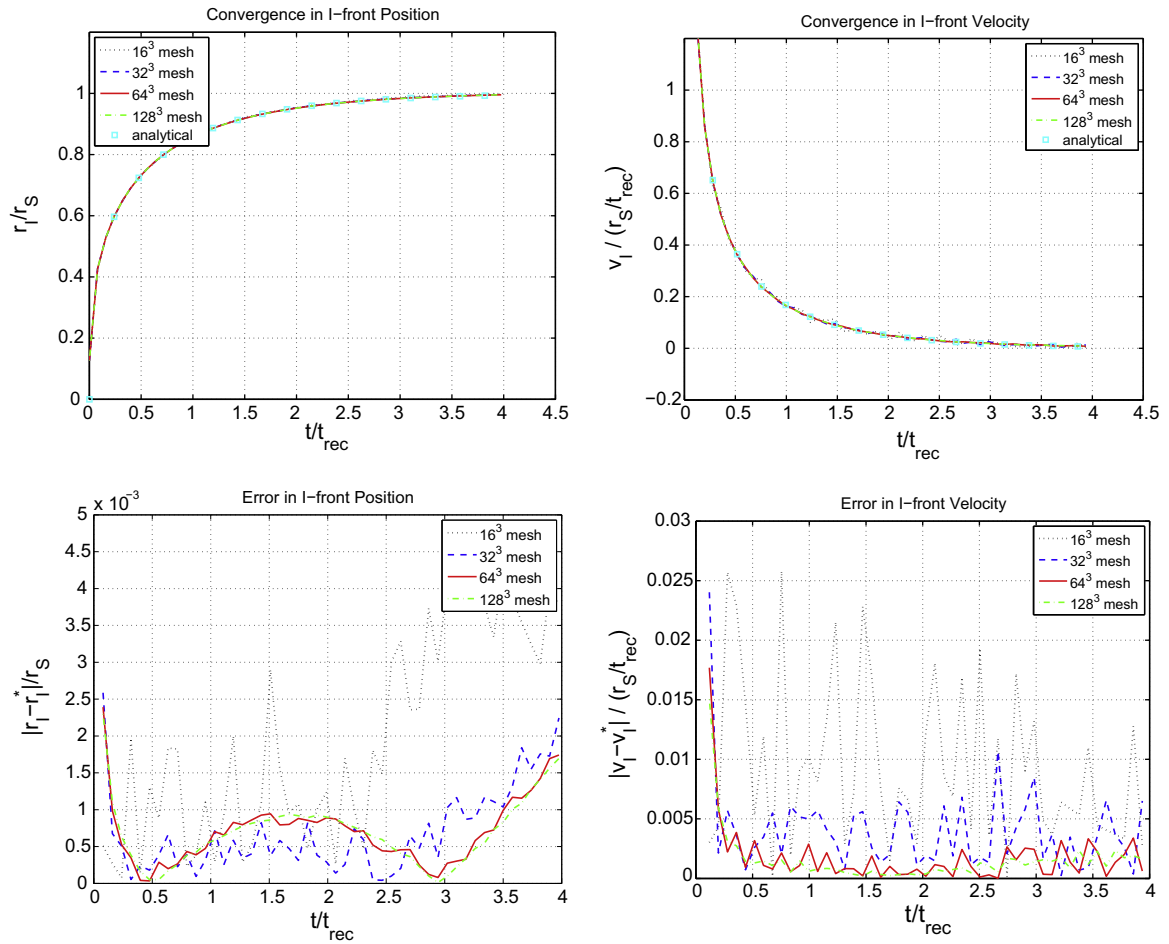


Fig. 8. Left: evolution of the I-front position, analytical solution, and relative error for varying resolutions. Right: computed I-front velocity and relative error.

