

Ly α COOLING EMISSION FROM GALAXY FORMATION

CLAUDE-ANDRÉ FAUCHER-GIGUÈRE^{1,2,5}, DUŠAN KERES^{1,2,6}, MARK DIJKSTRA^{1,3}, LARS HERNQUIST¹, AND MATIAS ZALDARRIAGA⁴

¹ Department of Astronomy, Harvard University, Cambridge, MA 02138, USA; cgiguere@berkeley.edu

² Department of Astronomy and Theoretical Astrophysics Center, University of California, Berkeley, CA 94720-3411, USA

³ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Strasse 1, 85748 Garching, Germany

⁴ School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540, USA

Received 2010 May 17; accepted 2010 October 2; published 2010 November 19

ABSTRACT

Recent numerical and analytical studies have shown that galaxies accrete most of their baryons via the cold mode, from streams with temperatures $T \sim 10^4\text{--}10^5$ K. At these temperatures, the streams should radiate primarily in the Ly α line and have therefore been proposed as a model to power the extended, high-redshift objects known as Ly α blobs, and may also be relevant for powering a range of less luminous Ly α sources. We introduce a new Ly α radiative transfer code, αRT , and calculate the transport of the Ly α emission from cold accretion in cosmological hydrodynamical simulations. In this paper, we describe our methodology, and address physical and numerical issues that are critical to making accurate predictions for the cooling luminosity, but that have been mostly neglected or treated simplistically so far. In particular, we highlight the importance of self-shielding and of properly treating sub-resolution models in numerical simulations. Most existing simulations do not self-consistently incorporate these effects, which can lead to order-of-magnitude errors in the predicted cooling luminosity. Using a combination of post-processing ionizing radiative transfer and re-simulation techniques, we develop an approximation to the consistent evolution of the self-shielded gas. We quantify the dependence of the Ly α cooling luminosity on halo mass at $z = 3$ for the simplified problem of pure gas accretion embedded in the cosmic radiation background and without feedback, and present radiative transfer results for a particular system. While pure cooling in massive halos (without additional energy input from star formation and active galactic nuclei) is in principle sufficient to produce $L_{\alpha} \sim 10^{43}\text{--}10^{44}$ erg s $^{-1}$ blobs, this requires including energy released in gas of density sufficient to form stars, but which is kept 100% gaseous in our optimistic estimates. Excluding emission from such dense gas yields lower luminosities by up to one to two orders of magnitude at high masses, making it difficult to explain the observed Ly α blobs with pure cooling. Resonant scattering produces diffuse Ly α halos, even for centrally concentrated emission, and broad double peaked line profiles. In particular, the emergent line widths are in general not representative of the velocity dispersion within galactic halos and cannot be directly used to infer host halo masses.

Key words: cosmology: theory – galaxies: evolution – galaxies: formation – galaxies: high-redshift – intergalactic medium – radiative transfer

Online-only material: color figures

1. INTRODUCTION

Steidel et al. (2000) discovered two extremely bright, large, and diffuse Ly α -emitting “blobs” in a narrowband survey of the SSA 22 proto-cluster region at $\langle z \rangle = 3.09$ (for the early detection of extended Ly α emission at $z \sim 2.4$, see also Francis et al. 1996, 2001; Keel et al. 1999). These two blobs, labeled LAB1 and LAB2, have physical extent $\gtrsim 140$ kpc, are more luminous ($L_{\text{Ly}\alpha} \sim 10^{44}$ erg s $^{-1}$) than typical line emitters at the same redshift by a factor $\sim 10\text{--}100$, and unlike similar halos around radio galaxies have no detectable radio continuum. Since their discovery, these blobs have become some of the most spectacular of a new class of sources that now counts several tens of members (e.g., Matsuda et al. 2004; Palunas et al. 2004; Francis et al. 2004; Dey et al. 2005; Nilsson et al. 2006; Saito et al. 2006; Smith & Jarvis 2007; Prescott et al. 2008, 2009; Ouchi et al. 2009; Yang et al. 2009, 2010) and whose nature remains unclear. Because detecting spatially extended Ly α emission usually requires deep narrowband imaging, which covers thin redshift slices, only a small volume the universe has been effectively surveyed for them to date; the blobs and their

fainter analogs are therefore likely to be numerous and cosmologically significant. In fact, existing observations and theoretical models indicate that they may be the sites of massive galaxy formation and may also display signatures of associated active galactic nuclei (AGNs) and/or supernova feedback. The Ly α blobs thus provide a unique opportunity of probing the processes driving and regulating galaxy formation, and may be intimately related to phenomena including proto-clusters, mergers, and submillimeter galaxies.

