

OPTIMIZED MULTI-FREQUENCY SPECTRA FOR APPLICATIONS IN RADIATIVE FEEDBACK AND COSMOLOGICAL REIONIZATION

JORDAN MIROCHA^{1,2}, STEPHEN SKORY^{1,2}, JACK O. BURNS^{1,2}, AND JOHN H. WISE³

¹ Center for Astrophysics and Space Astronomy, University of Colorado, Campus Box 389, Boulder, CO 80309, USA; jordan.mirocha@colorado.edu

² The NASA Lunar Science Institute, NASA Ames Research Center, Moffett Field, CA 94035, USA

³ Center for Relativistic Astrophysics, School of Physics, Georgia Institute of Technology, Atlanta, GA 30332, USA

Received 2012 April 9; accepted 2012 July 5; published 2012 August 20

ABSTRACT

The recent implementation of radiative transfer algorithms in numerous hydrodynamics codes has led to a dramatic improvement in studies of feedback in various astrophysical environments. However, because of methodological limitations and computational expense, the spectra of radiation sources are generally sampled at only a few evenly spaced discrete emission frequencies. Using one-dimensional radiative transfer calculations, we investigate the discrepancies in gas properties surrounding model stars and accreting black holes that arise solely due to spectral discretization. We find that even in the idealized case of a static and uniform density field, commonly used discretization schemes induce errors in the neutral fraction and temperature by factors of two to three on average, and by over an order of magnitude in certain column density regimes. The consequences are most severe for radiative feedback operating on large scales, dense clumps of gas, and media consisting of multiple chemical species. We have developed a method for optimally constructing discrete spectra, and show that for two test cases of interest, carefully chosen four-bin spectra can eliminate errors associated with frequency resolution to high precision. Applying these findings to a fully three-dimensional radiation-hydrodynamic simulation of the early universe, we find that the H II region around a primordial star is substantially altered in both size and morphology, corroborating the one-dimensional prediction that discrete spectral energy distributions can lead to sizable inaccuracies in the physical properties of a medium, and as a result, the subsequent evolution and observable signatures of objects embedded within it.

Key words: dark ages, reionization, first stars – methods: numerical – radiative transfer

Online-only material: color figures

1. INTRODUCTION

Energy injection by radiative processes fundamentally changes the evolution of astrophysical systems, whether it be in the context of star formation, galaxy evolution, or the growth of supermassive black holes (SMBHs). For instance, ultraviolet photons from the universe’s first stars (Population III (PopIII) stars; Abel et al. 2002) photodissociate the primary coolant (H₂) that first enabled their formation. Very recent radiation-hydrodynamic calculations of PopIII stars find that PopIII star masses may be limited by protostellar radiative feedback, perhaps explaining the lack of evidence for exotic pair instability supernovae in the early universe (Hosokawa et al. 2011). Conventional metal line cooling driven star formation can be affected by radiative feedback as well. Krumholz (2006) showed that photoheating around newly formed stars can strongly suppress fragmentation in surrounding protostellar clouds, while Dale et al. (2005) see both positive and negative feedback operating in radiation-hydrodynamic simulations of star cluster formation. Radiative feedback could also be a barrier to efficient black hole (BH) growth in the early universe (Alvarez et al. 2009), as X-rays from accreting BHs efficiently photoheat surrounding gas, leading to smaller Bondi–Hoyle accretion rates (Bondi & Hoyle 1944).

The mere presence of ionizing/dissociating photons ensures a change in the chemical and thermal state of a gas, though the magnitude of these changes hinges squarely on the number of photons propagating through the gas and their spectral energy distribution (SED). Holding the bolometric luminosity of a radiation source constant, even subtle changes in the SED

can lead to noticeable differences in the properties of the surrounding medium. For example, adjusting the X-ray power-law index of a BH accretion spectrum results in ionization fronts which differ by factors of ≈ 2 –3 in radius, and temperature profiles varying by 10^2 – 10^3 K on scales of several hundred kpc (Thomas & Zaroubi 2008). Simply truncating the emission of identical X-ray SEDs at harder energies (0.4 keV rather than 0.2 keV) causes a drastic reduction in heating, ionized fractions, and H₂ fractions surrounding “miniquasars” at high redshift (Kuhlen & Madau 2005).

Unfortunately, not all radiative transfer algorithms are able to represent radiation sources with continuous SEDs, or perhaps cannot afford the additional computational expense associated with the frequency dependence of the radiative transfer equation. The natural first step is to represent sources as monochromatic emitters, choosing an emission frequency characteristic of the full SED. Some authors have improved upon the monochromatic treatment using “multi-group” methods, which average SED properties and absorption cross-sections over one or more frequency bandpasses (Gnedin & Abel 2001; Aubert & Teyssier 2008), while others have sampled continuous SEDs at n_ν frequencies, which are generally evenly spaced bins (in linear or log-space) between the hydrogen ionization threshold and an upper frequency cutoff. In either case, there is no clear method of deciding how many frequency-averaged bandpasses or discrete emission frequencies are required for a given problem, and though the standard multi-group treatment is physically motivated, it does not guarantee that the photoionization and photoheating rates are adequately reproduced as a function of column density.

Frequency resolution has recently been studied in radiation-hydrodynamic settings by Wise & Abel (2011) and Whalen & Norman (2008). Wise & Abel (2011) find that for the expansion of an H II region around a 10^5 K blackbody source in a hydrogen-only medium, the density, temperature, velocity, and ionization profiles are well converged for $n_v \geq 4$. Use of a monochromatic spectrum for this problem introduces significant errors since all photons are absorbed at a characteristic column density, whereas multi-frequency treatments achieve some column density dependent behavior and can thus mimic the behavior of a truly continuous spectrum. Whalen & Norman (2008) studied the effects of frequency resolution in the setting of I-front instabilities, and did not achieve convergence until $n_v \geq 80$ (logarithmically spaced between 13.6 and 90 eV).

The convergence for the test of Wise & Abel (2011) using only four frequency bins is reassuring, though the prospects for convergence are less clear if one were interested in the absorption processes of multiple chemical species, ionization and heating due to X-rays and their energetic secondary photoelectrons (Shull & van Steenberg 1985; Furlanetto & Stoeve 2010), or inhomogeneous media. Kramer & Haiman (2008, hereafter KH08) briefly compared monochromatic and continuous treatments of absorbed power-law X-ray sources in a study of ionization front thickness around high- z quasars (the I-front thickness is a potentially powerful indirect probe of the ionizing spectrum of high- z quasars). The hydrogen and helium I-front thickness is expected to grow over the lifetime of a quasar given the discrepancy in evolution timescales between the largest and smallest scales. At small radii, photoionization equilibrium is reached quickly since ionizing photons are abundant, whereas geometrical dilution and attenuation of the initial radiation field slow ionization evolution considerably on large scales, effectively “stretching out” the I-fronts of hydrogen and helium with time. A monochromatic representation of the quasar SED leads to a reduction in this effect, but also leads to severe errors in the overall ionization structure (see Figure 3 of KH08). These errors are of the same order of magnitude as those resulting from the neglect of physical effects, such as ionization via helium recombination photons (KH08, Figure 6), or ionization from secondary electrons (KH08, Figure 7). These effects are likely important in studies of radiative feedback from stars and active galactic nuclei (AGNs), and most certainly in efforts to simulate cosmological reionization. An effort must be made to ensure that the SEDs used in numerical simulations accurately reflect the properties of their continuous analogs, especially if it is spectrum-dependent effects in which we are most interested.

We will focus on the following questions in this paper. How significant are the errors in the temperature and ionization state of a medium that arise solely due to the discretization of SEDs? How many frequencies are required to minimize such errors, where must they be positioned in frequency space, and how should their relative luminosities be apportioned? For what numerical methods is it possible to represent sources with continuous SEDs, or are there perhaps advantages in discretizing SEDs, even when it is not required by the algorithm of choice? Answers to these questions may lead to revised interpretations of previous studies which used discrete radiation fields, but more importantly, will reduce the guesswork involved in discretizing SEDs, and promote frequency resolution to the same status as spatial, temporal, and mass resolution, which are more easily selected on a problem-by-problem basis.

In Section 2, we will introduce the one-dimensional radiative transfer framework used to obtain the solutions presented in later

sections. In Section 3, we quantitatively assess the accuracy with which multi-frequency calculations reproduce the ionization and heating profiles of continuous SEDs. Section 4 is devoted to introducing a technique for optimally selecting discrete SED templates, and Section 5 will present the results obtained with this method, including applications to one-dimensional and fully three-dimensional radiation-hydrodynamic calculations. Discussion and conclusions can be found in Sections 6 and 7, respectively. Validation of the radiative transfer code used for this work and further details regarding the optimization algorithm can be found in the Appendix.

2. RADIATIVE TRANSFER FRAMEWORK

One dimensional radiative transfer calculations around point sources have been used to model cosmological reionization (Fukugita & Kawasaki 1994), the thickness of quasar ionization fronts (KH08), the time evolution of ionization and heating around first stars, galaxies, and quasars (Thomas & Zaroubi 2008; Venkatesan & Benson 2011), and their associated observable signatures. Given that our focus is on frequency resolution, it would be unnecessary to perform calculations in a more complex setting than this, with additional unrelated physics. As a result, our one-dimensional methods strongly resemble those used by previous authors, though for completeness, we will reiterate the aspects of these methods most pertinent to the problem at hand.