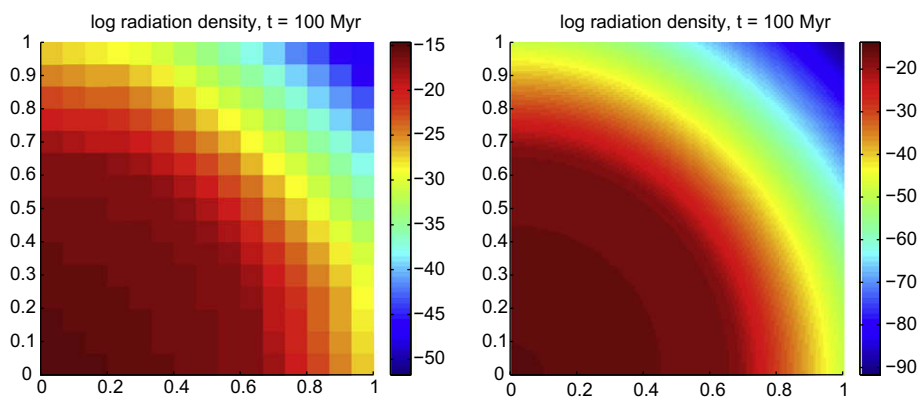


Fig. 9. Cross-sections of the radiation density through the ionization source at $t = 100$ Myr: 16³ grid (left) and 128³ grid (right). Note the convergence to the spherical analytical shape.

4.6. Cosmological radiative ionization

We now perform the same test as above, but within a cosmologically-expanding universe. The problem is originally due to Shapiro and Giroux [47], and combines cosmology, radiative transfer and chemical ionization. Here, the physics of interest

is again the expansion of a HII region in uniform gas around a single monochromatic ($h\nu = 13.6$ eV) ionizing source. Again, the ionization front should propagate quickly at first, approaching the Strömgen radius, but then should begin to lag behind as cosmological expansion drives the Strömgen radius outward faster than the I-front can propagate. Due to the cosmological expansion, the Strömgen radius changes in time, and is given by

$$r_s(t) = \left[\frac{3\dot{N}_\gamma}{4\pi\alpha_B n_H(t)^2} \right]^{1/3}, \tag{64}$$

where although the Hydrogen number density n_H is spatially static, it diminishes due to cosmological expansion by a factor of $a^{-3}(t)$. Defining the parameter $\lambda = \alpha_B n_{H,i}/H_0/(1+z_i)$, where the subscript i refers to the quantity at the initial redshift z_0 , the analytical solution is given by

$$r_i(t) = r_{s,i} \left(\lambda e^{-\tau(t)} \int_1^{a(t)} e^{\tau(\bar{a})} [1 - 2q_0 + 2q_0(1+z_i)/\bar{a}]^{-1/2} d\bar{a} \right)^{1/3}, \tag{65}$$

where

$$\tau(a) = \lambda [6q_0^2(1+z_i)^2]^{-1} [F(a) - F(1)], \tag{66}$$

$$F(a) = [2 - 4q_0 - 2q_0(1+z_i)/a][1 - 2q_0 + 2q_0(1+z_i)/a]^{1/2}. \tag{67}$$

Here, q_0 is the cosmological deceleration parameter and z_i is the initial redshift. We perform four of the tests provided in the original paper [47]: q_0 of 0.5 and 0.05, and z_i of 4 and 10. These correspond to the parameters:

- $(q_0, z_i) = (0.5, 4) : L_i \approx 80$ kpc, $H_0 = 0.5, \Omega_m = 1.0, \Omega_\Lambda = 0, \Omega_b = 0.2$
- $(q_0, z_i) = (0.05, 4) : L_i \approx 60$ kpc, $H_0 = 1, \Omega_m = 0.1, \Omega_\Lambda = 0, \Omega_b = 0.1$
- $(q_0, z_i) = (0.5, 10) : L_i \approx 36$ kpc, $H_0 = 0.5, \Omega_m = 1.0, \Omega_\Lambda = 0, \Omega_b = 0.2$
- $(q_0, z_i) = (0.05, 10) : L_i \approx 27$ kpc, $H_0 = 1, \Omega_m = 0.1, \Omega_\Lambda = 0, \Omega_b = 0.1$

where L_i is the initial box size, H_0 is the Hubble constant, Ω_m is the contribution of all non-relativistic matter to the gas energy density at $z = 0$, in units of the value required to close the universe, similarly Ω_Λ and Ω_b are the contributions of the cosmological constant and the baryonic matter to the energy density, respectively. These two types of cosmology result in slightly different functions for the expansion coefficient a . For the case of $q_0 = 0.05$, this value comes from equations (13-3) and (13-10) in [48]. For the case $q_0 = 0.5$, we use the standard formula $a = (1+z)^{-1}$. We begin all problems with an initial radiation energy density of $E = 10^{-35}$ erg cm⁻³ and an initial ionization fraction (HII/H) of 0. The initial density is dependent on q_0 , with $\rho_{b,i} = 1.175 \times 10^{-28}$ g cm⁻³ for $q_0 = 0.5$, and $\rho_{b,i} = 2.35 \times 10^{-28}$ g cm⁻³ for $q_0 = 0.05$. All simulations are run from the initial redshift z_i to $z = 0$. All other problem parameters are identical to those in Section 4.5. All implicit solver parameters are also identical to those in Section 4.5, but with desired solution accuracy $\tau_{tol} = 0.001$ and inexactness parameter $\delta_k = 10^{-13} \|f(U_k)\|$.

