MAGNETOHYDRODYNAMIC EVOLUTION OF H II REGIONS IN MOLECULAR CLOUDS: SIMULATION METHODOLOGY, TESTS, AND UNIFORM MEDIA

Mark R. Krumholz,[1] James M. Stone, and Thomas A. Gardiner
Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544
Received 2006 June 21; accepted 2007 August 21

ABSTRACT

We present a method for simulating the evolution of H II regions driven by point sources of ionizing radiation in magnetohydrodynamic media, implemented in the three-dimensional Athena MHD code. We compare simulations using our algorithm to analytic solutions and show that the method passes rigorous tests of accuracy and convergence. The tests reveal several conditions that an ionizing radiation hydrodynamic code must satisfy to reproduce analytic solutions. As a demonstration of our new method, we present the first three-dimensional, global simulation of an H II region expanding into a magnetized gas. The simulation shows that magnetic fields suppress sweeping up of gas perpendicular to magnetic field lines, leading to small density contrasts and extremely weak shocks at the leading edge of the H II region's expanding shell.

Subject headings: H II regions — instabilities — ISM: kinematics and dynamics — MHD —
methods: numerical — radiative transfer

Online material: mpeg animation

1. INTRODUCTION

Observations show that giant molecular clouds (GMCs) in Local Group galaxies convert at most a few percent of their mass into stars per cloud crossing time (Zuckerman & Evans 1974) and that clouds are typically destroyed in a few crossing times (Blitz et al. 2007), well before they have converted a significant fraction of their mass into stars. Observations also strongly support the idea that GMCs are gravitationally bound (Krumholz & McKee 2005; Blitz et al. 2007; Krumholz et al. 2006), so the fact that they survive for more than a crossing time but do not collapse entirely into stars strongly suggests that internal feedback plays a dominant role in GMC evolution. H II regions driven by newly formed massive stars are likely to be the dominant sources of energy injection and mass loss in clouds and are therefore critical to GMC evolution (McKee & Williams 1997; Williams & McKee 1997; Matzner 2002). Krumholz et al. (2006) show, using semianalytic models, that feedback from H II regions can quantitatively reproduce the observed lifetime, star formation rate, and star formation efficiency of GMCs. However, these calculations rely on simple analytic solutions for the evolution of spherically symmetric H II regions in nonmagnetic gas. Understanding the detailed evolution of GMCs will require a considerably more sophisticated numerical approach, and for this reason, three-dimensional (3D) simulation of the evolution of molecular clouds under the influence of internal sources of ionizing radiation is a critical problem in numerical astrophysics.

In this paper we have the dual purpose of presenting a new algorithm for simulation of H II regions in molecular clouds and using this algorithm to explore potential computational problems that arise in general for simulations of this type. We also demonstrate our new algorithm in a simple application, the expansion of an H II region into a uniform neutral gas in which the magnetic pressure greatly exceeds the thermal pressure, as it does in molecular clouds. In the past year several authors have considered the problem of simulating H II regions and have presented both algorithms (e.g., Arthur & Hoare 2006; Mellema et al. 2006b) and results on the evolution of H II regions in turbulent hydrodynamic

media (Dale et al. 2005; Mellema et al. 2006a; Mac Low et al. 2007). For a recent review, see Henney (2007). Several authors have also presented methods for the simulation of ionizing radiation hydrodynamics (IRHD) in a cosmological context (e.g., Abel & Wandelt 2002; Whalen & Norman 2006), which differs from the problem in the context of present-day molecular clouds primarily in the amount of cooling to which the gas is subjected and the conditions of the gas before it is ionized. Finally, researchers studying the evolution of ultracompact H II region planetary nebulae have presented algorithms and results for ionizing radiative transfer with hydrodynamics and magnetohydrodynamics in two and three dimensions under the simplifying assumption that the ionized gas has a perfectly sharp edge and is always in thermal and ionization equilibrium (e.g., Garcia-Segura & Franco 1996; Garcia-Segura 1997; García-Segura & López 2000).

Our algorithm improves on previous work in that it is the first to couple ionizing radiative transfer to magnetohydrodynamics (IRMHD) without imposing an assumption of thermal or ionization equilibrium. The inclusion of magnetic fields is important because observations indicate that the magnetic energy in molecular clouds is comparable to the kinetic and gravitational potential energies (Crutcher 1999, 2005; Heiles & Crutcher 2005). Thus, dynamical expansion of ionized regions may be significantly altered by magnetic confinement, an effect we wish to explore. Indeed, we show that even in the simple case of expansion of an H II region into a uniform, magnetized medium, the magnetic field produces qualitatively new phenomena. The Alfvén speeds of a few km s^{-1} that are typically found in molecular clouds reduce the strength of the shocks associated with expanding ionization fronts at early times, and at later times they turn off the shocks entirely, greatly reducing the collection of gas and possibly thereby reducing the amount of triggered star formation. A nonequilibrium treatment of the ionization structure is important because we find that a nonequilibrium treatment of the thermal and ionization structure of the ionized gas leads to modest but significant effects on quantities such as the expansion rate of H II regions, effects that are lost in the assumption of perfect radiative equilibrium.

Before exploring these new phenomena, however, we point out that there has yet to be a detailed study of potential computational problems and constraints that arise in 3D simulations of

---

[1] Hubble Fellow.

IRHD and IRMHD with the strong cooling and large temperature contrasts that are expected for modern-day (as opposed to primordial) interstellar chemistry. We have therefore performed a detailed comparison of simulations using our method to analytic solutions for computationally challenging problems, as a way of searching for potential difficulties that may arise in general IRHD methods. We discover several conditions that a simulation must satisfy in order to reproduce analytic results correctly. One must limit the amount by which the gas pressure is allowed to change between hydrodynamic updates; for a ray-tracing method, one must periodically rotate the orientation of the rays; and one must either resolve the ionization front or restrict the rate of cooling at the front to suppress excess cooling due to numerical mixing. We show that failure to meet these conditions produces quantitatively incorrect results.

The remainder of this paper proceeds as follows. In § 2 we describe the physical formulation of the problem that we adopt, including our approximations and assumptions. In § 3 we present our simulation algorithm. In § 4 we compare our code to analytic solutions and thereby demonstrate the existence of conditions that numerical methods must satisfy in order to reproduce analytic solutions correctly. We use our method to simulate the evolution of H ɪɪ regions in magnetized media in § 5, and finally we summarize and present conclusions in § 6.

2. PHYSICAL FORMULATION

We wish to simulate the evolution of a magnetized molecular gas subjected to an ionizing radiation field due to point sources within it. The n sources of ionization are located at positions $\boldsymbol{x}_n$, and each has an ionizing luminosity of s_n, in units of photons per unit time above the Lyman limit. We treat the ionizing sources as monochromatic and assume that they emit all their ionizing photons at frequencies near the ionization threshold of hydrogen. Since our focus is on the dynamics of the cloud rather than the detailed chemistry of the ionized gas, this approximation is quite reasonable.

We describe the gas by a total mass density ρ, density of neutral species ρ_n, velocity $\boldsymbol{v}$, and total energy density E. The gas is threaded by a magnetic field $\boldsymbol{B}$. The interface between an H ɪɪ region and a molecular cloud is a photodissociation region (PDR), within which there is a significant amount of atomic gas. In principle we should therefore separately track ionized, atomic, and molecular hydrogen, and possibly other species as well. However, PDRs do not have a significant effect on the large-scale dynamics of the ionized gas unless the ionizing source is very weak or the neutral gas is much less dense than is typical for molecular clouds, as we show in § 2.4. For this reason, we neglect atomic gas, and we assume that the density of molecular gas is ρ_n and the density of ionized gas is $\rho - \rho_n$.

2.1. Magnetohydrodynamics

We compute the evolution of our gas using the equations of multispecies ideal MHD, including radiative heating and cooling and chemical evolution terms. In conservative form, these are

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

$$\frac{\partial}{\partial t}(\rho \boldsymbol{v}) + \nabla \cdot (\rho \boldsymbol{vv} - \boldsymbol{BB}) + \nabla P^* = 0, \tag{2}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} + \nabla \cdot (\boldsymbol{vB} - \boldsymbol{Bv}) = 0, \tag{3}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot [(E + P^*)\boldsymbol{v} - \boldsymbol{B}(\boldsymbol{B} \cdot \boldsymbol{v})] = \mathcal{G} - \mathcal{L}, \tag{4}$$

$$\frac{\partial \rho_n}{\partial t} + \nabla \cdot (\rho_n \boldsymbol{v}) = \mathcal{R} - \mathcal{I}, \tag{5}$$

where $P^* \equiv P + (\boldsymbol{B} \cdot \boldsymbol{B})/2$ is the total pressure, P is the gas thermal pressure, $\mathcal{R}$ and $\mathcal{I}$ are the rates of recombination and ionization, measured in mass per unit time per unit volume, and $\mathcal{G}$ and $\mathcal{L}$ are the rates of radiative heating and cooling, measured in energy per unit time per unit volume. The total energy density E, not including atomic or molecular binding energies, is

$$E \equiv \epsilon + \rho \frac{\boldsymbol{v} \cdot \boldsymbol{v}}{2} + \frac{\boldsymbol{B} \cdot \boldsymbol{B}}{2}, \tag{6}$$

where ϵ is the gas thermal energy density. We adopt an ideal gas equation of state, so the pressure is $P = (\gamma - 1)\epsilon$ with $\gamma = 5/3$, corresponding to a monatomic gas. This choice of γ is appropriate because ionized gas is monatomic, and hydrogen in molecular clouds, while diatomic, is generally too cool to be able to access its rotational or vibrational degrees of freedom. It therefore acts as if it were monatomic for the purpose of dynamics. In addition to the evolution equations, the magnetic field is subject to the divergence-free constraint,

$$\nabla \cdot \boldsymbol{B} = 0. \tag{7}$$

Physically, equations (1), (2), and (4) represent the conservation of mass, momentum, and energy, equation (3) represents the approximation that the magnetic field is perfectly frozen in to the fluid, and equation (5) states that the mass of neutral gas is conserved by advection and is altered only by ionizations and recombinations. Equation (7) expresses the nonexistence of magnetic monopoles. Note that we have adopted a system of units in which the magnetic permeability $\mu = 1$; to convert to cgs units, we multiply our magnetic field strengths by a factor of $(4\pi)^{1/2}$.

2.2. Recombination and Ionization

To complete the evolution equations, we must specify the rates of the radiative processes: recombination, ionization, heating, and cooling. Recombination is the simplest of these. We adopt the on-the-spot approximation, whereby any recombination to the ground state of hydrogen is assumed to yield an ionizing photon with a very short mean free path, which will in turn cause another ionization at essentially the same location as the recombination (Osterbrock 1989). On the other hand, recombinations to excited states of hydrogen produce photons with long mean free paths, which are assumed to escape the H ɪɪ region. Thus, the recombination rate is approximately

$$\mathcal{R} = \mu_H \alpha^{(B)} n_e n_{H^+}, \tag{8}$$

where $\mu_H \approx 2.34 \times 10^{-24}$ g is the mean gas mass per hydrogen atom (for a gas of hydrogen and helium in the standard cosmic abundance), n_e and n_{H^+} are the number densities of electrons and hydrogen nuclei, respectively, and $\alpha^{(B)}$ is the case B recombination coefficient, $\alpha^{(B)} \approx 2.59 \times 10^{-13}(T/10^4 \text{ K})^{-0.7} \text{ cm}^{-3} \text{ s}^{-1}$ (Osterbrock 1989; Rijkhorst et al. 2006) at a gas temperature T. The H$^+$ and e^- number densities are $n_{H^+} = (\rho - \rho_n)/\mu_H$ and $n_e = n_{H^+} + \rho\alpha_C/(14\mu_H)$, where $\alpha_C \approx 3 \times 10^{-3}$ is the carbon abundance in the ISM (Sofia & Meyer 2001). This assumes that all carbon is singly ionized. Since the first ionization potential of carbon is smaller than that of hydrogen, this is likely to be the case in regions in which there is any ionized hydrogen present, which are the only regions for which we care about the recombination rate.

In the on-the-spot approximation, the ionization rate is relatively easy to compute as well, since we only need to solve the

radiative transfer equation along rays from ionizing sources. The ionizing flux at position $\boldsymbol{x}$ due to the nth ionizing source is

$$F_n = \frac{s_n}{4\pi|\boldsymbol{x} - \boldsymbol{x}_n|^2} e^{-\tau(\boldsymbol{x}, \boldsymbol{x}_n)}, \tag{9}$$

where

$$\tau(\boldsymbol{x}, \boldsymbol{x}_n) \approx \int_{\boldsymbol{x}_n}^{\boldsymbol{x}} (\sigma n_{\rm H} + \sigma_d n)\, d\ell \tag{10}$$

is the optical depth to ionizing photons along the path from $\boldsymbol{x}_n$ to $\boldsymbol{x}$, $n_{\rm H} = \rho_n/\mu_{\rm H}$ is the number density of neutral hydrogen atoms, $n = \rho/\mu_{\rm H}$ is the number density of hydrogen nuclei, $\sigma = 6.3 \times 10^{-18}$ cm^2 is the cross section for absorption of a photon at the ionization threshold by a neutral hydrogen atom, and $\sigma_d \approx 6.1 \times 10^{-22}$ cm^{-2} per hydrogen atom is the cross section for absorption of ionizing photons by dust grains (Bertoldi & Draine 1996). For the purposes of the tests that we report in this paper, which we will compare to analytic solutions derived for pure hydrogen, we set $\sigma_d = 0$. The photoionization rate is just the rate at which these photons are absorbed, so the total photoionization rate due to all sources at position $\boldsymbol{x}$ is

$$\mathcal{I}_{\rm ph} = \sigma \rho_n \sum_n \frac{s_n}{4\pi|\boldsymbol{x} - \boldsymbol{x}_n|^2} e^{-\tau(\boldsymbol{x}, \boldsymbol{x}_n)}. \tag{11}$$

We also compute the rate of ionizations due to collisions between neutral hydrogen atoms and electrons,

$$\mathcal{I}_{\rm coll} \approx k_{\rm coll} n_e \rho_n, \tag{12}$$

where the collisional ionization rate coefficient is (Tenorio-Tagle et al. 1986)

$$k_{\rm coll} \approx 5.84 \times 10^{-11} \sqrt{\frac{T}{\rm K}} e^{(-13.6\ {\rm eV})/(k_{\rm B}T)}\ {\rm cm^3\ s^{-1}}, \tag{13}$$

although for the gas densities and temperatures in the systems with which we are concerned here, this term is generally negligible. The total ionization rate is the sum of these two terms,

$$\mathcal{I} = \rho_n \left(\sigma \sum_n \frac{s_n}{4\pi|\boldsymbol{x} - \boldsymbol{x}_n|^2} e^{-\tau(\boldsymbol{x}, \boldsymbol{x}_n)} + k_{\rm coll} n_e \right). \tag{14}$$

It should be noted that Ritzerveld (2005) has recently pointed out that diffuse photons (i.e., those that have been absorbed and reemitted) can be as important as direct photons in determining the structure of H II regions, so the on-the-spot approximation that we adopt is potentially problematic. However, the importance of diffuse versus direct photons depends strongly on the degree of central concentration of the gas inside the H II region. Once an H II region becomes D type, which is the phase in which we are most interested from the standpoint of radiation hydrodynamic evolution, the gas inside the ionized region is nearly uniform in density. For this case, Ritzerveld (2005) finds that direct photons dominate diffuse ones over the majority of the H II region volume, and the on-the-spot approximation is reasonable. Nonetheless, the on-the-spot approximation does make the shadows that are cast by dense gas sharper than they should be in reality, and we must keep this limitation in mind.