In general, the chemical and thermal evolution of gas surrounding a radiation source is governed by a set of differential equations describing the number densities of all ions and the temperature of the gas. Assuming a medium consisting of hydrogen and helium only, we first solve for the abundances of each ion via

$$\frac{dn_{\text{H II}}}{dt} = (\Gamma_{\text{H I}} + \gamma_{\text{H I}} + \beta_{\text{H I}} n_e) n_{\text{H I}} - \alpha_{\text{H II}} n_e n_{\text{H II}} \quad (1)$$

$$\begin{aligned} \frac{dn_{\text{He II}}}{dt} = & (\Gamma_{\text{He I}} + \gamma_{\text{He I}} + \beta_{\text{He I}} n_e) n_{\text{He I}} + \alpha_{\text{He III}} n_e n_{\text{He III}} \\ & - (\beta_{\text{He II}} + \alpha_{\text{He II}} + \xi_{\text{He II}}) n_e n_{\text{He II}} \end{aligned} \quad (2)$$

$$\frac{dn_{\text{He III}}}{dt} = (\Gamma_{\text{He II}} + \gamma_{\text{He II}} + \beta_{\text{He II}} n_e) n_{\text{He II}} - \alpha_{\text{He III}} n_e n_{\text{He III}}. \quad (3)$$

Each of these equations represents the balance between ionizations of species H I, He I, and He II, and recombinations of H II, He II, and He III. Associating the index i with absorbing species, $i = \text{H I}, \text{He I}, \text{He II}$, and the index i' with ions, $i' = \text{H II}, \text{He II}, \text{He III}$, we define Γ_i as the photoionization rate coefficient, γ_i as the secondary ionization rate coefficient, $\alpha_{i'}$ ($\xi_{i'}$) as the case-B (dielectric) recombination rate coefficients, β_i as the collisional ionization rate coefficients, and $n_e = n_{\text{H II}} + n_{\text{He II}} + 2n_{\text{He III}}$ as the number density of electrons.

At each time step, we also solve for the temperature evolution, dT_k/dt , which is given by

$$\begin{aligned} \frac{3}{2} \frac{dT_k}{dt} \left(\frac{k_B T_k n_{\text{tot}}}{\mu} \right) = & f^{\text{heat}} \sum_i n_i \mathcal{H}_i - \sum_i \zeta_i n_e n_i - \sum_{i'} \eta_{i'} n_e n_{i'} \\ & - \sum_i \psi_i n_e n_i - \omega_{\text{He II}} n_e n_{\text{He II}}, \end{aligned} \quad (4)$$

where \mathcal{H}_i is the photoelectric heating rate coefficient (due to electrons previously bound to species i), $\omega_{\text{He II}}$ is the dielectric

recombination cooling coefficient, and ζ_i , η_i , and ψ_i are the collisional ionization, recombination, and collisional excitation cooling coefficients, respectively. The constants in Equation (4) are the total number density of baryons, $n_{\text{tot}} = n_{\text{H}} + n_{\text{He}} + n_{\text{e}}$, the mean molecular weight, μ , Boltzmann's constant, k_{B} , and the fraction of secondary electron energy deposited as heat, f^{heat} . We use the formulae in Appendix B of Fukugita & Kawasaki (1994) to compute the values of α_i , β_i , ξ_i , ζ_i , η_i , ψ_i , and $\omega_{\text{He II}}$.

The most critical aspect of propagating the radiation field in our one-dimensional simulations is computing the ionization (Γ_i , γ_i) and heating (\mathcal{H}_i) rate coefficients accurately. In order to directly relate our results to fully three-dimensional radiative transfer calculations, we have chosen to adopt a photon-conserving (PC) algorithm nearly identical to those employed by several widely used codes, like *C²Ray* (Mellema et al. 2006; Friedrich et al. 2012), and *Enzo* (Wise & Abel 2011). Our code is able to compute Γ_i , γ_i , and \mathcal{H}_i in a non-photon-conserving (NPC) fashion as well, to enable comparison with previous one-dimensional work such as Thomas & Zaroubi (2008). The two formalisms are equivalent in the limit of very optically thin cells, a condition that can be met easily in one-dimensional calculations but is rarely computationally feasible in three dimensions. For NPC methods, if the optical depth of an individual cell is substantial, the number of ionizations in that cell will *not* equal the number of photons absorbed for that cell, i.e., photon number will not be conserved. This problem was remedied by Abel et al. (1999), who inferred the number of photoionizations of species i in a cell from the radiation incident upon it and its optical depth

$$\Delta\tau_{i,v} = n_i \sigma_{i,v} \Delta r. \quad (5)$$

It is most straightforward to imagine our one-dimensional grid as a collection of concentric spherical shells, each having thickness Δr and volume $V_{\text{sh}}(r) = 4\pi[(r + \Delta r)^3 - r^3]/3$, where r is the distance between the origin and the inner interface of each shell. The ionization and heating rates can then be related to the number of absorptions in any given shell (thus preserving photon number), as

$$\Gamma_i = A_i \int_{\nu_i}^{\infty} I_{\nu} e^{-\tau_{\nu}} (1 - e^{-\Delta\tau_{i,v}}) \frac{d\nu}{h\nu} \quad (6)$$

$$\gamma_{ij} = A_j \int_{\nu_j}^{\infty} \left(\frac{\nu - \nu_j}{\nu_i} \right) I_{\nu} e^{-\tau_{\nu}} (1 - e^{-\Delta\tau_{j,v}}) \frac{d\nu}{h\nu} \quad (7)$$

$$\mathcal{H}_i = A_i \int_{\nu_i}^{\infty} (\nu - \nu_i) I_{\nu} e^{-\tau_{\nu}} (1 - e^{-\Delta\tau_{i,v}}) \frac{d\nu}{\nu}, \quad (8)$$

where we have defined the normalization constant $A_i \equiv L_{\text{bol}}/n_i V_{\text{sh}}(r)$, and denote the ionization threshold energy for species i as $h\nu_i$. I_{ν} represents the SED of radiation sources, and satisfies $\int_{\nu} I_{\nu} d\nu = 1$, such that $L_{\text{bol}} I_{\nu} = L_{\nu}$.

Equation (7) represents ionizations of species i due to fast secondary electrons from photoionizations of species j , which has number density n_j , and ionization threshold energy, $h\nu_j$. f_i^{ion} is the fraction of photoelectron energy deposited as ionizations of species i . In the remaining sections, we only include the effects of secondary electrons when considering X-ray sources, which emit photons in the range $10^2 \text{ eV} < E < 10^4 \text{ eV}$. In this regime, the values of f^{heat} and f_i^{ion} computed via the formulae of Shull & van Steenberg (1985) are sufficiently accurate, but for radiation at lower energies where f^{heat} and f_i^{ion} have a stronger

energy dependence, the fitting formulae of Ricotti et al. (2002) or the lookup tables of Furlanetto & Stoeve (2010) would be more appropriate. The total secondary ionization rate for a given species, γ_i , is the sum of ionizations due to the secondary electrons from all species, $\gamma_i = f_i^{\text{ion}} \sum_j \gamma_{ij} n_j / n_i$.

The optical depth, $\tau_{\nu} = \tau_{\nu}(r)$, in the above equations is the total optical depth at frequency ν due to all absorbing species, i.e.,

$$\begin{aligned} \tau_{\nu}(r) &= \sum_i \int_0^r \sigma_{i,\nu} n_i(r') dr' \\ &= \sum_i \sigma_{i,\nu} N_i(r), \end{aligned} \quad (9)$$

where N_i is the column density of species i at distance r from the source. We calculate the bound-free absorption cross-sections using the fits of Verner et al. (1996) throughout.

The values of Γ_i , γ_i , and \mathcal{H}_i are completely predetermined for a given radiation source, and as a result, can be tabulated as a function of column density to avoid evaluating the integrals in these expressions numerically “on-the-fly” as a simulation runs (e.g., Mellema et al. 2006; Thomas & Zaroubi 2008). Isolating the frequency-dependent components of Equations (6)–(8), we can define the integrals

$$\Phi_i(\tau_{\nu}) \equiv \int_{\nu_i}^{\infty} I_{\nu} e^{-\tau_{\nu}} \frac{d\nu}{h\nu} \quad (10)$$

$$\Psi_i(\tau_{\nu}) \equiv \int_{\nu_i}^{\infty} I_{\nu} e^{-\tau_{\nu}} d\nu, \quad (11)$$

allowing us to re-express the rate coefficients as

$$\Gamma_i = A_i [\Phi_i(\tau_{\nu}) - \Phi_i(\tau'_{i,v})] \quad (12)$$

$$\gamma_{ij} = \frac{A_j}{h\nu_i} \{ \Psi_j(\tau_{\nu}) - \Psi_j(\tau'_{j,v}) - h\nu_j [\Phi_j(\tau_{\nu}) - \Phi_j(\tau'_{j,v})] \} \quad (13)$$

$$\mathcal{H}_i = A_i \{ \Psi_i(\tau_{\nu}) - \Psi_i(\tau'_{i,v}) - h\nu_i [\Phi_i(\tau_{\nu}) - \Phi_i(\tau'_{i,v})] \}, \quad (14)$$

where $\tau'_{i,v} \equiv \tau_{\nu} + \Delta\tau_{i,v}$. Later references to “continuous SEDs” signify use of this technique, where the integral values Φ_i and Ψ_i are computed over a column density interval of interest a priori using a Gaussian quadrature technique, rather than on-the-fly via discrete summation.

Tabulating Equations (10) and (11) grants a significant speed-up computationally, but also forms the basis of our frequency resolution optimization strategy (Section 4). Note, however, that in general the dimensionality of these lookup tables is equal to the number of absorbing species (through $\Delta\tau_{i,v}$), so the tables for simulations including hydrogen only are one dimensional, while those including hydrogen and helium are three dimensional. If we chose to adopt the secondary electron treatment of Ricotti et al. (2002) or Furlanetto & Stoeve (2010), our lookup tables would inherit an additional dimension, as the secondary ionization and heating factors f_i^{ion} and f^{heat} would depend both on photon energy and the hydrogen ionized fraction, $x_{\text{H II}}$.