Although the most extreme Ly α blobs have received the most attention, there in fact exists a wide continuum of spatially extended Ly α sources at high redshift, for example the ones discovered by Saito et al. (2006) with line luminosities $\sim 10^{42}$ erg s $^{-1}$ and those discovered by Rauch et al. (2008), with line luminosities as low as $\sim 10^{39}$ erg s $^{-1}$. Understanding the nature of these fainter but more numerous sources is equally important to develop a physical picture of galaxy formation. In fact, the fainter sources likely probe different (perhaps earlier) stages of galaxy assembly.

A central puzzle for the Ly α blobs is that many of them do not appear to have a central source energetic enough to power their entire Ly α emission (e.g., Matsuda et al. 2004; Nilsson et al. 2006). Even for the blobs which do have energetically sufficient counterparts (e.g., Geach et al. 2005, 2007, 2009;

⁵ Miller Fellow.

⁶ Hubble Fellow.

Webb et al. 2009), it is unclear whether that energy can actually couple effectively to the Ly α emission. Submillimeter starbursts and obscured AGNs imply the presence of large quantities of dust, which acts to destroy Ly α photons particularly efficiently (Neufeld 1990). It is therefore uncertain whether Ly α photons produced by such dust-enshrouded sources can escape in significant amounts. Nevertheless, several different mechanisms have been proposed to power the blobs and can broadly be divided into three categories:

Embedded star formation or AGNs, possibly obscured from direct view by dust, could photoionize the surrounding hydrogen nebula (Møller & Warren 1998; Haiman & Rees 2001; Weidinger et al. 2004, 2005; Laursen & Sommer-Larsen 2007).

Superwinds driven by starburst supernovae could explain the observed sizes and kinematics of the blobs, with the Ly α emission being generated in the swept up material (Taniguchi & Shioya 2000; Taniguchi et al. 2001; Ohya et al. 2003; Mori et al. 2004; Wilman et al. 2005; Geach et al. 2005). AGNs could also drive similar winds (e.g., Murray et al. 2005; Springel et al. 2005; Di Matteo et al. 2005; Hopkins & Hernquist 2006).

Cooling radiation, emitted as gas, accretes onto forming galaxies, could produce luminous and extended structures. If a large fraction of the accreting gas has a temperature $T \sim 10^4$ – 10^5 K, then most of the cooling radiation could be Ly α (Fabian & Nulsen 1977; Katz & Gunn 1991; Hu 1992; Haiman et al. 2000; Fardal et al. 2001; Birnboim & Dekel 2003; Kereš et al. 2005; Furlanetto et al. 2005; Dijkstra et al. 2006; Yang et al. 2006; Dijkstra & Loeb 2009; Goerdt et al. 2010; Dayal et al. 2010).

In this work, motivated by recent progress on our understanding of how galaxies get their gas, we focus on the third possibility. The methods developed could, however, be applied to the first two classes of models as well, and we plan to extend our calculations to model those processes in the future.

In the classic sketch of galaxy formation (Rees & Ostriker 1977; Silk 1977; White & Rees 1978), gas falling into dark matter halos is shocked and heated to the virial temperature. For a galaxy with a mass similar to that of the Milky Way, the shocked gas attains a temperature $T_{\text{vir}} \sim 10^6$ K. In the dense inner regions of the halos, this gas efficiently radiates its thermal energy, loses its pressure support, and settles into compact disks where it can form stars. A wealth of recent work however suggests that this picture requires an important modification: most of the gas is never strongly shocked as it flows toward the central forming galaxy, but rather accretes in a “cold mode,” maintaining a temperature $T \lesssim 10^5$ K. Moreover, this cold accretion proceeds through dense filaments rather than in a spherically symmetric fashion.