In Fig. 10 we plot the scaled, spherically-averaged I-front position with respect to scaled redshift for each of the four tests (with axes identical to [47], Fig. 1a), as well as the corresponding plots for just the $z_i = 4$ tests along with their analytical solutions. These solutions all used a uniform 128³ spatial mesh. In Fig. 11 we plot the error in the computed I-front radius

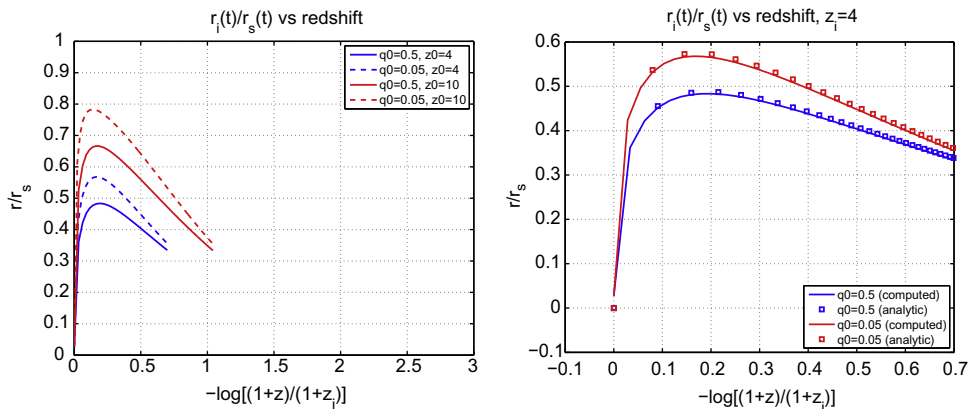


Fig. 10. Left: I-front radii vs. redshift for the four tests, $q_0 = 0.5$ and 0.05 , and $z_i = 4$ and 10 . Right: I-front radii vs. redshift for the $z_i = 4$ tests; analytical solution values are given by the open squares.

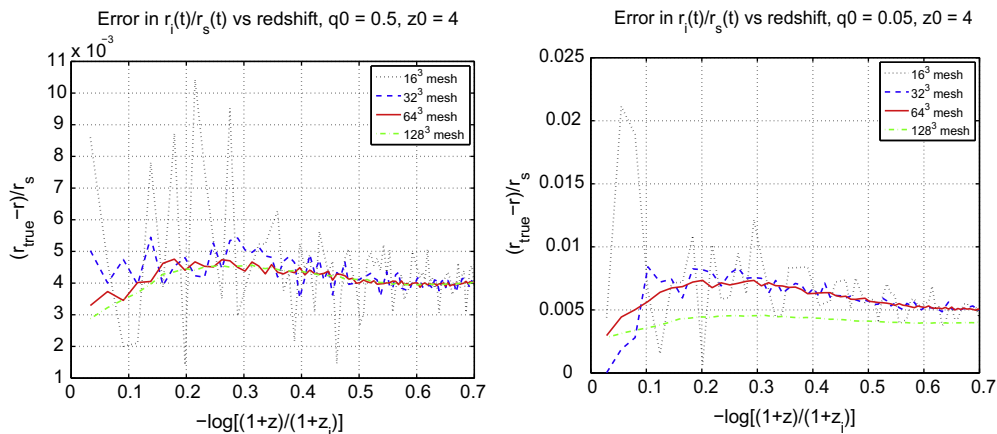


Fig. 11. Convergence of I-front radius vs. redshift for the two cases $q_0 = \{0.5, 0.05\}$ and $z_i = 4$ as the mesh is refined: spatial meshes shown are 16^3 (black dotted), 32^3 (blue dashed), 64^3 (red solid), and 128^3 (green dot-dashed). (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

for varying mesh sizes for the two cases of $q_0 = 0.5$ and $q_0 = 0.05$ with $z_i = 4$. Again, the accuracy in the computed I-front position improves as the mesh is resolved.