Equations (12)–(14) are completely general for PC algorithms, whether the source SEDs are discrete or continuous—the

only difference being for discrete SEDs, the integrals in Equations (10) and (11) become sums over the number of discrete emission frequencies, n_v . In practice, computing Γ_i , γ_i , and \mathcal{H}_i is more straightforward for sources with discrete SEDs, as we can simply count the number of ionizations caused by photons at each individual frequency, and convert this into the amount of excess electron kinetic energy available for further heating and ionization. When testing the accuracy of discrete solutions in later sections we employ this method, where radiation is emitted at n_v frequencies, with each frequency ν_n carrying a fraction I_n of the source’s bolometric luminosity. The photoionization and heating coefficients can then be expressed as

$$\Gamma_{i,n} = \frac{A_i I_n}{h \nu_n} e^{-\tau_{\nu_n}} (1 - e^{-\Delta\tau_{i,\nu_n}}) \quad (15)$$

$$\gamma_{ij,n} = \Gamma_{j,\nu_n} (\nu_n - \nu_j) / \nu_i \quad (16)$$

$$\mathcal{H}_{i,n} = \Gamma_{i,\nu_n} h (\nu_n - \nu_i). \quad (17)$$

The total rate coefficients can be found by summing each of these expressions over all frequencies, $n = 1, 2, 3, \dots, n_v$. These equations are identical to Equations (12)–(14) for the discrete SED case, but are perhaps more intuitive.

For simplicity, our current treatment neglects a few physical processes that are cosmological in origin, or simply do not rely on the radiation field directly. These include cooling via free–free emission and hydrogen and helium ionizations due to helium recombination photons (which depend on the gas kinetic temperature and electron density), and cosmological effects such as Hubble cooling, Compton cooling off cosmic microwave background (CMB) photons, and photoionization by Wien-tail CMB photons (which depend on kinetic temperature, redshift, and the Hubble parameter).

Two additional approximations are implicit in the remainder of this paper. They are (1) the infinite speed-of-light approximation and (2) the on-the-spot approximation (we use the case-B recombination coefficients in Equations (1)–(3)). The former approximation could be dubious for very bright sources in low-density media, while the latter is generally not a good assumption, as discussed at length in Cantalupo & Porciani (2011). As a result, the *absolute* accuracy of our solutions is not guaranteed in regimes where careful treatment of the speed of light and recombination photons is necessary, but this is acceptable since we only care about the *relative* differences among our solutions. The optimized SEDs of Section 5 will apply equally well to simulations including more ionization and/or heating/cooling processes, so long as they do not depend directly on the radiation field (e.g., ionization of H I and He I by helium recombination photons; Friedrich et al. 2012).

3. ASSESSING THE CONSEQUENCES OF DISCRETE RADIATION FIELDS

To quantify the differences between the ionization and temperature profiles around sources with continuous and discrete SEDs, we will simulate two test problems. First, the standard case of a 10^5 K blackbody in a hydrogen-only medium, and second, a power-law X-ray source in a medium consisting of both hydrogen and helium.

3.1. 10^5 K Blackbody

The 10^5 K blackbody problem has been studied extensively (e.g., test problem 2 in the Radiative Transfer Comparison Project; Iliev et al. 2006, hereafter RT06) due to its simplicity, and perhaps also because the surface temperatures of PopIII stars are expected to be $\sim 10^5$ K (Schaerer 2002). We adopt nearly the identical setup as in RT06, i.e., a uniform hydrogen-only medium with number density $n_H = 10^{-3} \text{ cm}^{-3}$, initial ionized fraction $x_{\text{HII}} = 1.2 \times 10^{-3}$, initial temperature $T_0 = 10^2$ K, and a 10^5 K blackbody with an ionizing photon luminosity of $\dot{Q} = 5 \times 10^{48} \text{ s}^{-1}$. The only difference between our simulations and RT06 is that we use a domain $L_{\text{box}} = 10$ kpc in size, rather than $L_{\text{box}} = 6.6$ kpc, to allow for a comparison of discrete and continuous solutions at slightly larger radii. We evolve the simulations for 500 Myr on a grid of 200 linearly spaced cells between $0.1 < r/\text{kpc} < 10$, ignoring the details of secondary ionization (i.e., all photoelectron energy is deposited as heat).

In Figure 1, we compare the ionization and temperature profiles around two 10^5 K “blackbody” sources of constant ionizing photon luminosity $\dot{Q} = 5 \times 10^{48} \text{ s}^{-1}$ —one a true blackbody emitter with a continuous SED spanning the range 13.6–100 eV (black lines), and the other with a monochromatic SED at $h\nu_1 = 29.6$ eV, the average energy of ionizing photons for this source (red lines). We can see the same qualitative results that have been pointed out by previous authors, namely, that monochromatic sources of radiation fail to ionize (top panels) and heat (lower panels) gas at large radii as significantly as continuous sources, since all photons are absorbed near a single characteristic column density, representing the point where $\tau_{\nu_1} \approx 1$, i.e., $N_{\text{char}} \sim \sigma_{\nu_1}^{-1}$. The relative error in the position of the ionization front, Δr_{IF} , where $r_{\text{IF}} \equiv r(x_{\text{H I}} = x_{\text{H II}} = 0.5)$, is 8% after 10 Myr, 10% after 100 Myr, and 11% after 500 Myr. In the optically thin regime, the monochromatic spectrum overestimates ionization by factors of two to three on average and up to an order of magnitude at all times, though the latter effect is primarily because the neutral fraction is a steeply declining function with decreasing radius, and the I-fronts of the two solutions are offset. Outside the I-front, the situation is more interesting as the gas is mostly neutral. After 100 Myr of evolution, the ionized fraction outside the I-front is underestimated by a factor of two on average, and by as much as a factor of six.

The temperature evolution, shown in the bottom panels of Figure 1, is significantly more troubling. The monochromatic source captures the temperature well within the ionization front where the gas is in photoionization equilibrium, but quickly diverges from the continuous solution outside. Like the ionization profiles, discrepancies grow with time. After 10 Myr of evolution, the monochromatic source underestimates the temperature at large radii by a factor of two on average, and by a factor of seven at the point of greatest discrepancy. After 100 (500) Myr, the discrete solution underestimates the temperature by up to a factor of 17 (41).

If considering the heating and ionization around a single PopIII star, the errors induced by monochromatic treatments may not be cause for concern upon first inspection since PopIII stars are expected to live only a few Myr, and we can see that errors are less significant at early times. However, the intergalactic medium (IGM) is subject to the ionization and heating caused by all sources, whose cumulative impact will be substantial even though the ionization and heating caused by individual sources may be very small. Globally, then, the IGM




Figure 1. Comparison of ionization (top) and temperature (bottom) profiles around a 10^5 K blackbody source after 10 Myr (left) and 100 Myr (right) using continuous (black) and monochromatic (red) SEDs. Solid lines in the top panels correspond to the neutral fraction (x_{HI}), while dashed lines correspond to the ionized fraction (x_{HII}). We apply these line color and line style conventions for all radial profiles presented in this paper.

(A color version of this figure is available in the online journal.)

is insensitive to individual stellar lifetimes, and instead evolves as it would if ionizing photons originated from a single, very luminous, very long lived object.

This manner of thinking has already materialized in the realm of large volume cosmological simulations, where “star particles” are generally as luminous as one or more star clusters, and “galaxy particles” behave in a way that is consistent with the integrated properties of an entire galactic stellar population (and perhaps active nucleus). Such approximations are necessary with limited spatial resolution, but more than adequate for studies of the IGM. Over time though, errors in gas properties due to poor frequency resolution will accrue, as it is the combined properties of all radiation sources which affect IGM properties, however short-lived each individual source may be.

3.2. Power-law X-Ray Source

To address the effects of discrete SEDs in environments where multiple chemical species are important and large attenuating columns are possible, we now turn our attention to a power-law X-ray source embedded in a 1 Mpc domain consisting of hydrogen and helium, with a primordial helium abundance (by mass) of $Y = 0.2477$.

Our selection of parameters for this problem is motivated by studies of high-redshift quasars, and particularly their role in the epoch of reionization (e.g., Venkatesan et al. 2001). X-rays have long mean free paths, and as a result are capable of ionizing and heating gas on very large (\sim Mpc) scales. Large-scale heating is responsible for driving the high-redshift all-sky 21 cm signal toward emission, and inducing fluctuations in 21 cm power spectra on large angular scales (for a review of 21 cm cosmology, see Furlanetto et al. 2006). An early X-ray background may also be important in interpreting the

optical depth to electron scattering of the CMB (e.g., Ricotti et al. 2005; Shull & Venkatesan 2008).