Although the importance of cold accretion in galaxy formation has only recently been demonstrated in high-resolution three-dimensional hydrodynamical simulations (e.g., Fardal et al. 2001; Katz et al. 2003; Kereš et al. 2005, 2009b; Ocvirk et al. 2008; Brooks et al. 2009; Dekel et al. 2009), Binney (1977) had argued on the basis of analytic models of proto-galaxy collapse that the amount of shock heating could be small for plausible physical conditions. Moreover, already in the first simulations of forming galaxies, most of the gas never heated above $T \sim 3 \times 10^4$ K (Katz & Gunn 1991) and the importance of filamentary structures was recognized by Katz & White (1993) and Katz et al. (1994). Birnboim & Dekel (2003) carried out a stability analysis, supported by one-dimensional hydrodynamical simulations (see also Birnboim et al. 2007; Dekel & Birnboim 2006), and found that when the radiative cooling is ef-

ficient compared with the infall rate, the post-shock gas becomes unstable and cannot support the shock. When applied to cosmology, their results agree well with those of three-dimensional simulations.

The Ly α emission from cold accretion has already been the subject of some studies. Fardal et al. (2001) first evaluated the Ly α cooling luminosity from hydrodynamical simulations and suggested that it could account for the Ly α blobs discovered by Steidel et al. (2000). Haiman et al. (2000) reached a similar conclusion using simplified analytic arguments. Also using simulations, Furlanetto et al. (2005) found that the Ly α cooling radiation from structure formation could account for some, but not all, of the luminosity of the Ly α blobs. Yang et al. (2006) studied both the hydrogen and helium cooling radiation using simulations, but found the hydrogen Ly α luminosity to be strongly dependent on the self-shielding correction applied. Recently, Dijkstra & Loeb (2009) developed an analytic model and suggested that cooling radiation from the cold mode could account for all the Ly α blobs under reasonable assumptions. Using adaptive mesh refinement (AMR) simulations, Goerdt et al. (2010) provided supporting evidence for this picture. In our discussion (Section 5), we will contrast our main results with those of Goerdt et al. (2010), concluding that the differences with theirs most likely originate from the treatments of self-shielding and sub-resolution modeling, which are a focus of our study.

No study focusing specifically on cooling emission has however combined realistic Ly α radiative transfer with hydrodynamical simulations before.⁷ Because Ly α photons resonantly scatter, the resultant morphologies, spatial extents, and spectra are strongly modified by radiative transfer effects (e.g., Dijkstra et al. 2006). Since the Ly α photons tend to follow paths of least resistance in space and frequency (Section 4), the spatial geometry and bulk velocity fields play critical roles in determining the radiation transport (for observational evidence of these effects, see, e.g., Kunth et al. 1998; Mas-Hesse et al. 2003). Fully three-dimensional calculations are therefore necessary to make realistic predictions. Moreover, both the existing analytical and numerical studies have limitations that could induce important errors in quantities as basic as the integrated Ly α luminosity of the cold streams. One such uncertainty arises from the exponential dependence of the Ly α emissivity on the gas temperature for the temperatures $T \sim 10^4$ K that are characteristic of cold accretion (Section 3). At present, most galaxy formation simulations do not self-consistently predict the temperature distribution within the streams. In fact, existing simulations usually do not follow the transport of the ultraviolet (UV) radiation that ionizes and heats dense gas. This seriously limits the predictive power of these calculations, since small errors in the temperatures can result in large errors in the Ly α cooling emission, and potentially grossly violate energy conservation. As we will demonstrate, models of sub-resolution physics in hydrodynamical simulations can also introduce large errors if not properly taken into account. Analytical studies based on energetic considerations are not as sensitive to the temperature of the cold streams (e.g., Dijkstra & Loeb 2009), but are not immune of uncertainties either, since they rely on assumptions regarding the efficiency of Ly α emission. Moreover, their simplified nature does not lend itself to detailed radiative transfer predictions.

⁷ Other authors have included a cooling component in radiative transfer calculations of Ly α -emitting galaxies (e.g., Tasitsiomi 2006a; Laursen & Sommer-Larsen 2007; Laursen et al. 2009a, 2009b), but have not explicitly separated out the signatures of pure cooling or investigated the important uncertainties in detail.

Resolving these issues is critical to relating the Ly α emission from cold accretion to observations.