4.7. Hydrodynamic radiative ionization

We now incorporate hydrodynamic motion into the mixture of physical processes, and examine a problem due to Whalen and Norman [7] that combines radiation, hydrodynamics and chemical ionization (but not cosmology). The problem is nearly identical to that from Section 4.5, but now in a dynamic medium (varying temperature, density and velocity). Again, the physics of interest is the expansion of a HII region in an initially uniform gas around a single ionizing source, though now the source emits $\dot{N}_\gamma = 5 \times 10^{48}$ photons per second with a frequency profile given by a $T_b = 10^5$ blackbody spectrum. Here, the ionization front should propagate quickly at first, slowing until it reaches the Strömgren radius (63), at which point the I-front transitions from radiation-driven (R-type) to dynamically-driven (D-type), and the high pressure of the ionized and heated gas inside the HII region continues to push the I-front out past the Strömgren sphere. The expansion will finally stall when a pressure equilibrium has been reached, at a radius $r_f = (2T_i/T_e)^{2/3} r_s$, where T_i is the temperature of the ionized gas behind the front, and T_e is the temperature of the ionized gas ahead of the front. Analytical models for the initial radiation-only phase predict

$$r_I^R = r_s [1 - e^{-t/t_{rec}}]^{1/3}, \quad (68)$$

where $t_{rec} = [\alpha_B(T_i)n_H]^{-1}$ is the recombination time (assumed constant in this phase). Analytical models for the subsequent pressure-only phase predict

$$r_I^P = r_s \left(1 + \frac{7c_s t}{4r_s}\right)^{4/7}, \quad (69)$$

where $c_s = \sqrt{p_i/\rho_i}$ is the sound speed in the ionized gas. We note that due to the dynamic nature of this problem, the true solution should lie between the two curves (68) and (69), since both radiation and gas pressure play a role throughout the dynamics, and neither T_i or T_e are in fact constant behind or beyond the I-front.

We use the following problem parameters: the initial gas temperature is set to $T = 10^2$ K; the initial radiation energy density is $E^0 = 10^{-20}$ erg cm $^{-3}$; the hydrogen is initially fully neutral (i.e. HII/H=0); the spatial domain is a 15 kpc box. We run for a simulation time of 1 Gyr, which is not long enough to reach final the final equilibrium r_f , but well past the transition from R-type to D-type. The implicit solver parameters are identical to those in Section 4.5, but with a linear solver parameter $\delta_k = 10^{-9} \|f(U_k)\|$.

Results from these tests are shown in Fig. 12, which plot the computed I-front position and neutral fractions for various spatial meshes, along with the “error” in these quantities. Since this test problem does not have true analytical solutions, we compute the “error” as the deviation in each solution from the most-refined 128^3 mesh solution.

We also use this problem to examine the effect of our operator-split solution strategy on the temporal accuracy of the solver. In Fig. 13 we plot the spherically-averaged temperature profile for a 128^3 spatial grid at 175 Myr, and the associated relative errors found through varying the time-step size. The error plot has been zoomed in around the heated region and front. We note that although both PPM and the described implicit sub-solver are both up to second-order accurate, the splitting reduces the resulting accuracy to slightly better than first-order in time. We also note that the adaptive time-stepping