2.3. Heating and Cooling

We can approximately determine the heating and cooling rates by breaking them up into heating and cooling associated with ionization and recombination of hydrogen atoms, and all other sources of heating and cooling. The heating rates due to ionization and recombination of hydrogen are quite simple in our monochromatic treatment of the ionizing radiation. In this approximation, each absorption of an ionizing photon delivers $e_\Gamma \approx 2.4$ eV of thermal energy to the gas (Whalen & Norman 2006), so the photoionization heating rate is

$$\mathcal{G}_{\rm ph} = e_\Gamma \sigma n_{\rm H} \sum_n \frac{s_n}{4\pi|\boldsymbol{x} - \boldsymbol{x}_n|^2} e^{-\tau(\boldsymbol{x}, \boldsymbol{x}_n)}. \tag{15}$$

Each recombination of a hydrogen atom allows photons to escape, leading to a loss of energy. A microphysical calculation of the resulting cooling rate gives $\mathcal{L}_{\rm rec} \approx \Lambda_{\rm rec} n_e n_{\rm H^+}$, where (Osterbrock 1989)

$$\Lambda_{\rm rec} \approx \left(6.1 \times 10^{-10}\ {\rm cm^3\ s^{-1}}\right) k_{\rm B} T \left(\frac{T}{\rm K}\right)^{-0.89} \tag{16}$$

for temperatures $T \gtrsim 100$ K, and recombination cooling falls to negligible values at lower temperatures.

For heating and cooling that are not directly due to ionization of hydrogen, we use simple optically thin heating and cooling curves. In molecular gas, we adopt the approximate cooling and heating functions of Koyama & Inutsuka (2002), and in partially ionized gas we compute the cooling rate following Osterbrock (1989). Our calculation includes the cooling by ion-electron collisions involving the first and second ionized states of O, N, and Ne, which are the dominant coolants in H II regions under solar metallicity conditions. We also include free-free cooling. Thus, the total heating and cooling rates are

$$\mathcal{G} = e_\Gamma \sigma n_{\rm H} \sum_n \frac{s_n}{4\pi|\boldsymbol{x} - \boldsymbol{x}_n|^2} e^{-\tau(\boldsymbol{x}, \boldsymbol{x}_n)} + n_{\rm H} \Gamma_{\rm KI}, \tag{17}$$

$$\mathcal{L} = \Lambda_{\rm KI}(T) n_{\rm H}^2 + \Lambda_{\rm rec}(T) n_e n_{\rm H^+} + \Lambda_{\rm ion-ff}(T) n_e n_{\rm H^+}, \tag{18}$$

where the subscripts "KI," "rec," and "ion-ff" indicate heating and cooling following the Koyama & Inutsuka curves, from recombinations, and from ion-neutral collisions and free-free emission in ionized gas. With these cooling curves, the equilibrium temperature in neutral gas at a density of $n_{\rm H} = 100$ cm^{-3} is approximately 11 K, and the equilibrium temperature in fully ionized gas is approximately 6400 K, with some weak density dependence.

2.4. Atomic Gas and PDRs

Here we justify our earlier statement that we can neglect PDRs when studying the large-scale dynamics of most H II regions. We do this using the analytic PDR models of Bertoldi & Draine (1996). Note that Bertoldi & Draine consider the case of convex fronts, in which a neutral cloud is bathed in ionizing radiation, whereas we are more concerned with concave fronts, for which the source is embedded within a molecular medium. To avoid complications arising from this, we will only consider the case of plane-parallel ionization fronts. However, the results should qualitatively generalize to our concave case.

First consider a system a short time after an ionizing source turns on, before a shock front forms. In this case, there is only a dissociation front and an ionization front. The propagation speed of each front is then determined solely by the rate at which ionizing or dissociating photons reach it. Let $s_{\rm Ly}$ and $s_{\rm FUV}$ be the stellar luminosities in units of the ionizing Lyman continuum ($\lambda < 912$ Å) and the dissociating FUV (912 Å $< \lambda < 1110$ Å) photons s^{-1} emitted by the ionizing source, and let $\tau_{\rm Ly}$ and $\tau_{\rm FUV}$ be the optical depths

from the source to the respective fronts. In this case an ionization front at a distance r_i from the source propagates at a velocity

$$v_i = \frac{s_{\mathrm{Ly}} e^{-\tau_{\mathrm{Ly}}}}{4\pi r_i^2 n_{\mathrm{H}}}, \tag{19}$$

where n_{H} is the number density of hydrogen nuclei outside the front. The dissociation front propagates at a rate

$$v_d \approx 0.15 \frac{s_{\mathrm{FUV}} e^{-\tau_{\mathrm{FUV}}}}{4\pi r_d^2 n_{\mathrm{H}}}, \tag{20}$$

where r_d is the front radius and the factor of 0.15 arises because only a minority of FUV photons actually produce dissociations. For stars of type O3 or earlier, i.e., those that drive significant H II regions, $s_{\mathrm{Ly}}/s_{\mathrm{FUV}} \approx 1$–2, so $v_i > v_d$ unless the Lyman continuum optical depth is large enough compared to the FUV optical depth that the fraction of Lyman continuum photons reaching the front is $\approx 8\%$–15% of the number of FUV photons reaching the front. This requires $\tau_{\mathrm{Ly}} - \tau_{\mathrm{FUV}} \gtrsim 1.9$–2.6.

Since the dust extinction cross section for FUV photons is larger than that for Lyman continuum photons by a factor of ≈ 1.25 (Bertoldi & Draine 1996), this can only happen if there is significant attenuation of Lyman continuum photons due to recombining gas within the H II region. At early times, before recombinations are significant compared to photoionizations, this means that $v_i > v_d$ and the ionization and dissociation fronts will be coincident. In this case, our neglect of the dissociation front is obviously not a problem.

As the ionization front radius approaches the Strömgren radius,

$$r_s = \left(\frac{3 s_{\mathrm{Ly}}}{4\pi \alpha^{(\mathrm{B})} n_{\mathrm{H}}^2} \right)^{1/3}, \tag{21}$$

recombinations within the ionized region produce a significant neutral population that raises τ_{Ly} but not τ_{FUV} and slows the expansion of the ionization front relative to the dissociation front. If this process continued indefinitely, then the dissociation front would eventually race ahead of the ionization front. However, slowing of the ionization front will also eventually allow a shock to form, converting it from R type to D type, and this will sweep gas up into a dense shell and modify the propagation speed of both the ionization and dissociation fronts. How much can the ionization front slow down before a shock forms? The condition for a shock to form is that the ionization front speed drops to of order the speed of a strong shock driven by the ionized gas, which is twice the ionized gas sound speed c_i. Thus, a shock forms when

$$\frac{s_{\mathrm{Ly}} e^{-\tau_{\mathrm{Ly}}}}{4\pi r_i^2 n_{\mathrm{H}}} \approx 2 c_i. \tag{22}$$

If we now approximate that $r_i \approx r_s$ once $\tau_{\mathrm{Ly}} \gtrsim 1$, it follows that

$$\tau_{\mathrm{Ly}} \approx \frac{1}{3} \ln \left(\frac{s_{\mathrm{Ly}} \alpha^{(\mathrm{B})^2} n_{\mathrm{H}}}{96\pi c_i^3} \right) \tag{23}$$

$$= 4.1 \ln \left(s_{49} n_2^2 \alpha_4^{(\mathrm{B})} c_6^{-3} \right), \tag{24}$$

where in the second step we have scaled to typical values for Galactic molecular clouds and H II regions: $s_{49} = s_{\mathrm{Ly}}/(10^{49} \mathrm{\ s}^{-1})$, $n_2 = n_{\mathrm{H}}/(100 \mathrm{\ cm}^{-3})$, $\alpha_4^{(\mathrm{B})} = \alpha^{(\mathrm{B})}/\alpha_{T=10^4 \mathrm{\ K}}^{(\mathrm{B})} = \alpha^{(\mathrm{B})}/(2.59 \times 10^{-13} \mathrm{\ cm}^{-3}$

s^{-1}), and $c_6 = c_i/(10^6 \mathrm{\ cm \ s}^{-1})$. Thus, the dissociation front will not begin moving faster than the ionization front until $\tau_{\mathrm{Ly}} \gtrsim 1.9$–2.6, but a shock will form once $\tau_{\mathrm{Ly}} \approx 4$.

Once a shock forms, we must consider how the width of the PDR compares to the width of the shocked layer to determine if the dissociation front can ever propagate ahead of the shock. In equilibrium, the column density of hydrogen atoms through a PDR is roughly (Bertoldi & Draine 1996)

$$N_{\mathrm{H}} \approx \sigma_{\mathrm{FUV}}^{-1} \left[1 + \left(2.7 + \phi_0^{-1} \right)^{4/3} \right]^{-1}, \tag{25}$$

where $\sigma_{\mathrm{FUV}} \approx 7.6 \times 10^{-22} \mathrm{\ cm}^{-2}$ is the attenuation cross section of dusty gas to FUV photons and the factor ϕ_0 is a dimensionless number describing the balance between H_2 formation and dissociation. Since the numerical factor involving ϕ_0 is at most 0.21, this implies that there is a maximum hydrogen column $N_{\mathrm{PDR}} \approx 2.8 \times 10^{20}$ through a PDR. Nonequilibrium PDRs are generally narrower than this. The column of hydrogen swept up in the dense shell around an ionized region shortly after it forms is roughly $n_{\mathrm{H}} r_s$. If we compare this to the maximum column density of a PDR, we find that the column of hydrogen $n_{\mathrm{H}} r_s$ present when the shock first forms exceeds the maximum column of a PDR whenever

$$s_{\mathrm{Ly}} \gtrsim \frac{4\pi}{3} \alpha^{(\mathrm{B})} \frac{N_{\mathrm{PDR}}^3}{n_{\mathrm{H}}} \tag{26}$$

$$= \left(2.3 \times 10^{47} \right) \frac{\alpha^{(\mathrm{B})}}{2.59 \times 10^{-13} \mathrm{\ cm}^{-3} \mathrm{\ s}^{-1}} \left(\frac{n_{\mathrm{H}}}{100 \mathrm{\ cm}^{-3}} \right)^{-1} \mathrm{\ s}^{-1}. \tag{27}$$

Thus, for $s_{\mathrm{Ly}} \gtrsim 10^{47} \mathrm{\ s}^{-1}$, which is the case for any strong ionizing source, the dissociation front will essentially always remain trapped between the ionization front and the shock front, except perhaps for a very brief interval when the front is changing from R type to D type and the Lyman continuum optical depth through the ionized region is about 2–4. As a result, the dissociation front will not have a significant effect on the dynamics of the ionized region, since the only gas it can affect will be gas that has already been shocked and is about to be ionized. The dissociation front will simply be passively carried along.

This can break down if we consider weak ionizing sources, or if we have a medium that is substantially less dense than $n_{\mathrm{H}} = 100 \mathrm{\ cm}^{-3}$, which is the typical mean density in a GMC. If either of these is the case, then a dissociation front may run ahead of the shock front and preheat the gas, thereby affecting the dynamics. In this case our neglect of the atomic gas becomes more problematic. However, even when preheating occurs, the dissociation front will only precede the shock front for a limited time, until the ionized region has swept up the critical column density N_{PDR}. Thus, even for weak ionizing sources or low-density regions, our approximation will only be invalid around the time when the front transitions from R type to D type.

Our simple analytic calculation is in good agreement with the detailed one-dimensional simulations of Hosokawa & Inutsuka (2005), who also find that the molecular hydrogen dissociation front around an expanding H II region becomes trapped between the shock and ionization fronts shortly after the front transitions from R type to D type.

3. SIMULATION ALGORITHM

3.1. Program Flow

We solve the evolution equations on a Cartesian grid with uniform cell spacing Δx using an operator-split approach, in which we

alternate conservative MHD updates with the source terms on the right-hand sides set to zero with radiation updates in which we add the source terms. In designing this update algorithm, we bear in mind two factors. First, due to the comparatively simple nature of the radiation update, computing a single explicit update using the radiation source terms is computationally much cheaper than computing a single magnetohydrodynamic update. For this reason, we choose to subcycle the radiation update relative to the MHD update. This effectively makes our radiation update step implicit, since during any given simulation time step we iterate the temperature and chemical state of each cell over many explicit update cycles.

Second, the overall update time step is restricted by both radiation and MHD. The time step must obey the Courant condition, but we also wish to impose a heating/cooling time step condition. If the energy in a cell changes too much between MHD updates, the corresponding large changes in pressure from one time step to the next may lead to inaccurate solutions. We therefore wish to limit the amount by which the energy in a cell can change between MHD updates, which in turn imposes a time step constraint associated with the rate of heating or cooling. Note that, in contrast, the MHD update is not directly affected by the neutral density ρ_n, so we need not impose a corresponding constraint on the overall simulation time step arising from changes in the chemical state of the gas.

In order to satisfy these constraints, we perform a two-step update procedure, illustrated schematically in Figure 1. At the end of time step n, in each cell we have a vector of quantities

$$q^n = \begin{pmatrix} \rho^n \\ p^n \\ B^n \\ E^n \\ \rho_n^n \end{pmatrix} \tag{28}$$

at time t^n. (Here p is the gas momentum, and we are glossing over the fact that in the Athena algorithm the magnetic field is actually stored at cell faces rather than cell centers.) We wish to determine the properties q^{n+1} at time t^{n+1}. We first update the gas energy E and neutral density ρ_n, using only the terms on the right-hand sides of equations (1)–(5). This update covers a time $\Delta t^n = \min(\Delta t_C, \Delta t_\epsilon)$, where Δt_C is the Courant time step for the current configuration and Δt_ϵ is a time step imposed by a requirement that the internal energy of a cell not change by too much in each time step. This gives us an intermediate vector $q^{n,*}$. Since we do not want to have to perform an MHD update every time we do a radiation update, the radiation update procedure is iterative. Therefore $E^{n,*}$ is not simply $E^n + \Delta t^n (\mathcal{G} - \mathcal{L})^n$, and similarly for $\rho_n^{n,*}$. We give a step-by-step description of the radiation update procedure in § 3.2, and we will defer formal definition of Δt_ϵ, $E^{n,*}$, and $\rho_n^{n,*}$ until then.