While supernovae and/or X-ray binaries could be important sources of hard photons in the early universe, we assume the source of X-rays is persistent—an accreting SMBH with mass $M_\bullet = 10^6 M_\odot$ and radiative efficiency of $\epsilon_\bullet = 10\%$, which leads to a bolometric luminosity of $\mathcal{L}_{\text{bol}} = \epsilon_\bullet \mathcal{L}_{\text{edd}} \simeq 1.26 \times 10^{43} \text{ erg s}^{-1}$. Here, $\mathcal{L}_{\text{edd}} = 4\pi G M_\bullet m_p c / \sigma_T$ is the Eddington luminosity, where m_p is the proton mass and σ_T the Thomson cross-section. The mass (and thus luminosity) of the SMBH is allowed to grow as it accretes,

$$M_\bullet(t) = M_\bullet(0) \exp \left[\frac{1 - \epsilon_\bullet}{\epsilon_\bullet} \left(\frac{t}{t_{\text{edd}}} \right) \right], \quad (18)$$

where $t_{\text{edd}} = 0.45 \text{ Gyr}$ is the e -folding timescale for SMBH growth (an Eddington, or Salpeter time). The SED is taken to be a power law of the form

$$I_\nu \propto \left(\frac{h\nu}{\text{keV}} \right)^{1-\alpha}, \quad (19)$$

where α is the spectral index. We adopt $\alpha = 1.5$, over the energy range 10^2 – 10^4 eV. The surrounding medium has a constant mass density of $\rho = 5.4 \times 10^{-28} \text{ g cm}^{-3}$ (cosmic mean at redshift $z = 10$), initial ionized fractions $x_{\text{HII}} = x_{\text{HeII}} = 10^{-4}$, $x_{\text{HeIII}} = 0$, and initial temperature $T_0 = 10^2 \text{ K}$. The domain for this problem is divided into 400 cells linearly spaced between $0.01 < r/\text{Mpc} < 1$, and is evolved for $\epsilon_\bullet t_{\text{edd}} = 45 \text{ Myr}$.

In Figure 2, we compare the hydrogen and helium ionization profiles for two X-ray sources having the same bolometric luminosity. One, a continuous power-law source as described




Figure 2. Comparison of hydrogen (top) and helium (bottom) ionization profiles around an $\alpha = 1.5$ power-law X-ray source after 45 Myr using continuous (black) and monochromatic (red) SEDs.

(A color version of this figure is available in the online journal.)

above, and the other a monochromatic source of 0.5 keV photons (a fiducial monochromatic emission energy). The monochromatic source underestimates the radii of both the hydrogen and helium ionization fronts by a factor of ~ 2.3 , and overestimates the hydrogen neutral fraction on average by a factor of three, and at most by a factor of 20 within the hydrogen I-front. The same general picture applies to helium, where errors in the neutral helium fraction are enormous since the He I–He II I-front is very sharp (as it was for hydrogen in the previous section), and $x_{\text{He II}}$ and $x_{\text{He III}}$ are in error by factors of 2–20 depending on radius.

Errors in the temperature profile are less extreme, as shown in Figure 3. On small scales, the monochromatic source captures the temperature quite well, but at large radii, the monochromatic source overestimates temperatures by a factor of two on average.

The disparity in the magnitude of ionization and temperature errors is a reflection of the strong frequency dependence of the bound–free absorption coefficients. Photoionization of hydrogen or helium by 0.5 keV photons is rare, but when it does occur, at least $\sim 90\%$ of the original photon energy is left to be deposited mostly as heat, unless the free electron density is very low. Because the ionization of hydrogen and helium by the monochromatic source is very inaccurate, errors in the free electron density will substantially alter the amount of secondary electron energy deposited as heat, rather than further ionization.

The consequences of miscalculating ionization and heating could affect efforts to model and interpret current and future 21 cm measurements, since the primary 21 cm observable, the differential brightness temperature (δT_b), depends on the hydrogen neutral fraction, UV radiation field, electron density, and the gas kinetic temperature (T_K) (Furlanetto et al. 2006).




Figure 3. Comparison of temperature profiles around an $\alpha = 1.5$ power-law X-ray source after 45 Myr using continuous (black) and monochromatic (red) SEDs.

(A color version of this figure is available in the online journal.)

Neglecting the presence of a Ly α background, the scaling

$$\delta T_b \propto T_K^{0.4} (1 + \delta) (1 + z)^{-1/2} \times \begin{cases} x_{\text{HI}} n_e, & n_e \gg n_{\text{HI}} \\ x_{\text{HI}}^2, & n_e \ll n_{\text{HI}} \end{cases} \quad (20)$$

holds approximately in regimes where $T_{\text{CMB}} \ll T_K \lesssim 10^4$ K.

In the immediate vicinity of radiation sources where gas is entirely ionized, $\delta T_b \rightarrow 0$ due to the leading x_{HI} term, but at large radii where the ionizing flux is weaker, the δT_b signatures of stars and quasars could vary significantly solely due to miscalculations of x_{HI} , n_e , and T_K . The above scalings have especially strong consequences for gas within a few Mpc of strong X-ray sources, where hydrogen is weakly ionized, temperatures are of the order of 10^2 – 10^3 K, and the free electron density is enhanced due to efficient ionization of helium by the hard radiation field. In the earliest stages of reionization where $T_K < T_{\text{CMB}}(z)$ and the Ly α background is important, errors in x_{HI} , n_e , and T_K will lead to errors in δT_b as well, though in a less straightforward way, since the spin temperature, T_S , must be computed carefully.

4. OPTIMIZATION STRATEGY

To avoid errors of the sort described in the previous section, we have developed a technique for optimally constructing discrete SEDs that preserves the ionization and heating properties of their continuous counterparts. Although ray-tracing algorithms are capable of tabulating the relevant ionization and heating quantities (Equations (10) and (11)), few codes have taken advantage of this, and have instead cast monochromatic rays (e.g., state of the art reionization simulations with $n_v = 5$; Trac et al. 2008). Monte Carlo codes (e.g., CRASH; Maselli et al. 2003) have been used to simulate reionization with $n_v \geq 20$ multi-frequency photon packets (Ciardi et al. 2012), though such a large number of frequencies may be computationally debilitating for some algorithms, or unnecessary depending on the problem of interest.

Even when the algorithm of choice is compatible with propagating continuous radiation fields via tabulation of Equations (10) and (11), it may not be computationally advantageous. The overhead alone can in fact be substantial, particularly in the case of source-dependent SEDs—for example, the SED of a stellar population as a function of age, or BH accretion spectra

Table 1
Optimal SEDs for 10^5 K Blackbody Sources

n_v	$n = 1$	$n = 2$	$n = 3$	$n = 4$
1	(29.61, 0.89)
2	(27.93, 0.68)	(62.04, 0.21)
3	(20.58, 0.39)	(40.75, 0.39)	(69.23, 0.11)	...
4	(17.98, 0.23)	(31.15, 0.36)	(49.09, 0.24)	(76.98, 0.06)

Notes. Each entry is the $(h\nu_n, I_n)$ pair for bin n . Energies are in units of eV, and normalizations are expressed as fraction of the bolometric luminosity.

that vary with mass or luminosity. Such situations would require a separate lookup table for Equations (10) and (11) at each age/mass/luminosity of interest for a given radiation source. In addition, there are algorithms for which propagating continuous radiation fields in large volumes become completely intractable, yet large volumes are a necessity for the science questions of interest (e.g., reionization). For more discussion on these issues, see Section 6.

As introduced in Section 2, our optimization strategy relies on the fact that the SED of a radiation source appears only in the quantities Φ_i and Ψ_i (see Equations (10) and (11)). If we can construct a discrete SED that reproduces the values of Φ_i and Ψ_i to a high degree of accuracy over a column density interval of interest, then the discrete radiation field is indistinguishable from its continuous counterpart, and we have successfully preserved the true radiative properties of the source.

For sources with discrete SEDs, Equations (10) and (11) become

$$\Phi'_i(\tau_{\nu_n}) \equiv \sum_{n=1}^{n_v} \frac{I_n}{h\nu_n} e^{-\tau_{\nu_n}} \quad (21)$$

$$\Psi'_i(\tau_{\nu_n}) \equiv \sum_{n=1}^{n_v} I_n e^{-\tau_{\nu_n}}, \quad (22)$$

where we have used primes to indicate that these quantities are computed by direct summation over $n = 1, 2, \dots, n_v$ frequencies, rather than by a continuous integral.

Ensuring that $\Phi_i = \Phi'_i$ and $\Psi_i = \Psi'_i$ is a minimization problem of dimensionality $2n_v$, since each additional frequency bin lends two degrees of freedom—its frequency (ν_n), and the fraction of the bolometric luminosity assigned to that frequency (I_n). Our goal is to minimize the difference between continuous and discrete solutions, i.e.,

$$\begin{aligned} \Phi_i - \Phi'_i &= 0 \\ \Psi_i - \Psi'_i &= 0. \end{aligned} \quad (23)$$

These functions span several orders of magnitude over a broad range in column density, making it more practical to seek solutions to

$$\begin{aligned} \log\left(\frac{\Phi_i}{\Phi'_i}\right) &= 0 \\ \log\left(\frac{\Psi_i}{\Psi'_i}\right) &= 0 \end{aligned} \quad (24)$$

which place equal emphasis on all column densities. Preserving the high column density behavior of Φ_i and Ψ_i is especially important for very luminous sources and/or environments with

Table 2
Optimal SEDs for $\alpha = 1.5$ Power-law X-Ray Sources

n_v	$n = 1$	$n = 2$	$n = 3$	$n = 4$
1	(999.98, 1.00)
2	(255.87, 0.17)	(2553.6, 0.83)
3	(171.93, 0.08)	(518.22, 0.14)	(3098.5, 0.78)	...
4	(146.11, 0.05)	(307.30, 0.07)	(704.56, 0.14)	(3564.2, 0.73)

Note. Same as Table 1 but for an $\alpha = 1.5$ power-law X-ray source.

dense clumps in the immediate vicinity of the source, since the actual photoionization and heating rates are a combination of Φ_i , Ψ_i , and the normalization factor $A_i \propto L_{\text{bol}}/r^2$.

For a given n_v and source SED, we solve Equation (24) using the optimization technique Simulated Annealing (Kirkpatrick et al. 1983; Černý 1985), which traverses our $2n_v$ dimensional parameter space in search of the frequency–normalization pairs (ν_n, I_n) that best reproduce the values of Φ_i and Ψ_i . We leave a more detailed description of the algorithm and our implementation of it to the Appendix.