Our ultimate goal is a systematic investigation of the Ly α emission from galaxy formation that is both detailed in its predictions, and robust. By detailed analysis, we envision predictions that can be directly compared with observations, and therefore require both three-dimensional hydrodynamical simulations and realistic radiative transfer. By robust, we mean that the predictions should be free of assumptions that introduce the kind of large uncertainties that existing studies are subject to. Due to the complexity of the problem, this ultimate goal is likely to require a long-term effort. The present paper is dedicated to laying down some of the foundations for this research program. Specifically, we present a new Ly α radiative transfer code, named αRT , and describe its application to the cooling radiation in cosmological simulations of galaxy formation (for other applications of Ly α radiative transfer codes to hydrodynamical simulations, see, e.g., Cantalupo et al. 2005; Tasitsiomi 2006a; Laursen & Sommer-Larsen 2007; Laursen et al. 2009a; Kollmeier et al. 2010; Zheng et al. 2010a). We pay particular attention to clarifying the physical and numerical uncertainties of these calculations, in particular with respect to the predicted Ly α luminosities, and illustrate the importance of radiative transfer effects. To do so, aside for calculating the Ly α luminosities of a sample of halos from a cosmological volume, we focus our radiative transfer calculations on a particular system at $z = 3$ and explore variations in both the emission and radiative transfer physics. We limit ourselves to the most basic physical problem of accreting halos embedded in a cosmic ionizing background and neglect feedback processes. Follow up studies will build on the results obtained here and investigate the properties of the Ly α emission as a function of halo mass and redshift, and will extend them by incorporating additional physics, including feedback (C.-A. Faucher-Giguère et al. 2010, in preparation).

We begin by describing our hydrodynamical simulations in Section 2. We address the emission of Ly α photons in Section 3 and explicitly demonstrate the sensitivity of the predicted Ly α cooling luminosity on assumptions regarding the thermal state of the self-shielded gas. Using a combination of post-processing ionizing radiative transfer and re-simulation techniques, we develop an approximation to consistently model the evolution of the dense gas in the hydrodynamical simulations, resulting in the most robust numerical predictions to date. In Section 4, we present the results of Ly α radiative transfer calculations for a particular system of total mass $M_h = 2.5 \times 10^{11} M_\odot$ at $z = 3$ and highlight the role of both the bulk velocity flows and of the resonant scatters in shaping the emergent morphology and spectrum. Finally, we discuss our results and conclude in Section 5. The Appendices document the radiative transfer code αRT introduced in this work, the ionizing radiative transfer method, and relevant analytical estimates.

Throughout, we assume a cosmology with $(\Omega_m, \Omega_b, \Omega_\Lambda, h, \sigma_8, n_s) = (0.28, 0.046, 0.72, 0.70, 0.82, 0.96)$, as inferred from the *Wilkinson Microwave Anisotropy Probe* (WMAP) five-year data in combination with baryon acoustic oscillations and supernovae (Komatsu et al. 2009). While some of our hydrodynamical simulations were run with slightly different parameters, none of our conclusions are sensitive to the details of the cosmology. We assume hydrogen and helium mass fractions of $X = 0.75$ and $Y = 0.25$ (e.g., Burles et al. 2001), the collisional ionization coefficients given in Katz et al. (1996), the Ly α collisional excitation coefficient and average number

Table 1
Symbols Used in This Work

Symbol	Definition
n_i	Number density of species i
N_i	Column density of species i
T	Gas temperature
τ_ν	Optical depth at frequency ν
ν_0	Ly α central frequency
$\Delta\nu_D$	Ly α Doppler width
x	Dimensionless frequency offset $(\nu - \nu_0)/\Delta\nu_D$
Γ_i	Photoionization rate of species i
$\Gamma_{i,c}$	Collisional ionization coefficient of species i
$C_{Ly\alpha}$	Ly α collisional excitation coefficient
ϵ_α	Ly α emissivity
$\alpha_i^{A,B}$	Case A, B recombination coefficients to species i
$x_{H\text{I}}, x_{H\text{II}}$	Fractions of hydrogen in H I, H II
$y_{He\text{I}}, y_{He\text{II}}, y_{He\text{III}}$	Fractions of helium in He I, He II, He III

of Ly α photons produced per recombination from Osterbrock & Ferland (2006), and the recombination coefficients in the appendix of Hui & Gnedin (1997). For convenience, some symbols used in this work are defined in Table 1.