Second, we perform an MHD update to gas quantities starting from state $q^{n,*}$, advancing the conservation equations (eqs. [1]–[5]) with the source terms on the right-hand sides set to zero. This advance covers the same time Δt^n as the radiation update. We perform this update using the standard Athena explicit MHD update procedure, which is described in Gardiner & Stone (2005, 2006), with the trivial addition of the passive scalar advection given in equation (5) for the neutral density. The overall program flow is independent of the details of the MHD update algorithm. For the runs we report in this paper, we configure Athena to use its directionally unsplit 6-solve upwind constrained-transport method, Roe's Riemann solver with H correction, and second-order spatial accuracy. We do not describe these methods in greater detail
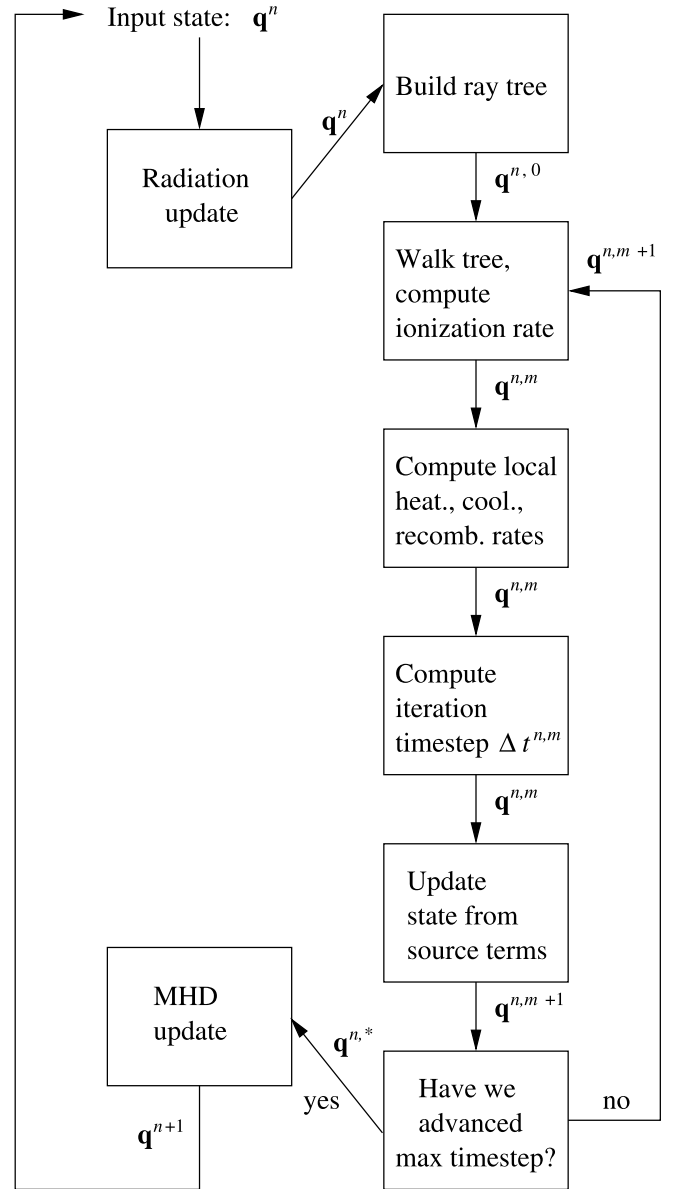

Fig. 1.—Schematic diagram of the main update loop of our method.

here, and instead we refer readers to the Gardiner & Stone papers. At the end of this step, we have updated q^n for all the terms in the evolution equations (eqs. [1]–[5]), so the final state is q^{n+1}, and we may begin the next time step.

It should be noted that our update cycle is quite similar to that of Abel et al. (1999) and Whalen & Norman (2006).

3.2. Radiation Update Method

Our radiation update algorithm combines elements of the approaches of Abel & Wandelt (2002) and Whalen & Norman (2006), together with some novel techniques that, through the tests we describe in § 4, we found to be necessary to achieve high accuracy. Our procedure to advance the input state q^n to the intermediate state $q^{n,*}$ is as follows.

We do two preparatory steps before beginning a radiation update. First, we compute the Courant time step Δt_C^n for the configuration q^n. Second, we construct a tree of rays, following the prescription of Abel & Wandelt (2002), for each radiation source. We will not review the full formalism of this method here, but a

summary of its relevant properties is that from each source, one sends out a tree of rays in directions that correspond to the centers of pixels in the HEALPix pixelization scheme (Górski et al. 2005). At its coarsest level, the tree contains 12 rays uniformly spaced over the unit sphere, but the scheme is hierarchical, such that any ray may be subdivided into four child rays an arbitrary number of times. After the lth such division, the unit sphere is discretized into $12(4^l)$ pixels, so that each pixel covers a solid angle $\Omega = (\pi/3)4^{-l}$ sr. As we follow rays outward, we subdivide them to guarantee that at a distance r from the source, $\Omega \leq (1/2)\Delta x^2/(4\pi r^2)$; i.e., that the solid angle that a given ray represents is always smaller than the solid angle subtended by a cell of the Cartesian grid by at least a factor of 2. We always divide each ray at least twice, so we never have fewer than 192 rays.

Along each ray we construct a list of the cells through which that ray passes and the length of the ray segment that lies within each cell. If we are running a calculation in parallel, so that the computational domain is subdivided into a series of processor domains that are contiguous rectangular blocks, we first construct the list of processor domains through which a ray passes. On each processor we only store information about rays for which either the ray, its parent, or one of its children intersects that processor's domain. This enables us to pass information from one processor to the next as we walk outward along the ray tree without storing the full tree on each processor.

If the source is moving, we repeat this procedure every time the source moves, which is generally every time we perform an MHD update. For a source that is fixed in space, we build a new tree for every fifth MHD time step. Each time we build the tree, we rotate it into a random orientation to minimize errors due to discretization in angle. We discuss this procedure, and why it is necessary, in more detail in § 4.3.3.

With these preparations done, we begin the iterative update procedure. We compute the ionizing flux passing through every cell by walking outward along the rays, following the procedure outlined in Abel & Wandelt (2002). Suppose a particular ray representing a solid angle Ω starts at a distance r_0 from the source, and that the flux of ionizing photons at this point is F_0. The ray passes through a sequence of I cells, and we define ℓ_i as the length of the ray segment that intersects the ith cell, $r_i = r_0 + \sum_{j=0}^{i-1}\ell_i$ as the distance from the source to cell i, and F_i as the radiation flux reaching cell i. The optical depth through cell i is $\tau_i = (\sigma n_H + \sigma_d n)\ell_i$, where n_H and n are the number densities of neutral hydrogen atoms and hydrogen nuclei in the cell. The flux reaching the next cell is

$$F_{i+1} = F_i e^{-\tau_i}\left(\frac{r_{i+1}}{r_i}\right)^2, \qquad (29)$$

and conservation of ionizing photons therefore implies that the number of ionizations per unit time in cell i due to photons moving along this ray is $F_i(1 - e^{-\tau_i})\Omega r_i^2$. The corresponding rates of mass photoionization per unit volume and photoionization heating per unit volume are

$$\mathcal{I}_{ph} = F_i(1 - e^{-\tau_i})\Omega r_i^2 \frac{\mu_H}{\Delta x^3}, \qquad (30)$$

$$\mathcal{G}_{ph} = F_i(1 - e^{-\tau_i})\Omega r_i^2 \frac{e_\Gamma}{\Delta x^3}. \qquad (31)$$

We use this procedure to trace the photons along each ray, starting from the 12 coarsest ones. Once we have reached the end of a ray, we pass the flux F_I that escapes the final cell on to its child rays and repeat this procedure for them. We continue until we either reach the edge of the computational domain or until such

a small fraction of the photons remain (we use 10^{-3}) that it is no longer worthwhile to continue following them. The total mass photoionization rate $\mathcal{I}_{ph}^{n,m}$ and photoionization heating rate $\mathcal{G}_{ph}^{n,m}$ in a cell at time $t^{n,m}$ is simply the sum of the rates contributed by all the rays that intersect it, as determined by equations (30) and (31), for the gas properties at that time ($\boldsymbol{q}^{n,m}$).

Once we have determined the rates of change of the neutral mass and energy due to photoionization, we determine the rates due to recombination, collisional ionization, and optically thin heating and cooling, using equations (8), (12), (16), (17), and (18). These calculations are purely local, and we simply perform them using the current state of the gas, $\boldsymbol{q}^{n,m}$. For reasons that we discuss in § 3.3, we set the heating and cooling rate due to molecular transitions to zero in cells with ionization fractions between 1% and 99%. Adding the rates from local processes together with the photoionization rates gives the total rate of change of the gas energy, $(\mathcal{G} - \mathcal{L})^{n,m}$, and the neutral density, $(\mathcal{R} - \mathcal{I})^{n,m}$, at time $t^{n,m}$.

Once we know the rates of change of the energy and the neutral density, we compute the time step for this iteration,

$$\Delta t^{n,m} = \frac{1}{10} \max\left[\frac{\epsilon^{n,m}}{(\mathcal{G} - \mathcal{L})^{n,m}}, \frac{\rho_n^{n,m}}{(\mathcal{R} - \mathcal{I})^{n,m}}, \frac{(\rho - \rho_n)^{n,m}}{(\mathcal{R} - \mathcal{I})^{n,m}}\right], \qquad (32)$$

where the maximum is over the three quantities in the parentheses computed over all cells. Thus, the iteration time step is chosen so that the maximum fractional change in the internal energy ϵ, neutral density ρ_n, or ion density $\rho - \rho_n$ in any cell is 10%. We also want to ensure that the total radiation time by which we advance is no larger than the Courant time step in the initial configuration, so if $t^{n,m} + \Delta t^{n,m} > t_n + \Delta t_C^n$, we set $\Delta t^{n,m} = t_n + \Delta t_C^n - t^{n,m}$.

Once we have fixed the time step $\Delta t^{n,m}$ for this iteration of the radiation update, we actually perform the update to find the new state

$$\boldsymbol{q}^{n,m+1} = \boldsymbol{q}^{n,m} + \Delta t^{n,m}\begin{pmatrix} 0 \\ \mathbf{0} \\ \mathbf{0} \\ \mathcal{G} - \mathcal{L} \\ \mathcal{R} - \mathcal{I} \end{pmatrix}^{n,m}. \qquad (33)$$

We terminate iteration if one of the following conditions is satisfied: (1) we have advanced over the total time determined by the Courant time step in the initial configuration, i.e., $t^{n,m+1} = t^n + \Delta t_C^n$; (2) we have advanced over a total time that is equal to or larger than the Courant time step in the *current* configuration, i.e., $t^{n,m+1} \geq t^n + \Delta t_C^{n,m}$; or (3) the internal energy of a cell has changed by more than a factor of f from the its initial value, i.e., $\max(\epsilon^{n,m}/\epsilon^n, \epsilon^n/\epsilon^{n,m}) > f$, where again the maximum is over both quantities and over all cells. If none of these conditions are met, we do another iteration, starting with a calculation of the photoionization rate by traversing the tree.

The state when iteration terminates is the intermediate state that we use as an input to the conservative MHD update. Thus, if iteration terminates after M cycles, then

$$\boldsymbol{q}^{n,*} = \boldsymbol{q}^n + \sum_{m=0}^{M-1} \Delta t^{n,m}\begin{pmatrix} 0 \\ \mathbf{0} \\ \mathbf{0} \\ \mathcal{G} - \mathcal{L} \\ \mathcal{R} - \mathcal{I} \end{pmatrix}^{n,m}. \qquad (34)$$

The time step for the MHD update is then set to $\Delta t = t^{n,M} - t^n$, the time for which we have computed the radiation update. If the iteration terminates due to condition (1), this will be the Courant time step of the initial state, $\Delta t = \Delta t_C^n$. If it terminates due to condition (2), this will be $\Delta t \approx \Delta t_C^{n,M}$, the Courant time step of the intermediate, postradiation state. If it terminates due to condition (3), the time step will be $\Delta t = \Delta t_\epsilon$, the time it takes for the internal energy to change by a factor of f in the cell that changes by the largest factor. In practice, the third condition is usually the most restrictive, since as the ionization front advances over a single cell, the internal energy of that cell changes by a factor of $\sim 10^3$, and we find that $f \ll 10^3$ is required to obtain good solutions. We present tests to determine what value of f produces acceptable results for simple radiation hydrodynamics problems in § 4.3.2.

3.3. *Cooling in Mixed Cells*

As part of our algorithm, we set Λ_{KI} and Γ_{KI}, the cooling and heating rates due to molecular processes, to zero in cells with ionization fractions between 1% and 99%. We do this in order to prevent a problem with overcooling in mixed-ionization cells. To understand why this is necessary, let us first consider what happens at an ionization front in reality. In the fully molecular gas outside the front, heating due to cosmic rays and interstellar UV and cooling due to molecular emission balance and the gas is effectively isothermal; in the fully ionized gas on the other side, heating by ionizing photons and cooling by recombination and forbidden line emission also balance to keep the gas isothermal. The ionized and molecular regions are separated by a PDR and an ionization front. The true thickness of the ionization front is of order the mean free path of an ionizing photon, which in a gas at a typical molecular cloud density of $n_H = 100$ cm^{-3} is $(\sigma n_H)^{-1} \approx 100$ AU. Once an expanding ionization front becomes D type and sweeps up a dense shell of material, the thickness will be even smaller. The ionization fraction will be between 1% and 99% only within this thin layer. Because the hot ionized gas on one side of the front and the cool molecular gas on the other side are physically separated and there is no mechanism to efficiently transport heat across the front, gas that gains energy from ionizing photons cannot lose it via molecular cooling.

Now, however, consider what happens in a simulation with finite resolution. If one simulates a region that is parsecs or more in size, the $\lesssim 100$ AU width of an ionization front will generally be much smaller than the size of a computational cell. As a result, the ionization front will be smeared out into a width comparable to the cell size, and this numerical mixing will supply a mechanism to transport heat from ionized gas into molecular gas. This can lead to a dramatic overestimate of the energy loss rate, because molecular emission is so efficient. In effect, the rate-limiting step in cooling will become how rapidly energy can be transported into neutral gas by numerical mixing. We demonstrate this effect in a real simulation in § 4.3.4.

Ideally, one would solve this problem by either resolving the width of the ionization front (or at least coming close enough to render the overcooling region too small to do much damage) or using an interface-tracking technique to follow the position of the front on a subgrid scale. However, the former option is prohibitively expensive in three dimensions (although it is feasible in two dimensions; e.g., Arthur & Hoare 2006), and the latter is extremely cumbersome in a three-dimensional simulation. Therefore, we adopt a much simpler solution: since we know physically that cells at an ionization front should contribute negligibly to the overall cooling rate, we simply set the molecular cooling rate in these cells to zero. This leads to a very minor underestimate of the cooling rate, but the error is small because in reality the front has a

tiny volume and therefore should contribute negligibly to the overall cooling rate. As we show in § 4.3.4, this procedure enables us to match analytic solutions to very high precision, whereas if we leave the cooling on in mixed cells, we find large errors relative to analytic solutions.