5. RESULTS

5.1. Optimal Discrete SEDs

We have obtained optimal SEDs for a 10^5 K blackbody emitting in the range 13.6–100 eV, and an $\alpha = 1.5$ power-law X-ray source with emission spanning the interval 10^2 – 10^4 eV. In each case, we set the upper column density limit for our optimization to be the column density of a fully neutral medium, i.e., $N_{\text{H I}}^{\text{max}} = n_{\text{H}} L_{\text{box}}$ and $N_{\text{He I}}^{\text{max}} = n_{\text{He}} L_{\text{box}}$, where we use L_{box} to denote the size of the domain, as in RT06. For the 10^5 K blackbody simulations, this works out to be $N_{\text{H I}}^{\text{max}} = 3.1 \times 10^{19} \text{ cm}^{-2}$, and for the power-law X-ray simulations, $N_{\text{H I}}^{\text{max}} \simeq \times 10^{22} \text{ cm}^{-2}$ and $N_{\text{He I}}^{\text{max}} \simeq \times 10^{21} \text{ cm}^{-2}$. For cosmological simulations with periodic boundary conditions, the upper column density limits would need to be chosen based on a maximum length scale of interest, or for radiative feedback focused simulations, by the column density of the densest objects of interest (damped Ly α systems, for example). Such choices are already made in ray-tracing calculations to limit computational expense. Generally, rays are terminated once the emission has been attenuated by a large factor.

The only situation in which we do not evaluate the full cost function is $n_v = 1$, where we instead optimize for the optically thin regime alone (i.e., only the first term of Equation (A2)), where Φ_i and Ψ_i are \sim constant with column density. In this case, the optimal solutions are simply those that preserve the bolometric luminosity of the source and the total number of ionizing photons, and can be verified analytically (Equations (10) and (11)). For the case of a hydrogen and helium medium, we have found that neglecting He II opacities mitigates the computational cost of the computation while resulting in no appreciable changes in our optimal SEDs and thus negligible changes in Φ' and Ψ' . The main results are summarized in Figures 6 and 7 and Tables 1 and 2, all results derived from $K = 2 \times 10^4$ and $K = 10^4$ Monte Carlo trials, for the 10^5 K blackbody and $\alpha = 1.5$ power-law source, respectively.

From Tables 1 and 2, it is clear that the optimal emission frequencies for both sources are not evenly spaced above the hydrogen or helium ionization thresholds, either in linear or log-space. In each case, the addition of a new frequency bin leads to a decrease in both the emission frequency and normalization




Figure 4. Top panels: comparison of Φ_{HI} and Φ'_{HI} (a) and Ψ_{HI} and Ψ'_{HI} (b) as a function of H I column density for a 10^5 K blackbody, showing the numerically computed continuous integral (solid black), best-fit composite four-bin discrete sum (blue crosses), and the contribution from each individual discrete frequency bin (dashed blue). Annotations represent the $(h\nu_n, I_n)$ pairs for each frequency group, drawn from Table 1. Bottom panels: percent error between discrete and continuous solutions. The solid blue line is the error for the four-bin optimal solution, while the errors induced by three-, two-, and one-bin solutions are shown in magenta, green, and red, respectively.

(A color version of this figure is available in the online journal.)

of all other bins. This signifies (1) the efficacy with which high energy photons photoionize and photoheat gas at large column densities (a regime inaccessible to lower energy photons which become optically thick at small columns), and (2) the increase in excess electron kinetic energy available for further ionization and heating with increasing photon energy. The former effect is most important for the blackbody source, which we can see in Figure 4. Not surprisingly, it is the lowest energy photons ($h\nu_1 = 17.98$ eV) in the $n_v = 4$ spectrum that are responsible for the ionization (through Φ) in the optically thin regime, while successively higher frequency bins become the primary agents of ionization as we move to higher column densities. The same trend does not hold completely in Figure 4(b), as in this case it is the second and third energy bins that provide the bulk of the heating (through Ψ) at low column densities.




Figure 5. Same as Figure 4 but for an $\alpha = 1.5$ power-law X-ray source. (A color version of this figure is available in the online journal.)

For the X-ray source, the second effect dominates, as the optical depth at any column density is small for most photons considered ($10^2 < h\nu < 10^4$ eV) over the entire domain. As shown in Figure 5, the photons responsible for the majority of the heating (through Ψ) over all column densities are those in the highest energy bin, the same photons which are the least effective at ionization. The trends and errors of Figure 5 are the same for Φ_i and Ψ as a function of helium column density.

In Figures 6 and 7, we show the probability distribution functions (PDFs) for the position and normalization of the optimal SED frequency bins obtained (drawn from Tables 1 and 2). Solutions are less tightly constrained as n_v is increased, as evidenced by a broadening in the distributions of frequency and normalization for each bin. This behavior is expected, given that each new bin contributes to the magnitude of Φ and Ψ in some region of column density space previously occupied by one or more other frequencies.

Holding I_n constant, a decrease in v_n will cause a negative vertical shift in the contribution of bin n to the magnitude of Φ , for example, but will simultaneously add power at larger column densities, since the turnover point for bin n occurs at $N_{\text{char}} \sim \sigma_{v_n}^{-1}$, and $\sigma_{v_n} \sim v^{-3}$. To avoid an increase in f , the

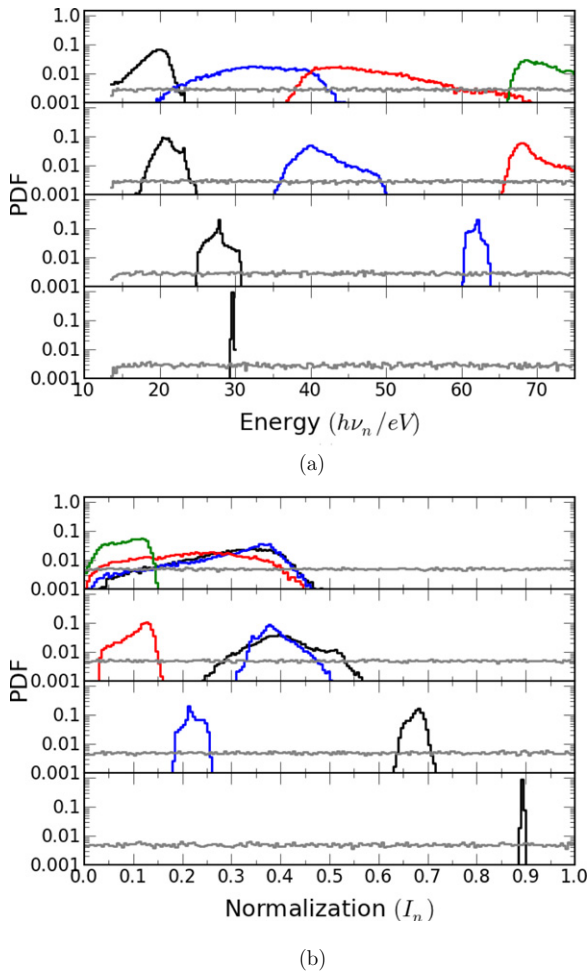


Figure 6. Emission energy (a) and normalization (b) probability distribution functions (PDFs) of optimized discrete 10^5 K blackbody spectrum using $n_v = 1, 2, 3, 4$ (from bottom to top). In each panel, the gray histogram denotes the initial guesses for all Monte Carlo trials, and the black, blue, red, and green histograms show the end point for the first, second, third, and fourth bins, respectively (ordered by increasing emission frequency).

(A color version of this figure is available in the online journal.)

power lost at small column densities has to be compensated for, either by a decrease in ν_{n-1} , or an increase in I_{n-1} , where $n-1$ denotes the bin with frequency $\nu_{n-1} < \nu_n$. As a result, there are degeneracies between all bins, and the magnitude of the degeneracy is greatest for bins positioned closest in frequency space. In order to tighten the PDFs for each optimal frequency bin, one or more terms would need to be added to f , in order to assign preference to one set of bins over another. For our purposes, any SED that minimizes f is just as good as any other, but additional terms in the cost function are certainly justifiable in the case of a ray-tracing calculation, where higher emission frequencies increase the computational cost of a calculation since their mean free paths are long. Adding a term to f that scales with ν_n would encourage optimal SEDs with the smallest emission frequencies possible, for example.

Optimization for $n_v > 4$ is certainly possible, though unnecessary in our case. At a given frequency, the transition from optically thin $\tau = 0$ to optically thick ($\tau \gtrsim 1$) in the functions Φ and Ψ occurs over an order of magnitude in column density (by definition, see Equation (9)). For both SEDs we have investigated, the column density regime of interest spans fewer than four orders of magnitude, motivating our




Figure 7. Same as Figure 6 but for an $\alpha = 1.5$ power-law X-ray source.

(A color version of this figure is available in the online journal.)

choice of $1 \leq n_v \leq 4$. We have performed optimizations with $n_v > 4$, but the addition of each additional bin when $n_v > \log_{10}(N_{\max}/N_{\min})$ reduces the error between Φ and Φ' , and Ψ and Ψ' much less significantly than additional bins when $n_v \leq \log_{10}(N_{\max}/N_{\min})$. For a given n_v , increasing N_{\max} will simply increase $\max|\Phi - \Phi'|$ and $\max|\Psi - \Psi'|$.

5.2. Confirmation with One-dimensional Calculations

To verify the solutions of the previous section, we ran simulations identical to those of Section 3 but with our optimal discrete SEDs. We compute Γ_i , γ_i , and \mathcal{H}_i via Equations (15)–(17) “on-the-fly,” rather than generating lookup tables of Φ_i and Ψ_i . As expected, accurate preservation of the quantities Φ_i and Ψ_i over the column density ranges of interest renders ionization and temperature profiles around sources of discrete radiation indistinguishable from their continuous counterparts.