2. SIMULATIONS

2.1. Code Details

We compute the hydrodynamics of forming galaxies in a Λ CDM universe using a modified version of the GADGET cosmological simulation code (Springel 2005). The calculation of the gravitational force uses a combination of the particle mesh algorithm (e.g., Hockney & Eastwood 1988) for large separations and the hierarchical tree algorithm (e.g., Barnes & Hut 1986; Hernquist 1987) at small distances. The gas dynamics is calculated using a smoothed particle hydrodynamics (SPH) algorithm (e.g., Lucy 1977; Gingold & Monaghan 1977) that conserves both energy and entropy (Springel & Hernquist 2002). The modifications with respect to the public version of the code include the treatment of cooling, the effects of a uniform ultraviolet background (UVB), and a multiphase star formation algorithm as in Springel & Hernquist (2003). Star formation is implemented by the stochastic spawning of collisionless star particles by the gas particles. In practice, star formation in the multiphase model occurs above a density threshold of $n_H = 0.13 \text{ cm}^{-3}$ and is calibrated to the observed Kennicutt (1998) law, although it plays only a tangential role in this work, which focuses the cooling emission. The thermal and ionization properties of the gas are calculated including all relevant processes in a plasma with primordial abundances of hydrogen and helium following Katz et al. (1996).

2.2. Simulation Parameters and Halo Identification

We use two types of simulations. To achieve high resolution, we “zoom in” on individual halos within a larger simulation box and only follow the local gas dynamics at the refined resolution. This is done by first running a dark matter only simulation and selecting halos of interests. The simulation is then rerun including gas particles, with 8 times the original mass resolution, in a Lagrangian volume surrounding the halo of interest (e.g., Katz & White 1993). In this work, we focus on zoom-in simulations of an individual halo (labeled A1) selected from a volume of side length $10 h^{-1}$ comoving Mpc. The A1 halo has a total mass $2.5 \times 10^{11} M_\odot$ at $z = 3$. As it is also important

Table 2
Hydrodynamical Simulations

Name	L (h^{-1} Mpc) ^a	N^b	ϵ (h^{-1} kpc) ^c	Description
z10n128_A1 ^d	10	2×256^3	0.8	$8 \times$ mass refinement zoom
z10n128_A1_noSF	10	2×256^3	0.8	No star formation
z10n128_A1_ssUV	10	2×256^3	0.8	UVB off at $n_H > 0.01 \text{ cm}^{-3}$
z10n128_A1_ssUV_noSF	10	2×256^3	0.8	UVB off at $n_H > 0.01 \text{ cm}^{-3}$, no SF
gdm40n512 ^e	40	2×512^3	1.6	Full box
gdm40n512_noSF	40	2×512^3	1.6	No star formation
gdm40n512_ssUV	40	2×512^3	1.6	UVB off at $n_H > 0.01 \text{ cm}^{-3}$
gdm40n512_ssUV_noSF	40	2×512^3	1.6	UVB off at $n_H > 0.01 \text{ cm}^{-3}$, no SF

Notes.

^a Comoving box side length.

^b Effective total number of dark matter+gas particles in the box, after zoom refinement.

^c Comoving Plummer equivalent gravitational softening length, after zoom refinement.

^d Dark matter, gas, and stellar particle masses are $4 \times 10^6 h^{-1} M_\odot$, $8 \times 10^5 h^{-1} M_\odot$, and $4 \times 10^5 h^{-1} M_\odot$, respectively.

^e Dark matter, gas, and stellar particle masses are $3 \times 10^7 h^{-1} M_\odot$, $6 \times 10^6 h^{-1} M_\odot$, and $3 \times 10^6 h^{-1} M_\odot$, respectively.

to understand the trends and variance between different halos, we simulated an entire cosmological volume consisting of a cubical box with a side length of $40 h^{-1}$ comoving Mpc. While the resolution in this volume is more limited by computational constraints, it provides us with a large number of halos of different masses. Table 2 lists the simulations used in this work and their parameters. The minimum gas smoothing length is set to 0.1 of the gravitational softening in all the simulations. Given the importance of self-shielding (Section 3), we rerun our simulations with exactly the same parameters, but with the UVB artificially turned off in regions exceeding a certain density (suffix *_ssUV*), to be discussed in Section 3.2. All the simulations are also rerun with the star formation model turned off. The simulations without star formation are identified by the additional suffix *_noSF*. The A1 zoom-in simulations assume a variant of the Haardt & Madau (1996) model of the UVB, while our cosmological simulations use the more recent model of Faucher-Giguère et al. (2009).