Our method does assume that the boundary between the ionized and molecular gas can be described well as a single, simple front. If instead the molecular and ionized gas are mixed into a two-phase medium with a complex interface that has structures that are small compared to a computational cell (for example, dense clumps of cold molecular gas surrounded by a hot, ionized interclump medium), then our approach of simply turning off molecular cooling will fail. In this case, one would either need to resolve the complex interface or use a detailed subgrid model of heating and cooling to account for the two-phase nature of the medium.

4. ACCURACY AND CONVERGENCE TESTS

Here we present tests of our code against analytic solutions. Since we are not aware of any analytic solutions for ionizing radiative transfer with magnetohydrodynamics in more than one dimension, all of the tests that we present in this section are hydrodynamic rather than magnetohydrodynamic; i.e., $\boldsymbol{B} = 0$. We present magnetohydrodynamic results in § 5.

4.1. *R-Type Ionization Fronts without Recombination*

A first, basic test of our implementation is the propagation of a spherical ionization front in a medium in which there is no hydrodynamic evolution and no recombinations. Physically, this is the type of expansion that should occur at the very earliest stages of an H II region's life, before there has been time for either hydrodynamic motions or for a significant number of recombinations to occur. In this case, a source with ionizing luminosity s placed in a uniform medium with density ρ produces an ionized region whose radius is simply set by equating the number of ionizing photons produced up to a time t with the number of hydrogen atoms within the radius r_i. Thus, we should find that

$$r_i = \left(\frac{3\mu_H s}{4\pi\rho}\right)^{1/3} t^{1/3}. \qquad (35)$$

This problem is effectively a test of photon conservation in our code. Since our radiation method is in principle exactly photon-conserving, we should be able to compute the position of the ionization front accurately to the resolution of our computational grid. We perform a test with a source that has an ionizing luminosity of $s = 4.0 \times 10^{49}$ ionizing photons s^{-1} that is placed in a uniform medium with a density of $\rho = 2.34 \times 10^{-22}$ g cm^{-3} ($n_H = 100$ cm^{-3}) and an initial temperature of 20 K. The source is at the center of a computational domain that runs from -10 pc to 10 pc in each direction. For this test we disable the MHD update in our code and set the rates for all radiative processes except photoionization to zero. We use a resolution of 64^3 cells.

Figure 2 shows the computed radius of the ionization front versus time in our test. As the plot shows, the computed solution agrees with the analytic prediction to better than 1%, which is an error that is much smaller than the size of a computational cell. Figure 3 shows the radial profile of the ionization fraction versus radius for every cell in our computational domain at a time of 10^{11} s (3.17 kyr), and Figure 4 shows the neutral fraction in an equatorial slice through the domain. As the plots show, the ionization front is entirely confined to cells at the analytically computed front radius at this time, which is located at 6.9 pc. Thus, our code passes this test quite well.
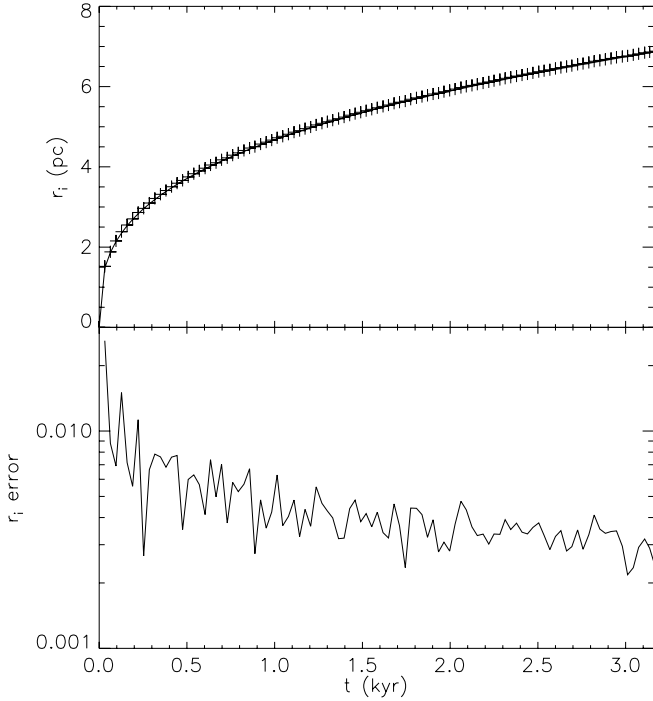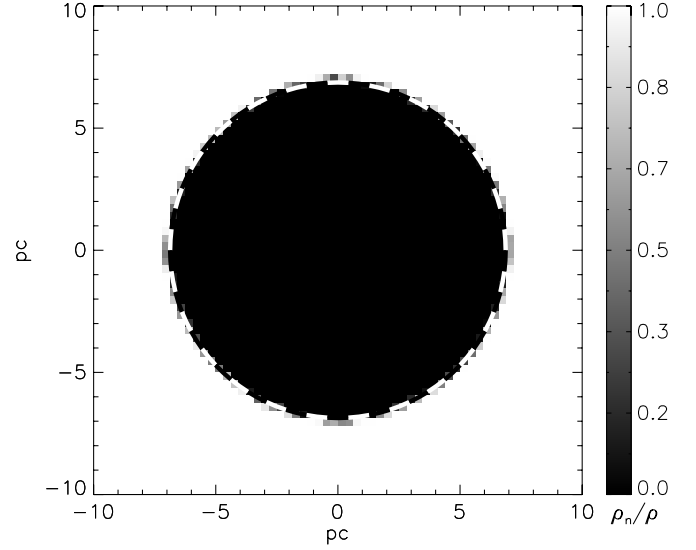
FIG. 4.—Ratio of neutral to total density ρ_n/ρ in an equatorial slice through our 3D computational domain in a simulation of an R-type ionization front without recombinations. The time shown is the same as that in Fig. 3. The dashed line shows the analytically computed front position.

4.2. R-Type Ionization Fronts with Recombination

A next test of our code is for the case in which there are still no motions, but recombinations do occur. Physically, this represents the next phase in the expansion of an R-type ionization front, when the recombination rate becomes competitive with the photoionization rate, but before the excess pressure in the ionized gas causes the H II region to begin hydrodynamically expanding. If the recombination coefficient $\alpha^{(B)}$ is constant, this problem has an analytic solution: at time t, the radius of the ionization front is

$$r_i = r_s \left(1 - e^{-t/t_{rec}} \right)^{1/3},$$ (36)

where

$$r_s = \left(\frac{3s\mu_H^2}{4\pi\alpha^{(B)}\rho^2} \right)^{1/3}$$ (37)

is the Strömgren radius and

$$t_{rec} = \frac{\mu_H}{\alpha^{(B)}\rho}$$ (38)

is the recombination timescale for the gas. At early times this solution matches the case of power-law expansion discussed in § 4.1, but as times $t \gtrsim t_{rec}$, the ionization front slows down due to recombinations. At times $t \gg t_{rec}$, the front radius approaches r_s and the front velocity drops to zero.

For any realistic cooling curve, for which the heating and cooling rates depend on the density, the radiation flux, and the ionization state, there will be some small temperature variation in time and space within the H II region, so $\alpha^{(B)}$ will not be exactly constant and the analytic solution will not apply exactly. To enable us to make comparisons with the analytic solution as well as possible, and therefore to make the strongest possible evaluation of the code, for the purposes of this test we set $\alpha^{(B)} = 2.59 \times 10^{-13}$ cm^{-3} s^{-1} independent of temperature.

We run this test using the same initial conditions and resolution as in § 4.1, which gives $r_s = 5.0$ pc and $t_{rec} = 1.2$ kyr for our
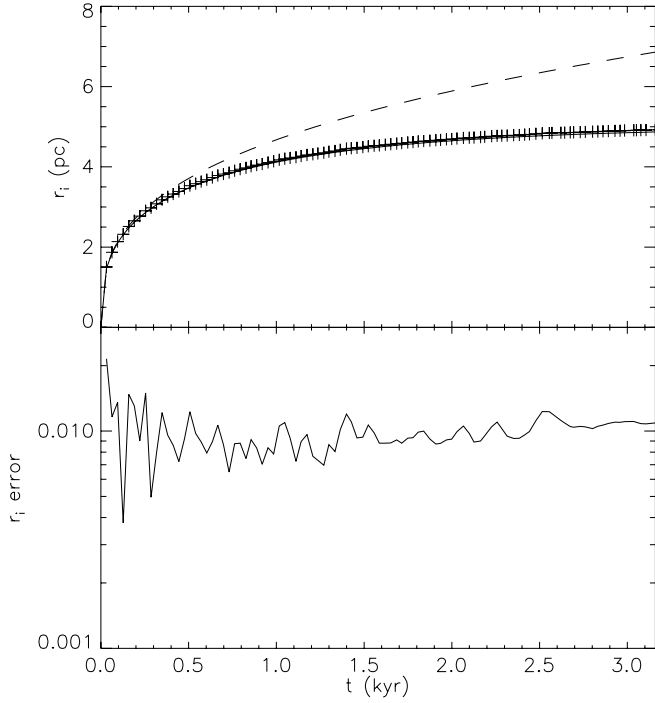

FIG. 2.—Ionization front radius r_i vs. time t (*top*), and error in simulated r_i relative to the analytic value vs. time (*bottom*), in a simulation of an R-type ionization front with no recombinations. The top panel shows the radius computed in the simulation (*plus signs*) and the analytically computed radius (*solid line*). We define the ionization front radius as the mean radius of the centers of all cells for which $0.01 < \rho_n/\rho < 0.99$. The error is defined as $|r_{sim} - r_{analyt}|/r_{analyt}$, where r_{sim} is the radius found in the simulation and r_{analyt} is the analytically computed value.

FIG. 3.—Ratio of neutral to total density ρ_n/ρ vs. radius r for a subset of cells in our computational domain at a time of 10^{11} s (3.17 kyr) in a simulation of an R-type ionization front without recombinations. The scale bar above the location of the ionization front shows a range of $\pm(\sqrt{3}/2)\Delta x$ about the analytically computed radius of the ionization front at this time. The minimum ratio of 10^{-20} in the ionized region is set by a numerical floor in our code.

chosen value of $\alpha^{(B)}$. For this test, we again disable the MHD update step in our code, but we allow the entire radiation module to operate normally. We run the test for 10^{11} s (3.17 kyr), which is more than $2t_{\rm rec}$. Figure 5 shows the position of the ionization front versus time, computed as in § 4.1, and Figure 6 shows the radial profile of the neutral density at a time of 10^{11} s. As we found in § 4.1, the ionization front is entirely confined to a radial extent of a single cell, so the front is as spherical as possible, given our Cartesian grid. The accuracy of the front position is ~1%, better than a single cell.

4.3. D-Type Ionization Fronts

The tests described in §§ 4.1 and 4.2 demonstrate that our radiation module produces correct results on its own, and the tests in Gardiner & Stone (2005, 2006) show that the MHD module works well. However, we must also test how the two modules work together. We therefore consider the problem of the hydrodynamic evolution of a uniform neutral medium with initial density ρ_n and temperature T_n within which there is an ionizing source. This problem has a well-known solution. At early times there is an R-type ionization front: the source ionizes the gas out to a radius $r_i \approx r_s$ and heats it to $T_i \approx 6400$ K, but the gas does not move, and the ionization front is not preceded by a shock. However, the ionized region is overpressured relative to the surrounding gas, and it expands on a sound crossing timescale, $t_s \equiv r_s/c_i$, where $c_i \approx 8 \times 10^5$ cm s^{-1} is the sound speed in the ionized gas. As it expands, it sweeps up a dense shell of neutral material, forming a D-type ionization front. The expansion is subsonic with respect to the ionized gas, so the ionized region has an almost uniform density and temperature. Ionization balance requires that this density must vary as $\rho_i \propto r_i^{-3/2}$, so the shell is driven outward by a
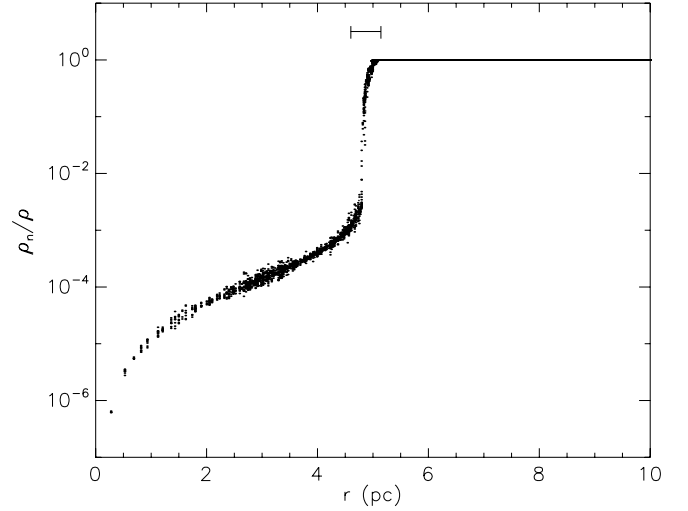

Fig. 6.—Ratio of neutral to total density ρ_n/ρ vs. radius r for a subsample of cells in our computational domain at a time of 10^{11} s (3.17 kyr) in a simulation of an R-type ionization front with recombinations. The scale bar above the location of the ionization front shows a range of $\pm(\sqrt{3}/2)\Delta x$ about the analytically computed radius of the ionization front at this time.

pressure of $p = \rho_i c_i^2 \propto r_i^{-3/2}$. In the limit in which the pressure in the ionized region that is driving the shell greatly exceeds the thermal pressure in the ambient medium and the neutral mass in the shell greatly exceeds the mass within the ionized region, conservation of momentum for the shell implies that after a time t it expands to a radius of (Spitzer 1978)

$$r_{\rm sh} = r_s \left(1 + \frac{7t}{4t_s}\right)^{4/7}. \tag{39}$$

The thickness of the shell relative to its radius is of order $(c_n/v_{\rm sh})^2$, where c_n is the isothermal sound speed in the neutral gas and $v_{\rm sh} = c_i[1 + (7/4)t/t_s]^{-3/7}$ is the velocity of the shell. As with the analytic solution for the R-type ionization front, this strictly applies only if the ionized gas sound speed c_i is constant. For a realistic cooling curve, we expect this to be nearly but not perfectly true.

We run simulations with an initial density of $\rho = 2.34 \times 10^{-22}$ g cm^{-3}, a temperature of $T = 11$ K, and an ionizing source of luminosity $s = 4.0 \times 10^{46}$ s^{-1}. The source is placed at the center of a box running from -10 pc to 10 pc in every direction. For these values of ρ and s, and assuming an ionized region temperature of roughly 8.0×10^3 K, we have $r_s = 0.5$ pc and $t_s = 0.061$ Myr. We expect that the similarity solution will apply for times $1 \ll t/t_s \ll 1000$. The lower limit on the time is imposed by the requirement that the mass in the swept-up shell be much larger than the mass in the interior of the H II region, and the upper time limit comes from the requirement that the pressure in the H II region greatly exceed the ambient pressure, or equivalently, that the thickness of the shell be much smaller than its radius. We run all our tests for 1 Myr, or $16.3t_s$.