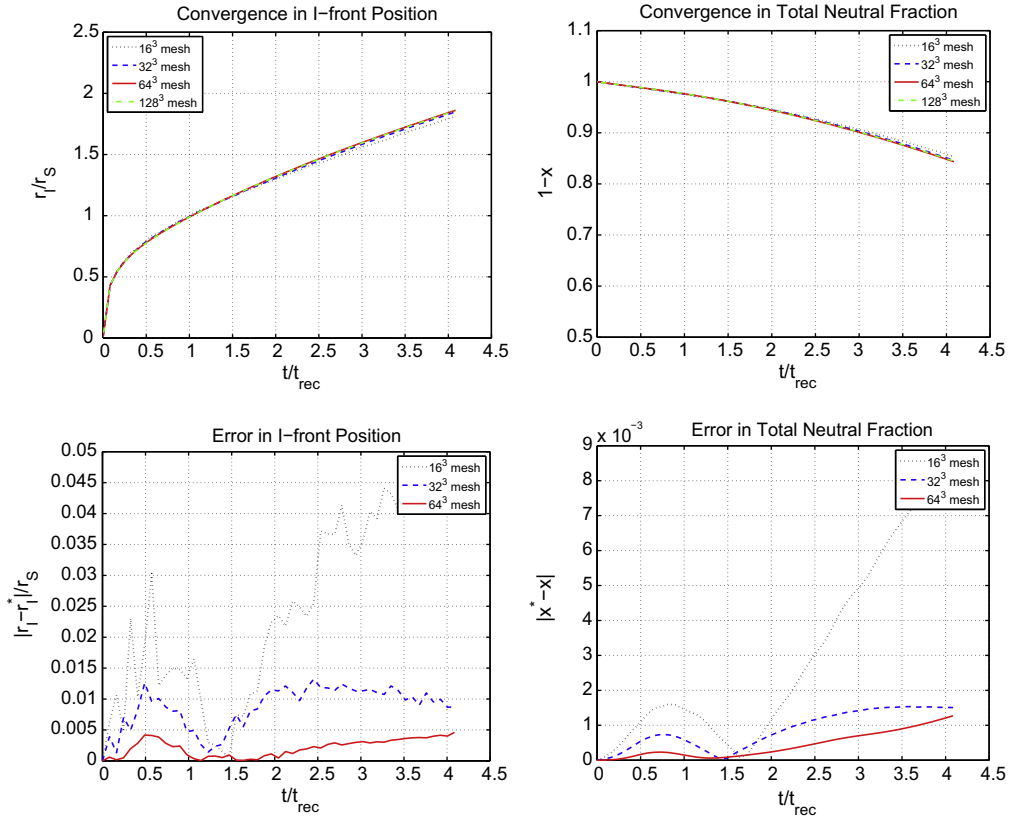


Fig. 12. Convergence of the hydrodynamic ionization test with mesh refinement. Left: overlay of computed I-front position (top) and error (bottom) for varying mesh sizes. Right: overlay of computed neutral fraction (top) and error for varying grids.

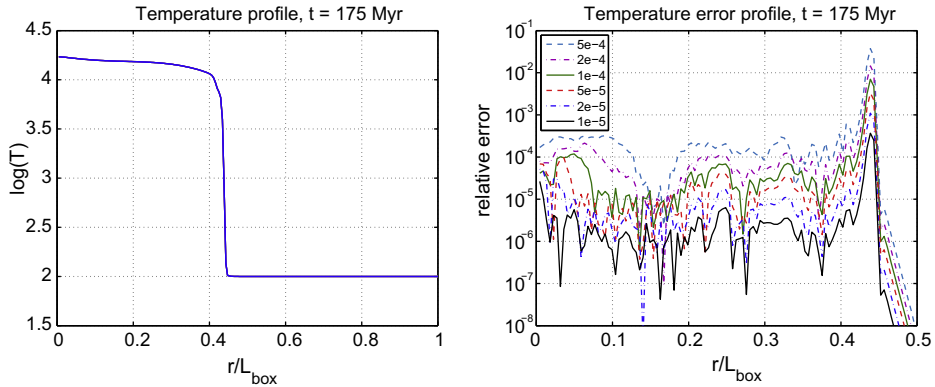


Fig. 13. Convergence of temperature profile as the time step size is refined. Left: the spherically-averaged temperature profile at 175 Myr. Right: relative errors in temperature for varying time step sizes. Average relative errors were $2.6e-4$, $1.0e-4$, $4.9e-5$, $2.4e-5$, $8.2e-6$, and $2.7e-6$, for the coarsest to finest time steps, respectively.

strategy from Section 3.2.3 results in average time steps of $6.0e-4$, $5.7e-4$, $4.0e-4$ and $2.7e-4$ for the 16^3 through 128^3 grids, respectively; whereas the hydrodynamic CFL-limited time steps for these same grids are $3.1e-3$, $1.4e-3$, $6.0e-4$ and $2.8e-4$. Hence the radiation and ionization dynamics drive the system for coarser meshes, while at finer mesh sizes the hydrodynamic CFL condition begins to dominate.

4.8. Weak scaling

As described in our introduction and throughout the description of our numerical methods, a key goal in introducing a fully-implicit solution mechanism for the stiff components in radiation, hydrodynamics and chemical ionization simulations

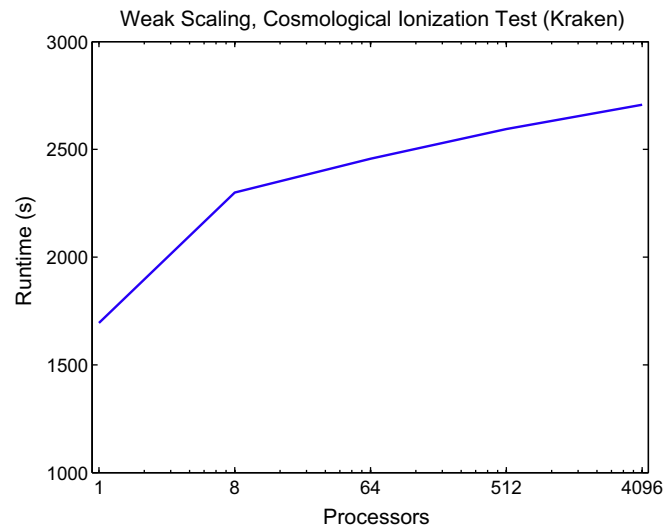


Fig. 14. Weak scaling results for the cosmological HII-region expansion test.