A friends-of-friends (FoF) algorithm (e.g., Davis et al. 1985) with linking length set to $b = 0.2$ in units of the mean inter-particle separation is used to identify the dark matter halos in the simulations. The total mass of the particles within each FoF group, M_{FoF} , corresponds approximately to the mass in a sphere of mean interior density 180 times the background matter density, M_{180b} (White 2002). We therefore define the virial radius of a halo as the radius of a sphere containing M_{180b} , $r_{\text{vir}} \equiv r_{180b} \approx [M_{\text{FoF}} / (240\pi\rho_u(z))]^{1/3}$, where $\rho_u(z) = \rho_{\text{crit}}\Omega_m(1+z)^3$ and ρ_{crit} is the critical density at $z = 0$. In what follows, we use M as a shorthand for $M_{\text{FoF}} \approx M_{180b}$. The center of each halo is determined as the point deepest in the gravitational potential. Since we run the FoF algorithm on the dark matter particles only, we multiply the returned masses by $\Omega_m/(\Omega_m - \Omega_b)$ to account for the baryons.

As we want to isolate the cold accretion cooling radiation, the simulations studied in this work do not include galactic winds or AGN feedback. It is of course likely that the results would be somewhat modified if these processes were included. We plan to investigate the effects of feedback and their observational manifestations in future work.

2.3. Ionization and Thermal Structure

Our basic SPH simulations, like most cosmological simulations to date, assume a UVB that is spatially uniform throughout the simulation volume and calculate the local ionization

state of the gas assuming photoionization equilibrium with this background. This approach misses the effects of self-shielding; where the optical depth to ionizing photons from the background is of order unity or more, the gas would in reality be exposed to an attenuated ionizing field. This has consequences for both the ionization state and the thermal properties of the gas. Indirectly, the gas dynamics is also affected by the modification of pressure forces.

The omission of self-shielding implies that (in the absence of local sources) dense gas in the simulations sees an ionizing flux stronger than in reality. As a result, the gas tends to be overionized. This is important for the Ly α radiative transfer problem for two reasons. First, the Ly α emission mechanisms which seed the Ly α photons depend not on the total gas density, but on the number densities of ions (Section 3.1). Second, the transport of Ly α photons depends on the *neutral* hydrogen distribution, as only this ion provides scattering opacity (Section 4).

Self-shielding also has important effects on the thermal evolution of the gas. As Fardal et al. (2001) pointed out, neglecting self-shielding in a simulation with a prescribed uniform UVB introduces an artificially high rate of photoheating in dense regions and thus results in overestimated temperatures in these regions. Moreover, the presence of a penetrating ionizing background suppresses the cooling function (e.g., Efstathiou 1992; Katz et al. 1996; Weinberg et al. 1997; Kereš et al. 2005; Wiersma et al. 2009), and these effects could be amplified by the lack of a dynamical response of the gas to cooling, which would tend to make it denser and hence to cool even more rapidly. Although these effects are critical to accurately predicting the Ly α cooling luminosity (Section 3), previous Ly α studies have either neglected them, made simplifying but not necessarily correct assumptions (Fardal et al. 2001; Goerdt et al. 2010), or have explored a range of prescriptions (e.g., Furlanetto et al. 2005; Yang et al. 2006).

We improve significantly over previous work by performing ionizing radiative transfer in post-processing to identify the self-shielded gas, rather than relying on simplified criteria (Section 3.2). Since knowing the distribution of neutral hydrogen is a fundamental component of Ly α radiative transfer, post-processing ionizing radiative transfer has previously been used by other groups for related problems (e.g., Cantalupo et al. 2005; Laursen et al. 2009a; Kollmeier et al. 2010; Zheng et al. 2010a), but never before in focused studies of cooling emission. As we

will show, the predicted Ly α cooling luminosity is very sensitive to the state of the self-shielded gas. In order to obtain more robust predictions, we rerun our hydrodynamical simulations with the ionizing background turned off in regions above a certain density threshold (informed by our ionizing radiative transfer calculations) as an approximation to the consistent treatment of self-shielding.