4.3.1. Comparison to the Analytic Solution

We begin with simulations at resolutions of 64^3, 128^3, and 256^3, using a factor of $f = 4$ for our heating time step constraint (i.e., we allow the internal energy of a cell to change by up to a factor of 4 between hydrodynamic updates). Figure 7 shows the radial distribution of densities, temperatures, neutral fractions, and isothermal sound speeds in the gas in the simulations after 1 Myr. As we expect, the density and temperature in the ionized region are
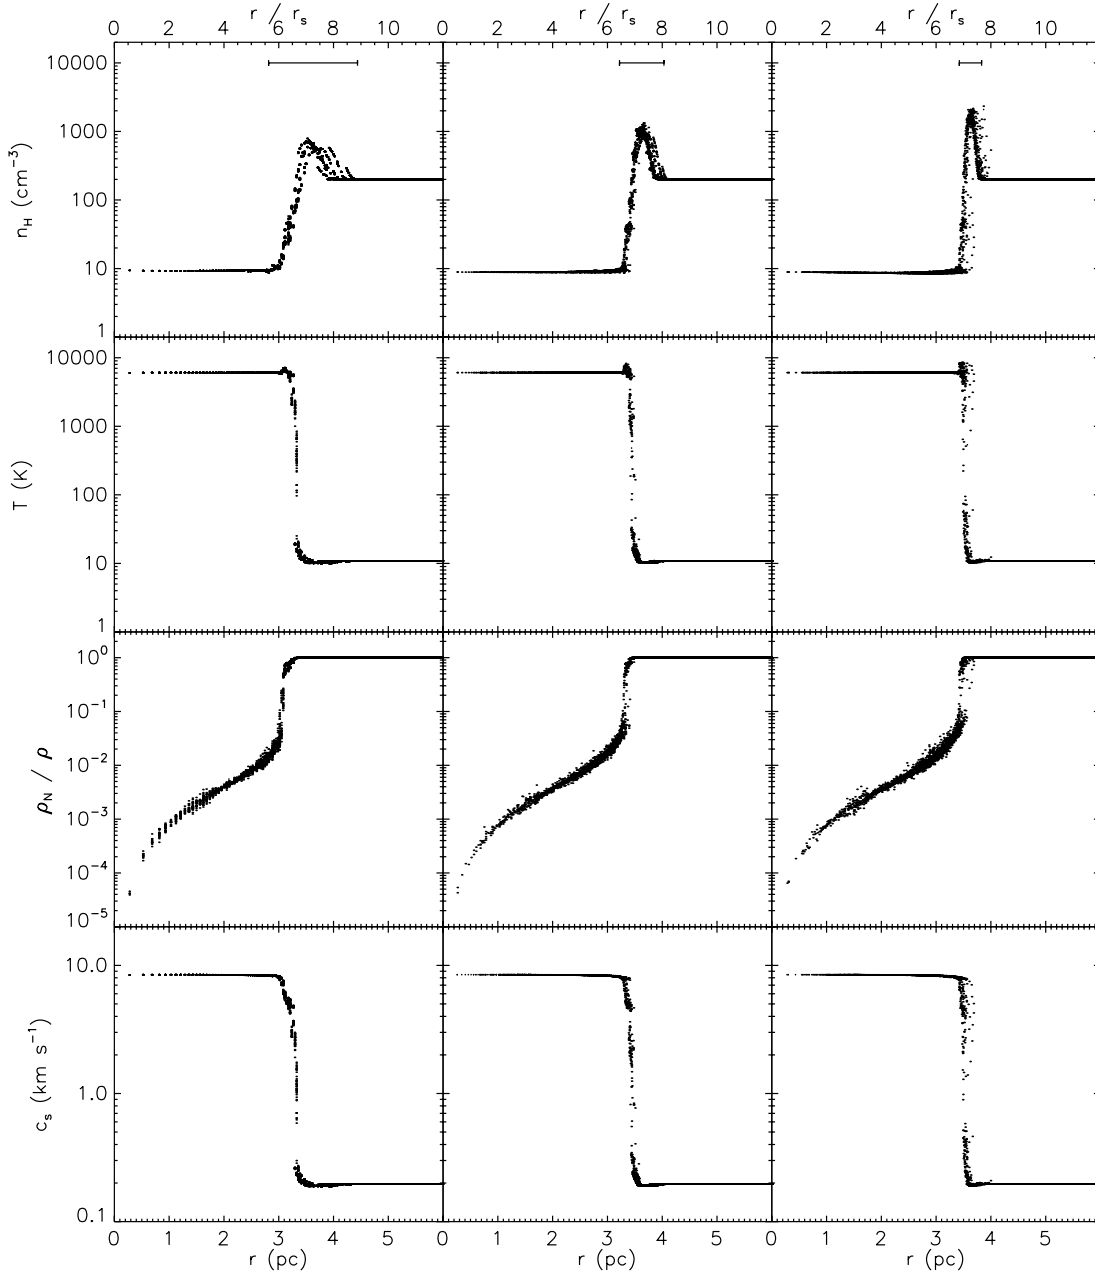
FIG. 7.—Hydrogen number density, temperature, neutral fraction, and isothermal sound speed vs. radius for a selection of computational cells in our simulation of a D-type ionization front, at resolutions of 64^3 (*left*), 128^3 (*middle*), and 256^3 (*right*). We show only a subsample of the grid points so as to minimize clutter and file size. The points are plotted at a time of 1 Myr ($17.7t_s$) after the start of the simulation. The scale bars in the top panels indicate a range of $\pm 3(\sqrt{3}/2)\Delta x$ about the analytically computed radius of the dense shell at this time, to indicate a width of three cells, about the expected width of a shock due to numerical viscosity.

almost constant. The density peaks at the position of the overdense shell, but the shell is spread over a few cells, so its width decreases with resolution. This spreading is to be expected. The true physical thickness of the shell at this time should be roughly 1% of $r_{\rm sh}$, smaller than the size of a single computational cell. However, our code has some numerical viscosity, particularly for motions that are not aligned with the computational grid, and this spreads shocks over ~3 cells. As Figure 8 shows, despite this spreading, the front remains round at each resolution.

We next compute the radius of the dense shell as a function of time in our simulation. We do this by identifying all the cells whose densities are at least 10% larger than the initial density and then taking the average of their radii. Using an overdensity cutoff different from 10% modifies the results in detail, but not qualitatively. We

compare this measurement to the analytic solution in Figure 9. As the plot shows, the radius of our simulated shell matches the analytically computed radius at late times quite well in all the runs. The error is always smaller than 1 cell in size, which is the best that can be expected when using a Cartesian grid with finite numerical viscosity. This error is ~3% of the analytically computed radius at the lowest resolution and less than 1% at the highest resolution. The error generally decreases with time, both because the ratio of the radius to a cell size is increasing and because the analytic solution is becoming a better and better approximation as time increases.

4.3.2. Heating Time Step Convergence Tests

We have shown that our method can match the analytic solution for the propagation of a D-type ionization front with a heating
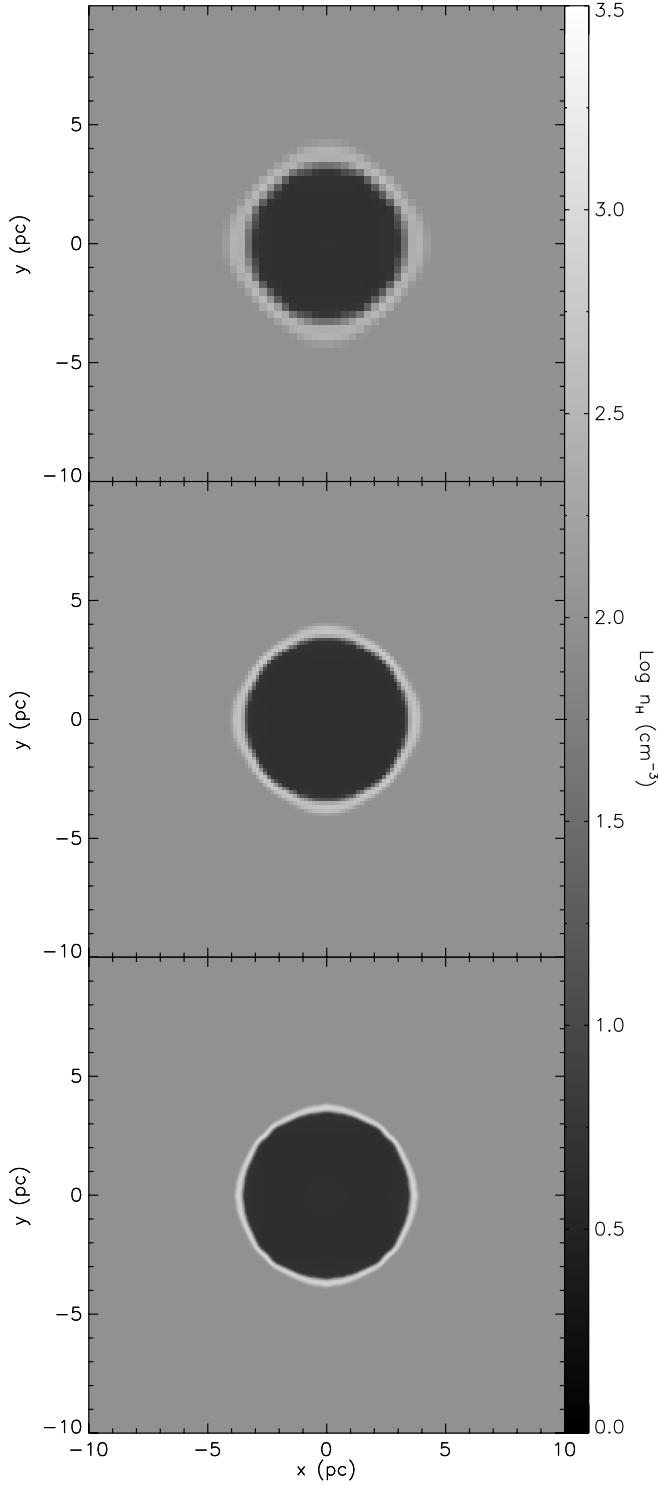
Fig. 8.—Slices through the *x-y* plane, showing the logarithm of the hydrogen number density in a simulation of a D-type ionization front 1 Myr (17.7t_s) after the start of the simulation, at resolutions of 64^3 (*top*), 128^3 (*middle*), and 256^3 (*bottom*). The slices in the *x-z* and *y-z* planes are nearly identical.
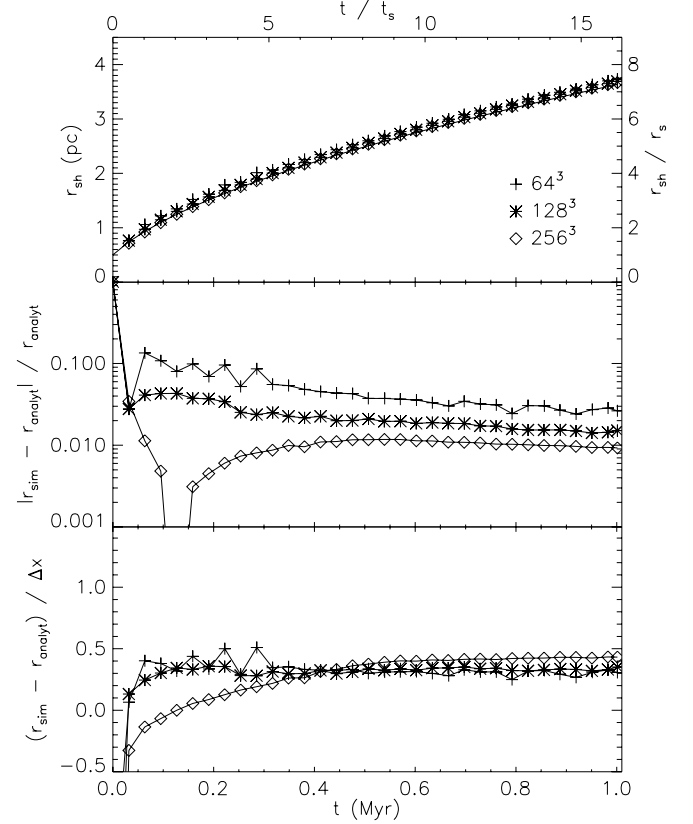

Fig. 9.—Radius of the dense shell around an H II region vs. time (*top*), computed analytically (*solid line*) and by simulations at resolutions of 64^3 (*plus signs*), 128^3 (*asterisks*), and 256^3 (*diamonds*). We also show the difference between the simulated and analytically computed shell radii, r_{sim} and r_{analyt}, normalized to the analytic radius (*middle*) and normalized to the size of a computational cell (*bottom*).

time step constraint of $f = 4$, meaning that we allow the pressure in a cell to change by at most a factor of 4 between hydrodynamic updates. However, we would like to explore a range of heating time step constraints to determine if the accuracy of the solution is affected by the choice of f. We therefore rerun our 128^3 simulation with $f = 2$, $f = 10$, and $f = 100$ to see if the quality of the solution changes significantly. The choice of $f = 100$ corresponds

to essentially no constraint on the hydrodynamic time step other than the ordinary Courant condition. We wish to use as large a value of f as possible, because that will minimize the computational cost by enabling us to take fewer of the more expensive MHD updates per radiation update. Of course, if f is too large, then the cost of the radiation update will dominate the total computation cost and there will be no additional advantage in increasing f. Where exactly this crossover occurs depends on the number of processors and the geometry of the ionization front.

Figure 10 shows the radius of the shell versus time as computed in our simulations with varying values of f and as computed analytically. The results show that the radius approaches the analytic solution somewhat for all values of f, but that the error is larger and the convergence slower for larger values of f. This suggests that numerical methods in which the temperature change between hydrodynamic updates is unconstrained (e.g., Dale et al. 2005; Mellema et al. 2006a, 2006b) or is limited only to values of $f \gg 10$ (e.g., Mac Low et al. 2007) may produce quantitatively incorrect results for the expansion rates of D-type ionization fronts, although the error is likely to be only a few cells. This problem may not affect methods that implicitly update the hydrodynamics and the radiation together (e.g., Miniati & Colella 2007), rather than in an operator-split fashion, but we are unaware of any such methods for ionizing radiation transport, as opposed to local heating and cooling functions.