Table 2

Cosmological weak scaling statistics.

Mesh	Processors	Time steps	Run time	Newton Its	CG Its	MG V-cycles
64^3	1	266	1694.38	322	914	2991
128^3	8	265	2299.60	274	799	2575
256^3	64	265	2456.58	268	787	2524
512^3	512	264	2594.50	265	780	2510
1024^3	4096	264	2707.30	265	780	2510

is the eventual scalability of such a solver to very large problem sizes. We therefore investigate the weak scaling of the implicit solver on the cosmological radiative ionization problem from Section 4.6. For these tests, we emulate the setup from the $q_0 = 0.5$ and $z_i = 4$ test, but here we place an ionizing source in the center of *each processor's* subgrid. Moreover, for these weak scaling tests we increase the domain size and mesh size proportionately to the number of processors, where each processor owns a 64^3 grid and an initial subgrid box size of 80 kpc. We then run problems that scale up from 1 to 4096 processors, resulting in spatial grids ranging from 64^3 to 1024^3 . Moreover, since we are investigating the scaling properties of the numerical methods, we shorten the simulation time to evolve from $z_i = 4$ to $z = 3$ in order to conserve on supercomputer resources, while retaining the portion of the simulation with the most rapidly-evolving dynamics. All runs were performed on the NSF Kraken machine (using 2 cores/node).

We show the runtimes associated with these tests in Fig. 14, and provide detailed statistics from each run in Table 2. We note that on this architecture, the solver demonstrates near-perfect scalability, with modest increases in runtime for parallel versus serial runs, and only marginal increases in solution time as the parallelism is increased. The reason for this is the near-constant number of iterations required by the nonlinear Newton solver, the outer CG linear solver, and the inner multigrid preconditioner. Therefore the increase in run time may be directly attributed to the ideal $\mathcal{O}(\log p)$ increase in runtimes typical of multigrid methods, allowing near-optimal scalability to the limits of modern supercomputer resources.

5. Conclusions

We have described an implicit formulation for coupling cosmological radiation transport, chemical ionization and gas energy feedback within Enzo hydrodynamics simulations. The formulation is based on an operator-splitting between the non-stiff hydrodynamics and stiff radiation-ionization-energy feedback physical processes, in which the stiff processes are solved within a fully-implicit Newton-Schur-Krylov-Multigrid framework.

Through numerous tests, we have demonstrated that this solver is accurate and stable, allowing simulations of a wide variety of physical environments from the laboratory scale to the astronomical and even cosmological scales. Moreover, through the choice of numerical methods that form the implicit solver, it demonstrates ideal scalability for such coupled physics simulations. In addition, this implicit formulation is highly extensible, and may easily be adjusted to allow new physical processes such as magnetic fields, multi-frequency radiation transfer, and additional chemical species. Finally, we are in the process of extending this approach to allow for adaptive spatial discretizations (AMR), which should require adjustments to only the inner multigrid linear solver.

Acknowledgments

The authors wish to acknowledge the many insightful discussions with Frank Graziani, Louis Howell, Carol Woodward, and Doug Swesty during the course of this work; as well as Rob Lowrie for providing semi-analytical solutions to (Section 4.4). The verification tests were performed on the Itanium 2 linux cluster *Thunder* at the Lawrence Livermore National Laboratory, and on the IBM Power 4 supercomputer *DataStar* at the San Diego Supercomputing Center. The weak scaling tests were performed on the Cray XT4 system *Kraken* at the National Institute for Computational Sciences. This work was supported in part by NSF grants AST-0708960 and AST-0808184, NASA ATP grant NNX08AH26G, and by UCSD-LLNL collaborative grant “LUSciD: Scientific Data Management” to MLN, and performed in part under the auspices of the US Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract W-7405-ENG-48.