In Figure 8, we compare ionization and heating around a 10^5 K blackbody after 100 Myr of evolution as in Section 3, showing the solution obtained with our optimal monochromatic (red) and four-bin (blue) SEDs. The continuous and four-bin solutions are indistinguishable.

In Figure 9, we perform the same analysis for the $\alpha = 1.5$ power-law simulations. Our optimal four-bin SED reproduces the hydrogen and helium ionization profiles (and thus electron density) and temperature of a continuous SED to high precision. The most noticeable errors are in the hydrogen neutral fraction within the hydrogen ionization front, where errors between




Figure 8. Comparison of ionization (top) and temperature (bottom) profiles around a 10^5 K blackbody source after 100 Myr showing the solutions obtained using continuous (black), monochromatic (red), and optimal four-bin discrete (blue circles/squares) SEDs.

(A color version of this figure is available in the online journal.)

four-bin and continuous solutions are still only $\sim 1\%$. Errors in $x_{\text{He III}}$ are negligible, justifying our neglect of $N_{\text{He II}}$ in the optimization process.

It should be noted that our optimal monochromatic SED for the X-ray source performs even more poorly than the fiducial 0.5 keV SED. This signifies a general problem with monochromatic emission for any spectrum with a hard component. Whereas the monochromatic optimization ($\tau_v = 0$) works quite well in the 10^5 K blackbody case since hydrogen absorbs UV photons readily, X-rays are not so readily absorbed by hydrogen and/or helium. As a result, the characteristic column density where most 1 keV photons are absorbed lies outside of our domain, leading to severe underionization (of all species) and underheating. The reason the 0.5 keV SED works better is because its characteristic absorption column is smaller, lying within our domain. We have experimented with relaxing the optically thin requirement for monochromatic optimization, and find that it is equally difficult to preserve ionization and heating profiles with emission at a single frequency.

5.3. Three-dimensional Radiation-hydrodynamic Simulations with Enzo

To study the impact of spectral discretization in a more complex setting, we ran RT06 test problem 2 with hydrodynamics, as well as two fully three-dimensional cosmological radiation-hydrodynamic simulations similar to those of Abel et al. (2007) and Alvarez et al. (2009), both with the *Enzo* code (Bryan &




Figure 9. Comparison of hydrogen and helium ionization (a), and temperature (b) profiles around a power-law X-ray source after 50 Myr showing the solutions obtained using continuous (black) and optimal four-bin discrete (blue symbols) SEDs.

(A color version of this figure is available in the online journal.)

Norman 1997; O’Shea et al. 2004).⁴ All analysis was performed with *yt* (Turk et al. 2011).

The results of the RT06 radiation-hydrodynamic test problem are shown in Figure 10, where we compare the solutions obtained using the four-bin SED employed by Wise & Abel (2011) in addition to our own (Table 1). The solutions are indistinguishable, which is expected given the relatively small range of column density explored in this problem.

The cosmological simulations follow the formation of a $100 M_{\odot}$ PopIII star, its brief 2.7 Myr lifetime in which it emits

⁴ Revision f4a8b5f5e6c5, modified to form only one star and use optimal SEDs.




Figure 10. Comparison of the four-bin solutions of Wise & Abel (2011; black) and our own (blue crosses) in a radiation-hydrodynamic simulation using the *Enzo* code. The setup is the same as in RT06 test problem 2, except hydrodynamics is included. (A color version of this figure is available in the online journal.)

1.2×10^{50} ionizing photons per second, and the X-ray emission resulting from accretion onto a remnant BH assumed to form via direct collapse after stellar death (as in Alvarez et al. 2009). The accretion rate, and thus luminosity assuming $\epsilon_{\bullet} = 10\%$, is the Bondi–Hoyle accretion rate of the cell in which the BH resides. The simulation volume is $0.25 \text{ Mpc } h^{-1}$ on a side, with 128^3 particles and cells on the root grid. A single nested grid occupies the inner one-eighth of the volume at twice the root grid resolution, where eight additional levels of adaptive-mesh refinement are allowed, yielding a peak spatial resolution of $0.23 \text{ pc } h^{-1}$.

We run two simulations, each identical to the other except for the choice of discrete SED. Our “control” simulation uses monochromatic SEDs—the PopIII star is a monochromatic source of $E = 29.6 \text{ eV}$ photons, while the X-ray source emits at $E = 2 \text{ keV}$. The second simulation employs the optimal four-bin SEDs found in Tables 1 and 2.

As shown in Figure 11, the magnitude of the errors between monochromatic and $n_v = 4$ solutions is even more significant in the cosmological problem than in the RT06 test problem, since the ionizing luminosity of the blackbody source considered is nearly two orders of magnitude larger (1.2×10^{50} versus $5 \times 10^{48} \text{ s}^{-1}$). For very luminous sources, even small errors in Φ and Ψ will become noticeable as characteristic timescales for photoionization and heating are short.

During the BH phase of evolution, there are more ways for the monochromatic and multi-frequency solutions to differ aside from the SEDs being employed. The accretion luminosity depends on local gas properties, which will be different in each simulation due to errors accrued during the PopIII star’s lifetime. Properties of the broader medium will of course vary for the same reason, leading to changes in how far soft X-rays are able to propagate before being absorbed. Throughout the 100 Myr of evolution, we simulate after the PopIII star’s death, the Bondi–Hoyle accretion rate and thus luminosity of the accreting BH is on average an order of magnitude smaller in the $n_v = 4$ simulation than for the monochromatic case. Errors in ionization and temperature exceeding an order of magnitude persist throughout the BH phase as well. Rather




Figure 11. Ratio of slices of the ionized fraction (a) and temperature (b) obtained using our optimized $n_v = 4$ blackbody SED ($x_{\text{HII},4}$, T_4) and the standard monochromatic SED ($x_{\text{HII},1}$, T_1). Both slices are 2.25 Myr after the formation of a Population III star. Contours (from center outward) correspond to hydrogen column densities of $N_{\text{H I}} = 2$ and $4 \times 10^{19} \text{ cm}^{-2}$.

(A color version of this figure is available in the online journal.)

than attempt to disentangle the BH phase induced errors from the preexisting errors, we simply emphasize that SED-induced errors will compound in feedback situations like this, since the initial conditions of each subsequent generation of objects will have been contaminated by errors associated with the previous one.

We cannot comment on the relative errors between monochromatic and multi-frequency treatments beyond the outermost column density contour, as our optimization extended only to $N_{\text{H I}} = 3.1 \times 10^{19} \text{ cm}^{-2}$. Future work focused on larger cosmological volumes, more luminous sources, and harder radiation fields will need to construct optimal SEDs valid beyond $N_{\text{H I}} = 10^{20} \text{ cm}^{-2}$, at least.

6. DISCUSSION

Algorithms developed for the purpose of studying point-source radiation (e.g., ray-tracing) are, in principle, capable of propagating continuous radiation fields, that is, tabulating Equations (10) and (11) and computing ionization and heating rates via Equations (12)–(14). The reason many have not taken this approach could be due to the additional computational overhead involved with using continuous SEDs—the quantities Φ_i and Ψ_i must be tabulated over the complete column density interval of interest. This includes column densities of all absorbing species, each of which must extend from the smallest expected column (i.e., the column density of a “fully ionized” cell—we adopted a minimum species fraction of $x_{\text{min}} = 10^{-5}$) up to the largest expected column (i.e., the column density of a fully neutral medium). The dimensionality of Φ_i and Ψ_i can be increased even further if, for example, energy-dependent secondary electron treatments (e.g., Ricotti et al. 2002; Furlanetto & Stoever 2010) or time-dependent SEDs are of interest.

For the simulations of Section 3.2, we generated three-dimensional lookup tables for Φ_i and Ψ_i covering the column density range $10^{11} < N_{\text{H I}} < 10^{21}$, and $10^{10} < N_{\text{He I}}, N_{\text{He II}} < 10^{20}$, sampling $N_{\text{H I}}$ at 200 points, and $N_{\text{He I}}$ and $N_{\text{He II}}$ with 100 points each, resulting in six three-dimensional tables, each consisting of 2×10^6 elements. We found that poorer sampling (e.g., tables of dimension $100 \times 50 \times 50$) leads to artificial “notches” in ionization and temperature profiles due to errors in the trilinear interpolation. In our case, $\Phi_{\text{H I}} = \Phi_{\text{He I}} = \Phi_{\text{He II}}$ and $\Psi_{\text{H I}} = \Psi_{\text{He I}} = \Psi_{\text{He II}}$ since all emission occurs above 10^2 eV , making the lower limit of integration for each quantity identical. In the general case, where emission extends all the way to the hydrogen ionization threshold, all six quantities would be unique. Generating these tables can take hundreds of CPU hours or more for a single SED depending on the number of column density elements. In addition, the radiative transfer solver requires additional modules to read in the lookup table, and perform interpolation four times per absorbing species per grid element (see Equations (12)–(14)). For sources with discrete SEDs, one can simply compute the photoionization rate for each neutral species, from which point the secondary ionization and heating rate coefficients are obtained in a simple algebraic fashion (see Equations (15)–(17)).