Note, however, that even with large values of f, the error does decrease in time. This result is easy to understand intuitively. Once the front is well into the D-type phase, the Courant time step is set by the sound speed in the ionized region, which is roughly constant.
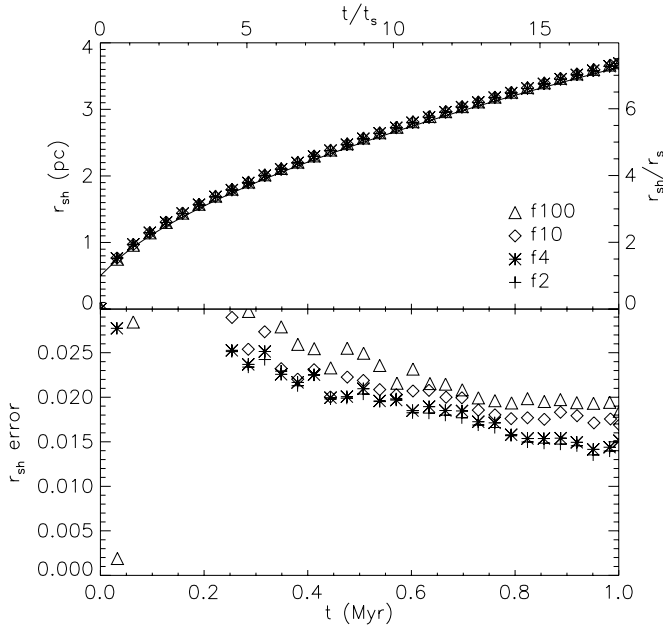
FIG. 10.—Radius of the dense shell around an H II region vs. time (*top*), computed analytically (*solid line*) and by simulations at a resolution of 128^3, with $f = 2$ (*plus signs*), $f = 4$ (*asterisks*), $f = 10$ (*diamonds*), and $f = 100$ (*triangles*). We also show the fractional error in the position of the shell relative to the similarity solution (*bottom*), defined as $|r_{sim} - r_{analyt}|/r_{analyt}$, where r_{sim} is the simulated radius and r_{analyt} is the analytically computed radius.

As the front sweeps up more material and slows down, the rate at which gas is heated from the initial temperature to the ionized temperature decreases, so on average, the fractional change in cell temperature per fixed Courant time step is a decreasing function of time. Thus, the error one makes by allowing a large change in gas temperature per Courant time step decreases with time. This means that codes with unconstrained temperature changes per hydrodynamic update are likely to be most inaccurate at early times, when fronts are moving at speeds close to the ionized sound speed, and most accurate at late times, when fronts move slowly.

4.3.3. *Ray Rotation*

To demonstrate why it is necessary to rotate the orientation of the rays in our calculation regularly, we repeat our 128^3, $f = 4$ test calculation without ever rotating the rays. The mean radial positions of the dense shell and the ionization front are essentially the same in the two calculations, but the shape of the ionization front is significantly different. Figure 11 shows the neutral fraction in slices through the simulations without and with ray rotation after 1 Myr of evolution. With rotation of the rays, as shown in the top panel, the ionization front boundary is smooth and round. Without rotation, as shown in the bottom panel, there are projections of partially neutral gas several cells into the ionized region, and the front boundary is much less smooth. If one is interested in studying instabilities or similar phenomena at the ionization front boundary, the presence of these fingers of neutral gas is potentially problematic, since they may serve as seeds for instability.

These projections are caused by differences in the ionization rate for cells at the same radius due to discretization in angle. When the orientation of the rays is constant over long times, small differences in the ionization rate at different angles can build up to produce the finger-like structures shown in the figure. One could reduce the discretization error by using a more refined ray tree; for example, by requiring that the solid angle subtended by a cell al-
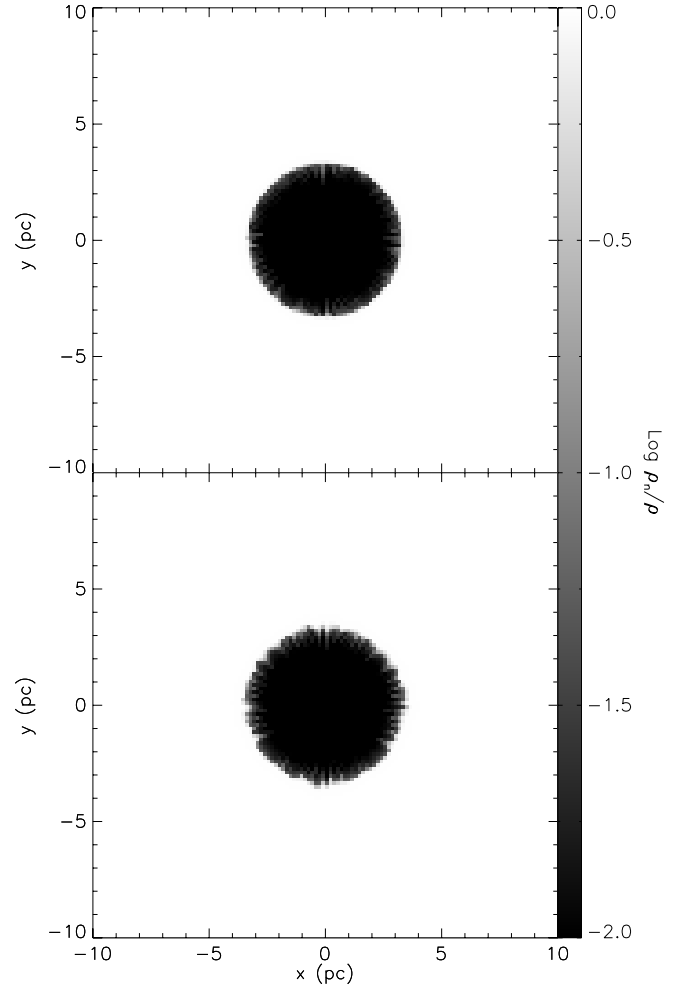

FIG. 11.—Slices through the *x-y* plane, showing the logarithm of the ratio of neutral density to total density in a simulation of a D-type ionization front 1 Myr ($17.7t_s$) after the start of the simulation, both with (*top*) and without (*bottom*) rotation of the rays.

ways be at least 4 times as large as the solid angle corresponding to a single ray, rather than twice as large, as we require. However, this is potentially expensive, since every factor-of-f increase in the angular resolution multiplies the time required to compute the photoionization rate by a factor of f^2, and the photoionization rate is computed during every radiation iteration. It is also overkill, since the angular discretization errors only build up over many time steps. Rotating the rays at intervals of five hydrodynamic time steps, as we have done for the simulation shown in the top panel of Figure 11, eliminates or vastly reduces the occurrence of neutral gas fingers in the ionized region. The computational cost of rebuilding the ray tree every five hydrodynamic time steps is completely negligible and, for parallel calculations, requires no additional interprocessor communication.

4.3.4. *Ionization Front Cooling*

We next explore the question of cooling in cells with mixed-ionization states. As discussed in § 3.3, we set the molecular cooling rate $\Lambda_{KI} = \Gamma_{KI} = 0$ in cells where the ionization fraction is between 1% and 99% in order to prevent excessive cooling due to numerical mixing of ionized and neutral gas. To demonstrate why this is necessary, we repeat our simulation of a D-type ionization front with $f = 4$, with cooling allowed in all cells regardless of ionization fraction.
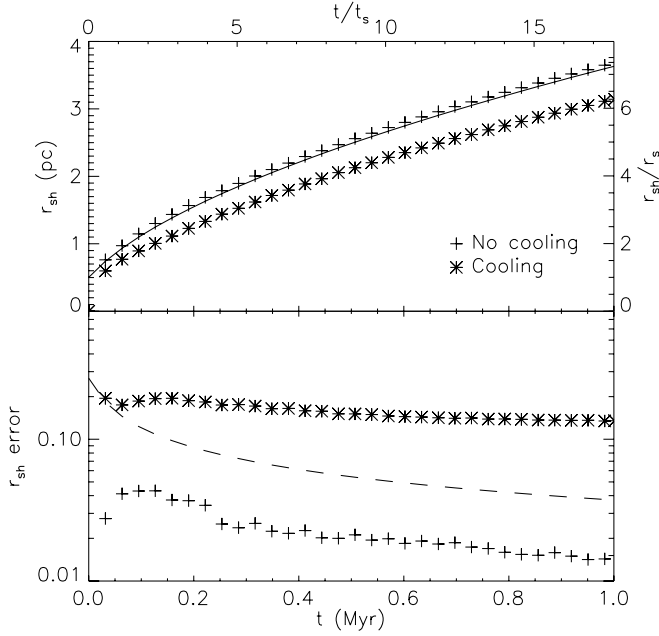
FIG. 12.—Radius of the dense shell around an H II region vs. time (*top*), computed both analytically (*solid line*) and by simulations at a resolution of 128^3 with cooling in mixed-ionization cells disabled (*plus signs*) and enabled (*asterisks*). We also show the fractional error in the position of the shell relative to the similarity solution (*bottom*), defined as $|r_{sim} - r_{analyt}|/r_{analyt}$, where r_{sim} is the simulated radius and r_{analyt} is the analytically computed radius. For comparison, we also show the fractional error corresponding to a difference of $(\sqrt{3}/2)\Delta x$ (*dashed line*), which is the distance from the center of a computational cell to its corner.

Figure 12 shows the result. The run in which all cells can cool lags behind the analytic solution even at early times, and the error does not decrease with time. The error at 1 Myr of evolution is roughly 14%, which is far larger than a single cell; the error in the slope of the r_{sh} versus t curve at 1 Myr is 10%. In contrast, the error with mixed-ionization cooling disabled is 1.5%, which corresponds roughly to the center-to-corner distance of a single cell. The slope of the curve differs from the analytic value by less than 1%.

The lag is a result of numerical mixing at the ionization front coupled with strong molecular cooling. To demonstrate this, we plot in Figure 13 both the total cooling rate $\mathcal{L} = \Lambda_{KI}n_H^2 + \Lambda_{rec}n_e n_{H^+} + \Lambda_{ion-ff}n_e n_{H^+}$ (Fig. 13a) and the molecular cooling rate $\Lambda_{KI}n_H^2$ (Fig. 13b) as a function of position from our run with cooling in mixed cells allowed, at a time of 1 Myr. As the plot shows, with cooling allowed in mixed cells, the energy loss rate from cells at the ionization front is much larger than that from cells either inside or outside the front.

One might think this is to be expected, since the cells at the ionization front are the ones furthest out of equilibrium and are the places where much of the energy from ionizing photons is deposited. However, 91% of the cooling in the region with an ionization fraction larger than 1% comes from molecular cooling rather than from recombination or forbidden line cooling, which are the processes that should dominate in ionized regions. Molecular cooling dominates even if we consider only the much more strongly ionized regions. Only if we limit our attention to gas for which the ionization fraction is greater than 74% do we find that molecular processes provide less than 50% of the total cooling rate. That molecular cooling is the dominant cooling process even in gas that is up to 74% ionized is a clear sign that the cooling rate is being artificially enhanced by numerical smearing of the ionization front, since for a real ionization front there
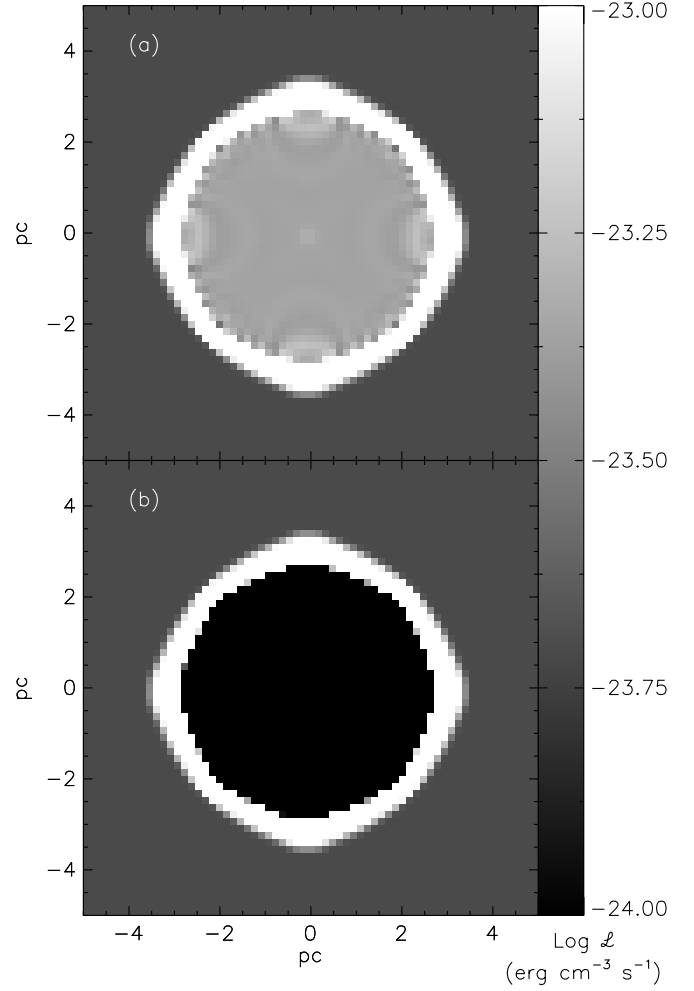

FIG. 13.—Slices through the *x-y* plane, showing (a) the total cooling rate and (b) the cooling rate due to molecular processes alone. Note that, in order to make the cooling rate in the fully ionized region visible, we have selected a color scale that saturates in the rapidly cooling ionization front. The maximum cooling rate reaches $>10^{-21}$ erg cm^{-3} s^{-1}.

should be no molecules to cool in regions that are significantly ionized.

This suggests that simulations can only compute the correct expansion rate for H II regions if they suppress the effects of numerical mixing at the ionization front. This may explain why some numerical simulations without this precaution produce results that lag the true solution significantly (e.g., Fig. 1 of Mac Low et al. 2007) unless they have enough resolution to at least marginally resolve the ionization front (e.g., Arthur & Hoare 2006), something that is not generally feasible in three dimensions. Note, however, that the strength of this effect depends on the existence of extremely rapid cooling processes for the non-ionized gas. If the chemistry is such that the neutral gas is not able to cool rapidly, the error will be much smaller. For example, with primordial cosmological chemistry, the lag is much smaller because the neutral gas cooling is much less efficient (D. Whalen 2006, private communication).

An important point is that the evolution in the case with over-cooling is *qualitatively* correct. There is an expanding, spherical H II region that drives a dense shell of swept-up material, and it follows a radius versus time curve that resembles the correct solution. It is only when one *quantitatively* checks the numerical method against an analytic solution that the problem becomes

clear. This highlights the importance of performing quantitative comparisons of numerical and analytic results.

5. H II REGION EVOLUTION IN UNIFORM MAGNETIZED MEDIA

Now that we have determined the constraints that our numerical method must obey from tests against analytic solutions in the hydrodynamic case, in this section we demonstrate our method in the simplest magnetohydrodynamic case. Several authors have considered the structure of magnetized ionization fronts in the past. However, these treatments have been limited to mathematically rigorous but one-dimensional analyses (e.g., Redman et al. 1998; Williams et al. 2000; Williams & Dyson 2001) or to simple, quasi-empirical models (e.g., Carlqvist et al. 2003; Ryutov et al. 2005) in more dimensions. Bertoldi (1989) gives an analytic model for the implosion of a spherical magnetized gas clump exposed to an external ionizing source, but his solution is limited to that case. To date, there has been no detailed simulation of the evolution of an H II region driven by a point source expanding into a magnetized medium in three dimensions.