References

- [1] R. Barkana, A. Loeb, The physics and early history of the intergalactic medium, Reports on Progress in Physics 70 (2007) 627–657.
- [2] M.L. Norman, G.L. Bryan, R. Harkness, J. Bordner, D.R. Reynolds, B. O’Shea, R. Wagner, in: Petascale Computing: Algorithms and Applications, Simulating Cosmological Evolution with Enzo, CRC Press, 2007.
- [3] V. Springel, The cosmological simulation code GADGET-2, MNRAS 364 (2005) 1105–1134.
- [4] I.T. Iliev, B. Ciardi, M.A. Alvarez, A. Maselli, A. Ferrara, N.Y. Gnedin, G. Mellema, T. Nakamoto, M.L. Norman, A.O. Razoumov, E.-J. Rijkhorst, J. Ritzerveld, P.R. Shapiro, H. Susa, M. Umemura, D.J. Whalen, Cosmological radiative transfer codes comparison project – I. The static density field tests, MNRAS 371 (2006) 1057–1086.
- [5] G. Mellema, I.T. Iliev, M.A. Alvarez, P.R. Shapiro, C²-ray: a new method for photon-conserving transport of ionizing radiation, New A 11 (2006) 374–395.
- [6] H. Susa, M. Umemura, Formation of dwarf galaxies during the cosmic reionization, ApJ 600 (2004) 1–16.
- [7] D. Whalen, M.L. Norman, A multistep algorithm for the radiation hydrodynamical transport of cosmological ionization fronts and ionized flows, ApJ 162 (Suppl.) (2006) 281–303.
- [8] E.-J. Rijkhorst, T. Plewa, A. Dubey, G. Mellema, Hybrid characteristics: 3D radiative transfer for parallel adaptive mesh refinement hydrodynamics, A&A 452 (2006) 907–920.
- [9] A. Maselli, A. Ferrara, B. Ciardi, CRASH: a radiative transfer scheme, MNRAS 345 (2003) 379–394.
- [10] B. Semelin, F. Combes, S. Baek, Lyman-alpha radiative transfer during the epoch of reionization: contribution to 21-cm signal fluctuations, A&A 474 (2007) 365–374.
- [11] N.Y. Gnedin, T. Abel, Multi-dimensional cosmological radiative transfer with a variable Eddington tensor formalism, New A 6 (2001) 437–455.
- [12] J.C. Hayes, M.L. Norman, Beyond flux-limited diffusion: parallel algorithms for multidimensional radiation hydrodynamics, ApJ 147 (Suppl.) (2003) 197–220.
- [13] P. Paschos, M.L. Norman, J.O. Bordner, R. Harkness, Late reheating of the IGM by quasars: a radiation hydrodynamical simulation of helium II reionization, ArXiv e-prints arXiv:0711.1904.
- [14] G.L. Bryan, M.L. Norman, J.M. Stone, R. Cen, J.P. Ostriker, A piecewise parabolic method for cosmological hydrodynamics, Comp. Phys. Commun. 89 (1995) 149–168.
- [15] P. Paschos, On the ionization and chemical evolution of the intergalactic medium, Ph.D. Thesis, University of Illinois at Urbana-Champaign, 2005.
- [16] T. Abel, P. Anninos, Y. Zhang, M.L. Norman, Modeling primordial gas in numerical cosmology, New A 2 (1997) 181–207.
- [17] L. Hui, N.Y. Gnedin, Equation of state of the photoionized intergalactic medium, MNRAS 292 (1997) 27–42.
- [18] R.W. Hockney, J.W. Eastwood (Eds.), Computer Simulation Using Particles, 1988.
- [19] M.L. Norman, G.L. Bryan, Cosmological adaptive mesh refinement, in: S.M. Miyama, K. Tomisaka, T. Hanawa (Eds.), Numerical Astrophysics, Astrophysics and Space Science Library, vol. 240, 1999, pp. 19–28.
- [20] B.W. O’Shea, G. Bryan, J.O. Bordner, M.L. Norman, T. Abel, R. Harkness, A. Kritsuk, in: Adaptive Mesh Refinement – Theory and Applications, Lecture Notes in Computational Science and Engineering, Introducing Enzo, An AMR Cosmology Application, Springer, 2004.
- [21] J.C. Hayes, M.L. Norman, R.A. Fiedler, J.O. Bordner, P.S. Li, S.E. Clark, A. ud-Doula, M.-M.M. Low, Simulating radiating and magnetized flows in multiple dimensions with ZEUS-MP, ApJ 165 (Suppl.) (2006) 188–228.