For high-resolution simulations focused on a single source of radiation (e.g., Kuhlen & Madau 2005; Alvarez et al. 2009), the additional effort required to accommodate continuous radiation fields seems well worth it to ensure that the ionization and thermal state of the gas is captured accurately. However, in large-scale simulations of cosmic reionization, which may spawn hundreds of thousands or perhaps millions of radiat-

ing “star particles” (depending on the simulation volume, resolution, etc.), ray-tracing methods are certainly not the most computationally advantageous algorithm. This is because the computational cost of a ray-tracing calculation scales with the number of radiation sources *and* the number of frequency bins in each source SED (though the former cost can be mitigated by merging nearby radiation sources; Trac & Cen 2007; Okamoto et al. 2012). If photons with long mean free paths are of interest, the simulation will be even more expensive since rays must be followed to larger distances, i.e., more ray segments and iterations of the numerical solver are required. An appealing option is to instead use moment-based methods such as the Variable Eddington Tensor approach (e.g., Gnedin & Abel 2001; Petkova & Springel 2009), flux-limited diffusion (e.g., Reynolds et al. 2009), or other variations (González et al. 2007; Aubert & Teyssier 2008; Finlator et al. 2009), as the computational cost of such algorithms is independent of the number of radiation sources and the mean free paths of photons, scaling only with the number of frequency bins in each source spectrum.

As discussed in Section 1, multi-group schemes common in the literature are an improvement over fiducial discrete SEDs, though it is not generally clear how many bandpasses are required for a given problem, or where they should lie in frequency space. Moreover, multi-group radiation suffers from the same problem as discrete polychromatic emission: photons at each frequency are absorbed near a characteristic column density, N_{char} . Computing new spectrum-weighted absorption cross-sections, $\bar{\sigma}_n$, for each frequency group merely shifts the location of N_{char} .

In principle, our minimization technique could be used to optimally select which bandpasses should be used for a multi-group algorithm, though in practice it would be much more computationally expensive. Rather than varying the location (v_n) or normalization (I_n) of frequency bin n on each Monte Carlo step, one would instead vary the position of bandpass edges, which would change the mean photon energy in each bandpass ($h\bar{v}_n$) and spectrum-weighted cross-section, $\bar{\sigma}_n$ (e.g., Aubert & Teyssier 2008). Because $h\bar{v}_n$ and $\bar{\sigma}_n$ are integral quantities, they would need to be computed numerically on each Monte Carlo step, and thus hundreds of thousands of times for a single optimization.

7. CONCLUSIONS

We have shown that the manner in which a discrete SED is constructed can induce substantial errors in simulation results, both in the ionization and temperature profiles around stars and quasars. But, these errors can be avoided to a large degree using only four discrete emission frequencies if source SEDs are designed via the methods of Section 4. Discrete SEDs constructed in a simple way (e.g., bins linearly spaced in frequency) will perform more poorly than optimally selected SEDs with the same number of bins, since it is the column density interval of interest that dictates the range of photon energies required, and the power to which each is assigned.

In general, discrete SED treatments fail to ionize and/or heat gas at large column densities, i.e., large physical scales or environments with dense clumps of gas. This has strong implications for simulations dedicated to understanding the magnitude and mode of radiative feedback on gas surrounding radiation sources. Current questions of this sort include whether or not radiation stimulates or suppresses further star formation in nearby protostellar clouds, and if radiative feedback can stifle the growth of SMBHs at high redshift.

As expected, extending our one-dimensional work to three dimensions produces ionized regions around a first star and remnant BH that deviate significantly in ionized fraction, temperature, size, and morphology. Such findings have implications in radiative feedback, but also in studies of both hydrogen and helium reionization. Certainly, miscalculations of the ionization state of gas surrounding galaxies in the early universe will lead to errors in the volume averaged neutral fraction, volume filling factor of ionized gas, and the optical depth of the CMB to electron scattering (τ_e). As we demonstrated in Section 3, such errors also introduce uncertainties in the interpretation of future 21 cm measurements, since the primary observable quantity (δT_b) depends directly on the hydrogen neutral fraction, electron density, and gas kinetic temperature.

Our optimizations in this work are by no means comprehensive, having selected two commonly used radiation sources (UV blackbody and X-ray power law) as test cases to demonstrate the method. However, optimization for more complex spectra is straightforward, and any new optimizations run will be made publicly available by the authors. The minimization code and one-dimensional radiative transfer codes are both available upon request. We leave more detailed investigations of reionization and radiative feedback, including multiple radiation sources and multi-frequency radiation transport, to future work.

The authors thank Steven Furlanetto and Daniel Reynolds for feedback on earlier versions of this draft, as well as the anonymous referee for a thorough review and many helpful suggestions. The LUNAR consortium (<http://lunar.colorado.edu>), headquartered at the University of Colorado, is funded by the NASA Lunar Science Institute (via Cooperative Agreement NNA09DB30A) to investigate concepts for astrophysical observatories on the Moon. This work used the JANUS supercomputer, which is supported by the National Science Foundation (award number CNS-0821794) and the University of Colorado Boulder. The JANUS supercomputer is a joint effort of the University of Colorado Boulder, the University of Colorado Denver, and the National Center for Atmospheric Research.

APPENDIX A

OPTIMIZATION VIA SIMULATED ANNEALING

To solve Equation (24), we employ the Monte Carlo method of Simulated Annealing (Kirkpatrick et al. 1983; Černý 1985). For a given source and n_v , we run K Monte Carlo trials, each consisting of L steps, aimed at determining the optimal values of I_n and ν_n for n_v frequency bins. We do not require the bolometric luminosity of sources to be conserved (i.e., $\sum_{n=1}^{n_v} I_n \neq 1$ is allowed), since some photons may traverse the entire one-dimensional “volume” without ionizing a single atom, or some fraction of the luminosity may be emitted below the hydrogen ionization threshold. Inclusion of such photons would be computational effort wasted in a fully three-dimensional ray-tracing calculation, for example, since their mean free paths are very long, and once absorbed they may contribute negligibly to ionization and heating.

Each random walk begins with randomly generated values of ν_n distributed between the hydrogen ionization threshold and the maximum emission frequency in the spectrum, and randomly generated values of I_n that sum to unity. Subsequent steps vary the energy or normalization of (randomly chosen) frequency bin n . In order to steer each random walk toward the global

minimum, we first evaluate the quantity

$$P = \exp[-(f_{k,l} - f_{k,l-1})/T_{SA}], \quad (\text{A1})$$

where $k = 0, 1, 2, \dots, K$ represents the current step in the current random walk, l , where $l = 0, 1, 2, \dots, L$, and f is the “cost function,” a measure of how good our current solution is. We adopt a cost function which is the sum of errors in Φ_i and Ψ_i over the column density range of interest. For each species (i), and each integral quantity (Φ , Ψ), we add the maximum deviation from continuous and discrete solutions in the optically thin limit (first term in Equation (A2)), the maximum deviation over the entire column density range (second term in Equation (A2)), and the average deviation over the entire column density range (final term in Equation (A2)), all in dex, i.e.,

$$\begin{aligned} f_{k,l} = & \sum_i \sum_{\Lambda=\Phi,\Psi} \left\{ \max \left[\log \left(\frac{\Lambda_i}{\Lambda'_i(\nu_{k,l}, I_{k,l})} \right) \right]_{\tau=0} \right. \\ & + \max \left[\log \left(\frac{\Lambda_i}{\Lambda'_i(\nu_{k,l}, I_{k,l})} \right) \right]_{\tau>0} \\ & \left. + \left\langle \log \left(\frac{\Lambda_i}{\Lambda'_i(\nu_{k,l}, I_{k,l})} \right) \right\rangle_{\tau>0} \right\}. \quad (\text{A2}) \end{aligned}$$

At each step in a given random walk, we also generate a random number, $q \in [0, 1]$, that will determine whether we keep our current guess, $(\nu_{k,l}, I_{k,l})$, or revert to our previous guess, $(\nu_{k,l-1}, I_{k,l-1})$. The condition for keeping our current guess is $P \geq q$.

The key aspect of this analysis is how we vary the control parameter T_{SA} , which is called the temperature in analogy with Boltzmann’s equation (we add the subscript SA to distinguish the gas kinetic temperature from this unphysical Simulated Annealing temperature). Equation (A1) tells us that regardless of the value of T_{SA} , if $f_{k,l} < f_{k,l-1}$ (i.e., our most recent guess is better than the last), then $P \geq 1$, and we have a 100% chance of keeping our current guess. In other words, our method of controlling the T_{SA} only effects how we deal with bad guesses—decreasing the temperature means we become less tolerant of bad guesses. There are many ways of doing this (Press et al. 1992), but for simplicity we adopt the following technique. Every s/n_v steps per frequency bin, we take

$$T \rightarrow \lambda T, \quad (\text{A3})$$

where λ is an experimentally determined quantity of the order of unity. For all results presented here, we have adopted $\lambda = 0.98$, and $s/n_v = 10$. We change the number of steps per random walk depending on the dimensionality, $2n_v$. We have found through experimentation that a good rule of thumb is $L = 5000$ steps per trial, K , per frequency bin n_v for our choice of λ and s/n_v . These control parameters are fairly conservative—further experimentation with them may yield converged solutions for fewer trials, K , and steps, L .

APPENDIX B

CODE VERIFICATION

Our one-dimensional radiative transfer code solves Equations (1)–(4) using the implicit Euler method for integration and a Newton–Raphson technique for root finding. Each simulation is initialized on a grid of N_c cells between L_0 and L_{box} ,




Figure 12. Test 1: (a) comparison of the numerical (dashed) and analytic (solid) solutions for the position of an expanding ionization front as a function of time in a hydrogen-only, isothermal medium (RT06 problem 1; top), and the ratio of the calculated and analytic solutions as a function of time and grid resolution (bottom). The numerical solution displayed in the top panel is from the highest resolution simulation (800 grid cells, i.e., $\Delta x = L_{\text{box}}/800$). (b) Radial profiles of the neutral (solid) and ionized (dashed) fractions at $t = 10, 100$, and 500 Myr.

(A color version of this figure is available in the online journal.)

such that the finest resolution element is $\Delta x = (L_{\text{box}} - L_0)/N_c$, or simply $\Delta x = 1/N_c$ in code units. Gas inside of the start radius, L_0 , contributes no optical depth, and Equations (1)–(4) are not solved. For the purposes of this section, we chose to use N_c linearly spaced cells between L_0 and L_{box} , though our code allows arbitrarily structured grids.