We consider a neutral medium with an initially uniform density ρ_n and an isothermal sound speed c_n, threaded by a magnetic field $\boldsymbol{B}_n$ that is oriented in the x-direction. We place a source of constant ionizing luminosity s at the origin and set it to begin radiating at time $t = 0$. We restrict ourselves to the case most relevant to ionizing sources in magnetized molecular clouds, in which the thermal pressure in ionized gas at the initial density greatly exceeds the magnetic pressure, which in turn greatly exceeds the thermal pressure in neutral gas at the initial density. Thus, $c_i \gg v_A \gg c_n$, where v_A is the Alfvén speed in the neutral gas and c_i is the ionized gas sound speed.

5.1. Analytic Evaluation

We can sketch out an approximate evolutionary scenario for this configuration using simple analytic estimates. At times $t \lesssim t_s$, the gas has not yet had time to move in response to photoionization heating. Since photons do not feel the magnetic field, the evolution should be identical to that of a nonmagnetized R-type ionization front. At $t \sim t_s$, the ionized gas begins to expand due to its high thermal pressure. Since by assumption $c_i \gg v_A$, the magnetic field in the gas is initially unable to resist this expansion. Thus, as in the nonmagnetic case, a D-type ionization front forms, gas moves radially outward, and a spherical dense shell forms. Ionization balance requires that when the shell bounding the H II region reaches a radius r_{sh}, the gas that remains in its interior and is not in the shell wall must have started within a distance

$$r_0 = \sqrt{r_s r_{sh}} = r_s \left(1 + \frac{7t}{4t_s}\right)^{2/7} \qquad (40)$$

of the ionizing source. Since the magnetic field is frozen into the gas, the magnetic flux passing through the interior of the ionized region is therefore

$$\Phi_i = \pi B_n r_0^2 = \pi B_n r_s^2 \left(1 + \frac{7t}{4t_s}\right)^{4/7}, \qquad (41)$$

where B_n is the magnitude of the magnetic field in the neutral gas outside the shell. In contrast, the total flux passing through the ionized region and the dense shell combined is simply the initial flux that passed through a circle of radius r_{sh},

$$\Phi_{sh} = \pi B_n r_{sh}^2 = \pi B_n r_s^2 \left(1 + \frac{7t}{4t_s}\right)^{8/7}. \qquad (42)$$

Since the magnetic flux through the H II region interior rises as $t^{4/7}$ at late times, while the total flux through the H II region interior and the dense shell varies as $t^{8/7}$, the fraction of magnetic field lines passing through the H II region that also pass through the shell interior must decline with time as $t^{-4/7}$. Physically, this occurs because magnetic field lines are being dragged with the gas out of the H II region interior and concentrated in the dense shell.

As time progresses, the magnetic field changes the evolution in two ways. First, magnetic pressure and tension oppose the expansion of the shell perpendicular to the magnetic field, causing it to deform and become aspherical. This effect should become significant when magnetic pressure is comparable to thermal pressure in the ionized gas; i.e., when $\rho_n v_A^2 \sim \rho_i c_i^2$. Since $\rho_i \approx (r_{sh}/r_s)^{-3/2} \rho_n$ before magnetic effects are significant, we should see significant deformation of the H II region when it reaches a radius of order $r_{sh} \approx (c_i/v_A)^{4/3} r_s$ and the shell velocity is $v_{sh} \approx v_A$. Second, magnetic pressure in the swept-up shell will limit compression in the shock perpendicular to the magnetic field to a factor of order v_{sh}/v_A. This will make the shell thinner along the field and thicker perpendicular to it, an effect that will become noticeable when $v_{sh} \sim v_A$. Thus, we expect both magnetic effects to become significant when the region has expanded to a radius of roughly

$$r_m \equiv \left(\frac{c_i}{v_A}\right)^{4/3} r_s, \qquad (43)$$

which we define to be the magnetic critical radius r_m for the H II region. If we assume that $r_m \gg r_s$, this happens at roughly a time

$$t_m \equiv \frac{4}{7} \left(\frac{c_i}{v_A}\right)^{7/3} t_s \qquad (44)$$

after the start of the evolution.

The evolution will then enter a new stage in which the thermal pressure inside the ionized region is small compared to the magnetic pressure, but both still greatly exceed the thermal pressure in the neutral gas. Since $|\boldsymbol{B}|^2 \gg \rho_i c_i^2$, the gas will be able to move only along field lines, and those field lines will be approximately straight. Since expansion is subsonic with respect to the ionized gas and we are concerned with times $\gg t_s$, the density in the ionized gas along each field line is constant. Along the field, gas motions are unrestricted, and the evolution will be a normal D-type ionization front as seen in nonmagnetized gas. Perpendicular to the field, magnetic pressure effects become more and more dominant as time goes on. The ionized gas continues to escape along field lines, so the pressure driving expansion perpendicular to the field continues to fall, which causes the front to slow and the density contrast to decrease perpendicular to the field. At the same time, everywhere that the magnetic field is at all oblique to the front, a slow mode will move into the neutral medium. This mode will remove magnetic support from the front by bending the field lines so that they become perpendicular to the shock (Williams et al. 2000). This loss of magnetic support will cause the shell to transition from thick and magnetic pressure-dominated to thin and without much magnetic support over a larger and larger solid angle as time goes on, until only the circle exactly perpendicular to the initial magnetic field is not occupied by a dense shell.

5.2. Simulation

We simulate an initially neutral uniform medium with a density of $\rho_n = 2.34 \times 10^{-22}$ g cm^{-3} ($n_H = 100$ cm^{-3}) and a temperature of $T_n = 11$ K, threaded by a magnetic field $\boldsymbol{B} = 14.2\hat{\boldsymbol{x}}$ μG

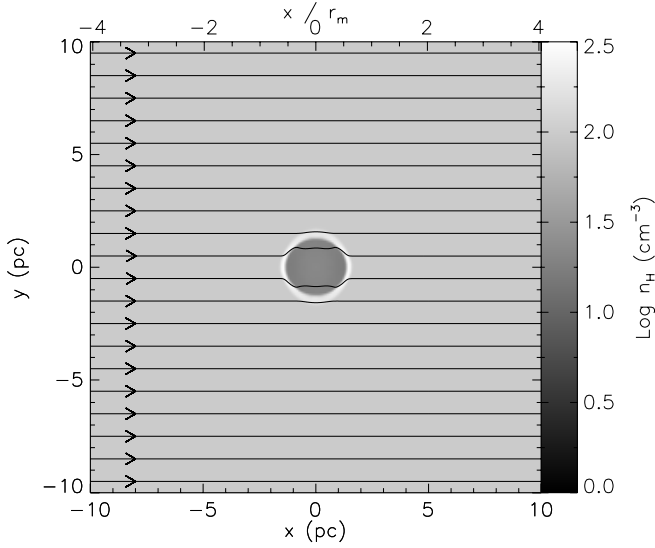FIG. 14.—Slice through the x-y plane, showing the logarithm of the hydrogen number density in a simulation of an ionization front expanding into a magnetized medium after a time of roughly $t_m/3 = 0.177$ Myr. The lines represent magnetic field lines. The density of the lines outside of the H II region corresponds to a magnetic field strength of 14.2 μG. [*This figure is a frame from an mpeg animation in the electronic edition of the Journal.*]

(4.0×10^{-6} in the units we use in this paper). This gives an Alfvén speed of 2.6 km s^{-1} and a sound speed of 0.20 km s^{-1} in the gas at the initial time. The computational domain runs from -10 pc to 10 pc in every direction, and there is an ionizing source at the origin with a luminosity of $s = 4.0 \times 10^{46}$ s^{-1}, giving an initial Strömgren radius of $r_s = 0.5$ pc and a sound crossing time of $t_s = 0.056$ Myr, if we assume a sound speed of $c_i = 8.7$ km s^{-1} in the ionized region. For these parameters, the magnetic critical radius is $r_m = 6.0 r_s = 2.5$ pc and the magnetic critical time is $t_m = 9.4 t_s = 0.53$ Myr. The simulation uses a resolution of 256^3 cells and a heating/cooling time step constraint of $f = 4$. We run the simulation for a time of $3 t_m = 1.58$ Myr.

5.2.1. Structure at $t \ll t_m$

Figures 14 and 15 show slices through the computational domain at a time of $t \approx t_m/3 = 0.177$ Myr. As expected, the H II region is almost spherical, although there is already a slight asymmetry visible. The outer boundary of the dense shell is nearly perfectly round and is at essentially the same radius as we would expect for a nonmagnetized region at the same time. However, the shell is slightly thicker in the direction perpendicular to the field than in the direction along the field, due to magnetic pressure opposing compression. Also as expected, because magnetic field lines are being advected with the gas, the magnetic field inside the H II region is much weaker than the background field, and the magnetic field in the dense shell is much stronger than the background field.

5.2.2. Structure at $t \sim t_m$

Figures 16 and 17 show the same slices as in Figures 14 and 15, but at a time $t \approx t_m = 0.53$ Myr. As expected, the magnetic effects have now become pronounced. The inner, ionized region is strongly prolate, such that it is roughly a factor of 2 longer along the field than perpendicular to it. Along the field, the dense shell is still only a few cells thick, the same as in the nonmagnetic case, but in the directions perpendicular to the field, its thickness is not much smaller than its radius, again as we expect. However, the
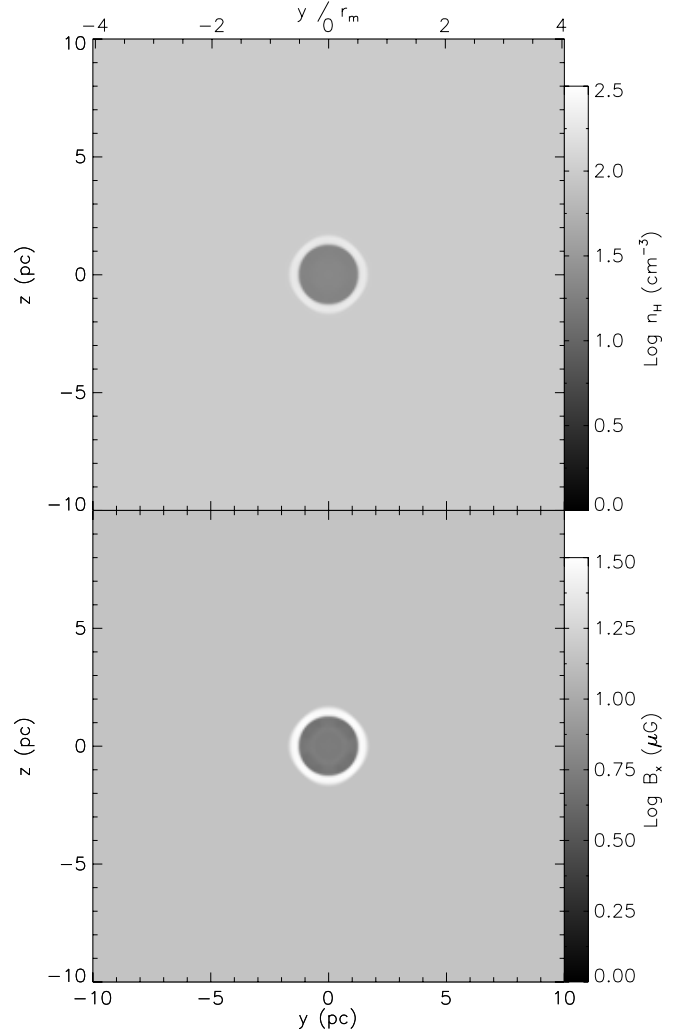

FIG. 15.—Slices through the y-z plane, showing the logarithm of the hydrogen number density (*top*) and the logarithm of the magnetic field strength (*bottom*), in a simulation of an ionization front expanding into a magnetized medium after a time of roughly $t_m/3 = 0.177$ Myr.
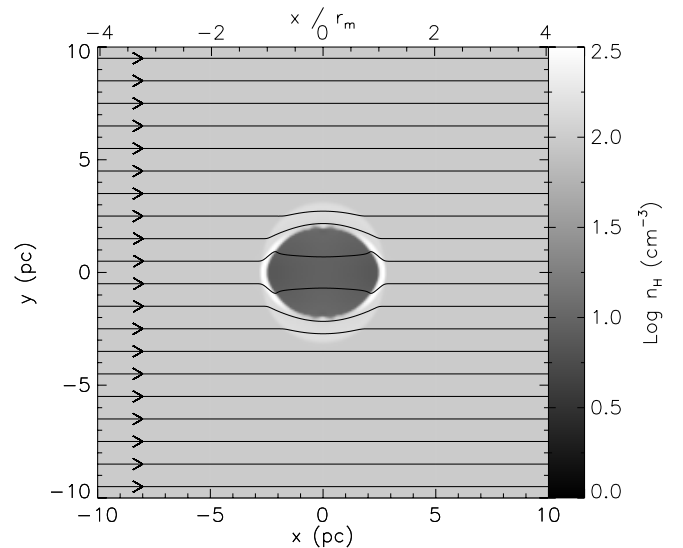

FIG. 16.—Same as Fig. 14, but after a time of roughly $t_m = 0.53$ Myr. [*This figure is a frame from an mpeg animation in the electronic edition of the Journal.*]