- [22] L.H. Howell, J.A. Greenough, Radiation diffusion for multi-fluid Eulerian hydrodynamics with adaptive mesh refinement, J. Comp. Phys. 184 (2003) 53–78.
- [23] J. Black, The physical state of primordial intergalactic clouds, Mon. Not. R. Astr. Soc. 197 (1981) 553–563.
- [24] D.E. Osterbrock, Astrophysics of Gaseous Nebulae and Active Galactic Nuclei, University Science Books, Mill Valley, California, 1989.
- [25] R. Cen, A hydrodynamic approach to cosmology: Methodology, ApJ 78 (Suppl.) (1992) 341.
- [26] A.O. Razoumov, M.L. Norman, T. Abel, D. Scott, Cosmological hydrogen reionization with three-dimensional radiative transfer, ApJ 572 (2002) 695–704.
- [27] C.D. Levermore, Relating Eddington factors to flux limiters, J. Quant. Spectrosc. Radiat. Transfer 31 (2) (1984) 149–160.
- [28] C.D. Levermore, G.C. Pomraning, A flux-limited diffusion theory, ApJ 248 (1981) 321–334.
- [29] P. Colella, P.R. Woodward, The piecewise parabolic method (PPM) for gas-dynamical simulations, J. Comp. Phys. 54 (1984) 174–201.
- [30] Enzo code project page. <<http://lca.ucsd.edu/portal/software/enzo>>.
- [31] J.C. Simo, F. Armero, Unconditional stability and long-term behavior of transient algorithms for the incompressible Navier–Stokes and Euler equations, CMAME 111 (1994) 111–154.
- [32] O. Østerby, Five ways of reducing the Crank–Nicolson oscillations, BIT Numer. Math. 43 (2003) 811–822.
- [33] C.T. Kelley, Iterative methods for linear and nonlinear equations, Frontiers Appl. Math., vol. 16, SIAM, 1995.
- [34] D.A. Knoll, D.E. Keyes, Jacobian-free Newton–Krylov methods: a survey of approaches and applications, J. Comp. Phys. 193 (2004) 357–397.
- [35] J.E. Dennis, R.B. Schnabel, Numerical Methods for Unconstrained Optimization and Nonlinear Equations, SIAM, Philadelphia, 1996.
- [36] J.M. Ortega, W.C. Rheinboldt, Iterative Solution of Nonlinear Equations in Several Variables, SIAM, 2000.
- [37] S.C. Eisenstat, H.F. Walker, Choosing the forcing terms in an inexact Newton method, SIAM J. Sci. Comp. 17 (1996) 16–32.
- [38] M. Weiser, A. Schiela, P. Deuffhard, Asymptotic mesh independence of Newton’s method revisited, SIAM J. Numer. Anal. 42 (2005) 1830–1845.
- [39] P.N. Brown, C.S. Woodward, Preconditioning strategies for fully implicit radiation diffusion with material-energy transfer, SIAM J. Sci. Comp. 23 (2001) 499–516.
- [40] R.D. Falgout, U.M. Yang, in: Computational Science – ICCS 2002 Part III, Lecture Notes in Computer Science, Hypre: A Library of High Performance Preconditioners, vol. 2331, Springer-Verlag, 2002, pp. 632–641.
- [41] HYPRE code project page. <<http://www.llnl.gov/CASC/hypre/software.html>>.
- [42] A.H. Baker, R.D. Falgout, U.M. Yang, An assumed partition algorithm for determining processor inter-communication, Parallel Comput. 32 (2006) 319–414.
- [43] N.J. Turner, J.M. Stone, A module for radiation hydrodynamic calculations with ZEUS-2D using flux-limited diffusion, ApJ 135 (Suppl.) (2001) 95–107, doi:10.1086/321779.

- [44] G.C. Pomraning, The non-equilibrium Marshak wave problem, *JQSRT* 21 (1979) 249–261.
- [45] B. Su, G.L. Olson, Benchmark results for the non-equilibrium Marshak diffusion problem, *JQSRT* 56 (1996) 337–351.
- [46] R.B. Lowrie, J.D. Edwards, Radiative shock solutions with grey nonequilibrium diffusion, *Shock Waves* 18 (2) (2008) 129–143.
- [47] P.R. Shapiro, M.L. Giroux, Cosmological HII regions and the photoionization of the intergalactic medium, *ApJ* 321 (1987) L107–L112.
- [48] P.J.E. Peebles, *Principles of Physical Cosmology*, Princeton University Press, 1993.