In order to track the propagation of ionization fronts accurately, we limit the time step based on a maximum neutral fraction change as introduced in Shapiro et al. (2004)

$$\Delta t_i = \epsilon_{\text{ion}} \frac{n_i}{|dn_i/dt|}, \quad (\text{B1})$$

where we include all absorbing species, $i = \text{H I}, \text{He I}, \text{He II}$, and set $\Delta t = \min(\Delta t_i)$. We additionally require that the time step increase by a factor of two at most, as in Wise & Abel (2011). For all simulations presented in this work, we have set $\epsilon_{\text{ion}} = 0.05$.

The primary solver implemented in our code assumes that the speed of light is infinite. Such an algorithm is appealing for two main reasons, aside from the fact that it is a very good approximation for the problems presented in this work. First, treating the speed of light explicitly introduces additional computational overhead as “photon packages” must be launched from the radiation source at each time step and tracked until they exit the domain. In the earliest stages of I-front propagation, the time step can be very small (as required by Equation (B1)), meaning the total number of photon packages, N_p , will be much larger than the total number of grid cells, N_c . Whereas $c = \infty$ treatments only require Equations (1)–(4) to be solved once per cell, finite speed-of-light treatments require this system of equations to be solved for each photon package. At later times, when $N_p < N_c$, solving the ion and heat equations is cheaper for finite speed-of-light treatments, though this offers no real advantage since the majority of the computational expense is at early times when I-front propagation is fastest. We have also included a finite c solver to accommodate a broader class of problems that may be of interest in future work.

The second advantage of assuming $c = \infty$ is that it allows the code to be efficiently parallelized. If $c = \infty$, cells in the domain can be solved in arbitrary order by a single

processor, or simultaneously by a network of processors, since the radiation incident on any cell is predetermined at the outset of each individual time step. Previous authors have ensured causality by solving cell k before cell $k + 1$ at time t (where increasing k corresponds to increasing r), but this is not in fact necessary—causality is ensured by the monotonicity of column density with distance. In other words, when $c = \infty$, N_i does not change within any given time step, and so the column density (and thus radiative flux) to cell k is less than the column density (and flux) to cell $k + 1$, meaning the solution of Equations (1)–(4) in cell $k + 1$ is completely independent of the properties of cell k at time $t + \Delta t$.

To demonstrate the functionality of the code, we repeat tests 1 and 2 from the Radiative Transfer Comparison Project (Iliev et al. 2006, hereafter referred to as RT06) on a grid of 200 linearly spaced cells. Test 1 is the expansion of an H II region in a hydrogen-only, isothermal medium surrounding a monochromatic source of 13.6 eV photons. We adopt the same parameters used in RT06: constant temperature $T = 10^4$ K, uniform hydrogen number density $n_{\text{H}} = 10^{-3} \text{ cm}^{-3}$, ionized fraction $x = 1.2 \times 10^{-3}$, in a box $L_{\text{box}} = 6.6$ kpc in size, and with photon luminosity $\dot{Q} = 5 \times 10^{48} \text{ s}^{-1}$. The classical analytic solution for the radius of an ionization front is

$$r_{\text{IF}}(t) = r_s(1 - e^{-t/t_{\text{rec}}})^{1/3}, \quad (\text{B2})$$

where r_s is the Strömgren radius,

$$r_s = \left(\frac{3\dot{Q}}{4\pi\alpha_{\text{H II}}n_{\text{H}}^2} \right)^{1/3}, \quad (\text{B3})$$

and the recombination time, t_{rec} , is defined as

$$t_{\text{rec}} \equiv \frac{1}{\alpha_{\text{H II}}n_{\text{H I}}}. \quad (\text{B4})$$

This solution is approximate even in isothermal media, given that it assumes a constant neutral hydrogen density, $n_{\text{H I}}$. More accurate analytic solutions exist (Osterbrock & Ferland 2006),




Figure 13. Test 2: (a) radial profiles of the neutral (solid) and ionized (dashed) fractions at $t = 10, 100$, and 500 Myr. (b) Radial profiles of the kinetic temperature at $t = 10, 100$, and 500 Myr (solid, dashed, and dotted lines, respectively).

and predict a departure from the classical solution at $t/t_{\text{rec}} \simeq 1$, which grows to a $\sim 5\%$ difference by $t/t_{\text{rec}} \simeq 4$. Our numerical solution (see Figure 12(a)) captures this behavior very well. In Figure 12(b), we show radial profiles of the ionized and neutral fractions at three stages of the I-front expansion, which are again in very good agreement with the calculations presented in RT06.

Test 2 is the same as Test 1, except now the temperature is allowed to evolve according to Equation (4), and the monochromatic radiation source is replaced by a 10^5 K blackbody spectrum. Radial profiles of the neutral and ionized fractions and temperature can be seen in Figure 13. Again, our numerical solutions are in very good agreement with previous work.

REFERENCES

- Abel, T., Bryan, G. L., & Norman, M. L. 2002, *Science*, **295**, 93
 Abel, T., Norman, M. L., & Madau, P. 1999, *ApJ*, **523**, 66
 Abel, T., Wise, J. H., & Bryan, G. L. 2007, *ApJ*, **659**, L87
 Alvarez, M. A., Wise, J. H., & Abel, T. 2009, *ApJ*, **701**, L133
 Aubert, D., & Teyssier, R. 2008, *MNRAS*, **387**, 295
 Bondi, H., & Hoyle, F. 1944, *MNRAS*, **104**, 273
 Bryan, G. L., & Norman, M. L. 1997, arXiv:astro-ph/9710187
 Cantalupo, S., & Porciani, C. 2011, *MNRAS*, **411**, 1678
 Černý, V. 1985, *J. Optim. Theory Appl.*, **45**, 41
 Ciardi, B., Bolton, J. S., Maselli, A., & Graziani, L. 2012, *MNRAS*, **423**, 558
 Dale, J. E., Bonnell, I. A., Clarke, C. J., & Bate, M. R. 2005, *MNRAS*, **358**, 291
 Finlator, K., Özel, F., & Davé, R. 2009, *MNRAS*, **393**, 1090
 Friedrich, M. M., Mellema, G., Iliev, I. T., & Shapiro, P. R. 2012, *MNRAS*, **421**, 2232
 Fukugita, M., & Kawasaki, M. 1994, *MNRAS*, **269**, 563
 Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Phys. Rep.*, **433**, 181
 Furlanetto, S. R., & Stoever, S. J. 2010, *MNRAS*, **404**, 1869
 Gnedin, N. Y., & Abel, T. 2001, *New Astron.*, **6**, 437
 González, M., Audit, E., & Huynh, P. 2007, *A&A*, **464**, 429
 Hosokawa, T., Omukai, K., Yoshida, N., & Yorke, H. W. 2011, *Science*, **334**, 1250
 Iliev, I. T., Ciardi, B., Alvarez, M. A., et al. 2006, *MNRAS*, **371**, 1057
 Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. 1983, *Science*, **220**, 671
 Kramer, R. H., & Haiman, Z. 2008, *MNRAS*, **385**, 1561
 Krumholz, M. R. 2006, *ApJ*, **641**, L45
 Kuhlen, M., & Madau, P. 2005, *MNRAS*, **363**, 1069
 Maselli, A., Ferrara, A., & Ciardi, B. 2003, *MNRAS*, **345**, 379
 Mellema, G., Iliev, I. T., Alvarez, M. A., & Shapiro, P. R. 2006, *New Astron.*, **11**, 374
 Okamoto, T., Yoshikawa, K., & Umemura, M. 2012, *MNRAS*, **419**, 2855
 O'Shea, B. W., Bryan, G., Bordner, J., et al. 2004, arXiv:astro-ph/0403044
 Osterbrock, D. E., & Ferland, G. J. 2006, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei* (2nd ed.; Mill Valley, CA: Univ. Science Books)
 Petkova, M., & Springel, V. 2009, *MNRAS*, **396**, 1383
 Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1992, *Numerical Recipes in FORTRAN. The Art of Scientific Computing* (2nd ed.; Cambridge: Cambridge Univ. Press)
 Reynolds, D. R., Hayes, J. C., Paschos, P., & Norman, M. L. 2009, *J. Comput. Phys.*, **228**, 6833
 Ricotti, M., Gnedin, N. Y., & Shull, J. M. 2002, *ApJ*, **575**, 33
 Ricotti, M., Ostriker, J. P., & Gnedin, N. Y. 2005, *MNRAS*, **357**, 207
 Schaerer, D. 2002, *A&A*, **382**, 28
 Shapiro, P. R., Iliev, I. T., & Raga, A. C. 2004, *MNRAS*, **348**, 753
 Shull, J. M., & van Steenberg, M. E. 1985, *ApJ*, **298**, 268
 Shull, J. M., & Venkatesan, A. 2008, *ApJ*, **685**, 1
 Thomas, R. M., & Zaroubi, S. 2008, *MNRAS*, **384**, 1080
 Trac, H., & Cen, R. 2007, *ApJ*, **671**, 1
 Trac, H., Cen, R., & Loeb, A. 2008, *ApJ*, **689**, L81
 Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *ApJS*, **192**, 9
 Venkatesan, A., & Benson, A. 2011, *MNRAS*, **417**, 2264
 Venkatesan, A., Giroux, M. L., & Shull, J. M. 2001, *ApJ*, **563**, 1
 Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, *ApJ*, **465**, 487
 Whalen, D., & Norman, M. L. 2008, *ApJ*, **673**, 664
 Wise, J. H., & Abel, T. 2011, *MNRAS*, **414**, 3458