Fig. 18.—Same as Fig. 14, but after 1.58 Myr ($3t_m$). [*This figure is a frame from an mpeg animation in the electronic edition of the Journal.*]
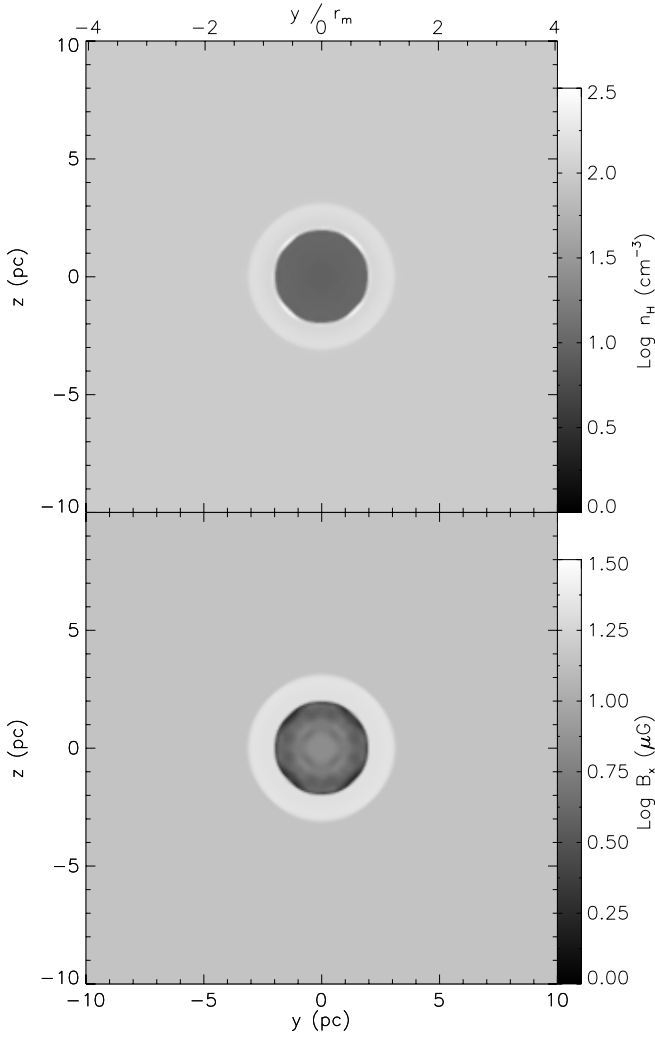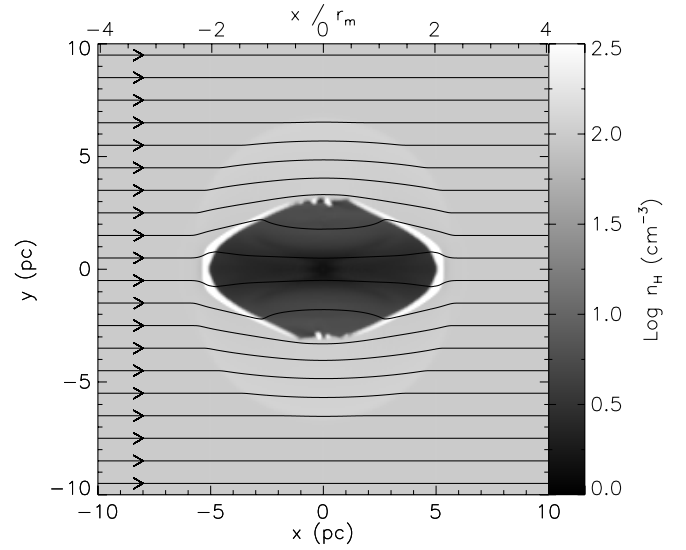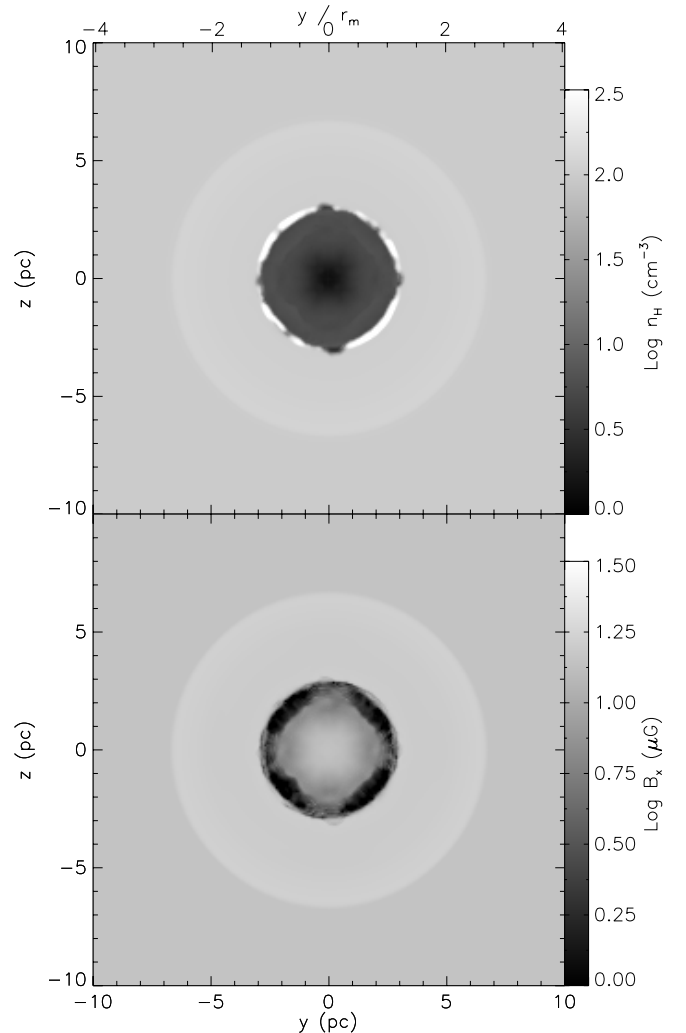
Fig. 17.—Same as Fig. 15, but after a time of roughly $t_m = 0.53$ Myr.

outer edge of the expanding shell is still not too far from spherical, and in the direction along the field, its radius is close to r_m, the radius we would expect were there no field present.

Interestingly, the shell extends somewhat further in the direction across the field than in the direction along it, so the outer radius of the swept-up shell in the y- and z-directions is noticeably larger than r_m. This is likely because at $t = t_m$ the expansion speed is comparable to the Alfvén speed and is slower than the fast magnetosonic speed. Across the field, signals from the shell can propagate by fast magnetosonic waves more rapidly than the shell can expand along the field. This creates a density disturbance, carried by the fast mode, ahead of the radius that the shell would have reached had there been no magnetic field. Since fast waves cannot propagate parallel to the magnetic field, the leading edge of the disturbance has advanced less in the x-direction than in the y- or z-directions.

5.2.3. Structure at $t \gg t_m$

Figures 18 and 19 show the simulation at our final time slice, $t = 1.58$ Myr $= 3t_m$. In the evolution after time t_m, the inner ionized region becomes more prolate, expanding more rapidly along the field than in the directions perpendicular to it. Along the field, in the x-direction, there continues to be a dense shell of swept-up material.

Fig. 19.—Same as Fig. 15, but after 1.58 Myr ($3t_m$).

In the y- and z-directions, which are perpendicular to the field lines, the thick shell of gas bounded between the ionization front and the fast-mode front has also continued to expand. However, its overdensity is decreasing, and at this time it is only 120 cm^{-3}, so it is within 20% of the background density and is smaller than the density at time t_m. The shell expansion velocity perpendicular to the field lines is much smaller than the Alfvén speed in the neutral gas and is comparable to the sound speed, so the leading edge of the shell is no longer a shock. Instead, the shell is just a small overdensity that is gradually relaxing away.

The boundary between the neutral and ionized regions in the y- and z-directions appears to be somewhat corrugated, but this may well be a numerical artifact, particularly since Figure 19 shows that the structure at the ionization front is aligned with the computational grid. It may be that for ionization fronts that are nearly static, as this one is due to magnetic confinement, rebuilding the ray tree once every five time steps is insufficient to prevent the growth of structures.

In between the x-direction and the y-direction, there is a region in which the field is oblique relative to the ionization front. As predicted by Williams et al. (2000), this region shows two types of behavior, depending on how close the field is to being parallel or perpendicular to the front. Closer to the x-direction, the field been bent such that it is close to perpendicular to the front, and a dense shell has formed at the ionization front. Closer to the y-direction, the field is not bent and enters the front at an oblique angle, and there is no shock or dense shell. The solid angle over which a dense shell exists appears to be increasing with time.

While expansion of the ionized region perpendicular to the field has slowed, the expansion velocity along the field lines is actually larger than or about the same as it was at time t_m. To understand the origin of this effect, consider the jump conditions that govern the expansion of the dense shell in the x-direction at late times, when the expansion is strongly subsonic with respect to the ionized gas, but the thermal pressure in the ionized gas is still much greater than the thermal pressure in the neutral gas. In the rest frame of the shell, this means that the ram pressure of neutral gas flowing into the shell must balance the thermal pressure of ionized gas downstream from it. Thus, $\rho_n v_{\rm sh}^2 \propto \rho_i c_i^2$, such that the shell expansion velocity obeys $v_{\rm sh} \propto c_i (\rho_i / \rho_n)^{1/2}$. Thus far, our analysis could apply just as well to a nonmagnetized H II region. The difference that the magnetic field makes is that it greatly reduces the expansion rate of the ionized region in two directions. This prevents the density from dropping as fast as it does in the hydrodynamic case, which keeps the expansion velocity larger at later times.

However, while the decline in ionized gas density and pressure is slower than in the hydrodynamic case, since gas continues to flow out of the ionized region along field lines into the end caps, the pressure inside the H II region is always decreasing after time t_m. By $3t_m$, the H II region is in the limit where the magnetic pressure inside the ionization front is higher than the thermal pressure. As a result, the field lines have begun to straighten out inside the ionized region, so they are kinked only at the dense end caps. The magnetic energy density inside the ionized region is actually larger than it was at time t_m, since for $t \gtrsim t_m$ the field lines that were pushed into the dense shell start to move back into the ionized region interior.

6. DISCUSSION AND CONCLUSION

We have demonstrated an algorithm for computing the evolution of magnetized molecular gases that are subjected to internal sources of ionizing radiation, which is potentially applicable to molecular clouds. In testing our algorithm, we have discovered three conditions that are likely to apply to ionizing radiation hydrodynamic and magnetohydrodynamic codes in general. First, to achieve maximum accuracy, the update time step must be limited so that the temperature in cells does not change by more than a factor of $f < 10$ between hydrodynamic or magnetohydrodynamic updates. Larger values of f produce small but significant errors in the expansion rates of D-type ionization fronts. Second, when using a ray-tracing approach to compute the ionizing radiative transfer, one should rotate the orientation of the rays periodically to avoid a build-up of errors caused by the discretization of angles around the ionizing source. Failure to obey this condition results in fronts that should be spherical developing aspherical features, and in more complex calculations this could potentially seed instabilities. Third, and most significantly, one must avoid overcooling caused by numerical smearing of the ionization front. This can be handled most easily by suppressing cooling in cells with mixed-ionization fractions. Failure to obey this condition leads to an unphysical loss of energy from expanding H II regions that causes them to lag behind analytic solutions by tens of percent. A calculation that satisfies these three constraints can reproduce the analytic solution for the expansion of a D-type ionization front to an accuracy of 1% for at least ~ 20 ionized sound-crossing times.

Using our algorithm, we report the first three-dimensional simulations of the expansion of an H II region into a magnetized gas. We show that the presence of a magnetic field distorts the H II region and greatly reduces the strength of the shock and the density contrast in directions perpendicular to the magnetic field. This leads to the formation of an H II region that is bounded by a dense shell of swept-up gas in the direction along the field, but not perpendicular to it. The absence of a dense shell over much of the solid angle means that, in the presence of strong, ordered magnetic fields, H II regions may not be able to collect and compress as much gas as one might expect from purely hydrodynamic estimates. This may reduce the efficiency of triggered star formation from H II regions.

REFERENCES

Abel, T., Norman, M. L., & Madau, P. 1999, ApJ, 523, 66
Abel, T., & Wandelt, B. D. 2002, MNRAS, 330, L53
Arthur, S. J., & Hoare, M. G. 2006, ApJS, 165, 283
Bertoldi, F. 1989, ApJ, 346, 735
Bertoldi, F., & Draine, B. T. 1996, ApJ, 458, 222
Blitz, L., Fukui, Y., Kawamura, A., Leroy, A., Mizuno, N., & Rosolowsky, E. 2007, in Protostars and Planets V, ed. B. Reipurth, D. Jewitt, & K. Keil (Tucson: Univ. Arizona Press), 81

Carlqvist, P., Gahm, G. F., & Kristen, H. 2003, A&A, 403, 399
Crutcher, R. M. 1999, ApJ, 520, 706
———. 2005, in IAU Symp. 227, Massive Star Birth: A Crossroads of Astrophysics, ed. R. Cesaroni, M. Felli, E. Churchwell, & M. Walmsley (Cambridge: Cambridge Univ. Press), 98
Dale, J. E., Bonnell, I. A., Clarke, C. J., & Bate, M. R. 2005, MNRAS, 358, 291
Garcia-Segura, G. 1997, ApJ, 489, L189
Garcia-Segura, G., & Franco, J. 1996, ApJ, 469, 171
García-Segura, G., & López, J. A. 2000, ApJ, 544, 336
Gardiner, T. A., & Stone, J. M. 2005, J. Comput. Phys., 205, 509
———. 2006, J. Comput. Phys., submitted
Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, ApJ, 622, 759
Heiles, C., & Crutcher, M. 2005, in Cosmic Magnetic Fields, ed. R. Wielebinski & R. Beck (Berlin: Springer), 137
Henney, W. J. 2007, in Diffuse Matter from Star Forming Regions to Active Galaxies, ed. T. W. Hartquist, J. M. Pittard, & S. A. E. G. Falle (Dordrecht: Springer), 103
Hosokawa, T., & Inutsuka, S.-I. 2005, ApJ, 623, 917
Koyama, H., & Inutsuka, S. 2002, ApJ, 564, L97
Krumholz, M. R., Matzner, C. D., & McKee, C. F. 2006, ApJ, 653, 361
Krumholz, M. R., & McKee, C. F. 2005, ApJ, 630, 250
Mac Low, M.-M., Toraskar, J., Oishi, J. S., & Abel, T. 2007, ApJ, 668, 980
Matzner, C. D. 2002, ApJ, 566, 302

McKee, C. F., & Williams, J. P. 1997, ApJ, 476, 144
Mellema, G., Arthur, S. J., Henney, W. J., Iliev, I. T., & Shapiro, P. R. 2006a, ApJ, 647, 397
Mellema, G., Iliev, I. T., Alvarez, M. A., & Shapiro, P. R. 2006b, NewA, 11, 374
Miniati, F., & Colella, P. 2007, J. Comput. Phys., 224, 519
Osterbrock, D. E. 1989, Astrophysics of Gaseous Nebulae and Active Galactic Nuclei (Mill Valley: University Science Books)
Redman, M. P., Williams, R. J. R., Dyson, J. E., Hartquist, T. W., & Fernandez, B. R. 1998, A&A, 331, 1099
Rijkhorst, E.-J., Plewa, T., Dubey, A., & Mellema, G. 2006, A&A, 452, 907
Ritzerveld, J. 2005, A&A, 439, L23
Ryutov, D. D., Kane, J. O., Mizuta, A., Pound, M. W., & Remington, B. A. 2005, Ap&SS, 298, 183
Sofia, U. J., & Meyer, D. M. 2001, ApJ, 554, L221
Spitzer, L. 1978, Physical Processes in the Interstellar Medium (New York: Wiley)
Tenorio-Tagle, G., Bodenheimer, P., Lin, D. N. C., & Noriega-Crespo, A. 1986, MNRAS, 221, 635
Whalen, D., & Norman, M. L. 2006, ApJS, 162, 281
Williams, J. P., & McKee, C. F. 1997, ApJ, 476, 166
Williams, R. J. R., & Dyson, J. E. 2001, MNRAS, 325, 293
Williams, R. J. R., Dyson, J. E., & Hartquist, T. W. 2000, MNRAS, 314, 315
Zuckerman, B., & Evans, N. J. 1974, ApJ, 192, L149