For the moment, we pause to discuss the star-forming gas, which also requires special treatment.

2.4. The Multiphase ISM

The star-forming gas particles in the multiphase model carry effective ionic densities and temperatures that are mass-weighted averages of the hot and cold components (Springel & Hernquist 2003). While most of the mass in this model is in the $T = 1000$ K cold component,⁸ the hot component is in general much hotter ($T = 10^5$ – 10^8 K). This results in high effective temperatures with simultaneously large neutral fractions, and therefore in high collisional excitation rates (Section 3.1.2). As we will show, using the effective multiphase temperatures and ionic densities for the star-forming particles yields artificially high cooling luminosities, owing to the nonlinearity of the emissivity function with respect to density and temperature. Moreover, in the multiphase model, supernovae are responsible for pressurizing the interstellar medium (ISM) and are therefore an additional source of energy.

To obtain realistic results uncontaminated by feedback energy, it is necessary to exclude the star-forming particles from the cooling luminosity calculations. We explore two ways of doing this. First, we use the simulations with star formation, but ignore all the multiphase particles in the luminosity calculation. For this case, we assume that the star-forming particles are effectively optically thin for the purpose of transporting the Ly α photons. In reality, dust may destroy a portion of those photons, but we do not model this effect for this extreme prescription. This case approximates a multiphase medium in which cold neutral clumps embedded in a hot medium are either so compact that their covering factor is negligible, or in which they simply reflect the Ly α photons and contribute only a small effective Ly α optical depth (e.g., Neufeld 1991; Hansen & Oh 2006). A potential worry with this approach, however, is that it misses the cooling that occurs in particles with density above the star formation threshold. In our second approach, we make sure to capture all the cooling by using identical simulations but with the star formation model turned off, in which case we can simply sum over all the particles. The gas is allowed to become arbitrarily dense (as permitted by the resolution) and the absence of ionizing and mechanical feedback from embedded stars makes the density peaks very optically thick to Ly α photons. By comparison with the case of optically thin star-forming gas, this prescription therefore allows us to also quantify the effects of the opacity provided by star-forming particles on the transport problem.

We intentionally do not model the Ly α photons produced by stars in this work in order to separate out the properties of pure cooling emission; we briefly discuss their importance in the discussion (Section 5) and in Appendix A.

⁸ The terminology with respect to temperature is somewhat discrepant in the contexts of galaxy formation and of the ISM. While the $T = 10^4$ K gas is termed cold in galaxy formation and throughout most of this paper, it is usually qualified as warm in the context of the ISM to distinguish it from the much cooler, star-forming molecular gas.

3. Ly α EMISSION

3.1. Emission Processes

3.1.1. Recombination

Ionizing radiation (either from the cosmic background, local star formation, or an AGN) can photoionize gas that recombines and produces Ly α photons. Collisions in gas of sufficient density and temperature can also ionize hydrogen and be followed by the re-emission of Ly α photons via recombination. We group these two processes in “recombination emission.” Recombination emission produces Ly α photons at a rate (in units of $\text{ph s}^{-1} \text{cm}^{-3}$)

$$\epsilon_{\alpha}^{\text{ph,rec}} = f_{\alpha,\text{rec}} \alpha_{\text{H I}}^{\text{B}}(T) n_{\text{H II}} n_e, \quad (1)$$

where $\alpha_{\text{H I}}^{\text{B}}(T) \propto T^{-0.7}$ is the hydrogen case B recombination coefficient and $f_{\alpha,\text{rec}}$ is the average number of Ly α photons produced per case B recombination. For gas at $T = 10^4$ K that is optically thick to Lyman series transitions (so that higher order Lyman series recombination photons can ultimately be degraded into a Ly α photon), $f_{\alpha,\text{rec}} = 0.68$ (Osterbrock & Ferland 2006). This fraction is only weakly dependent on temperature and so we assume this constant value throughout.

3.1.2. Collisional Excitation

Collisions can also excite the Ly α line without ionizing hydrogen. The Ly α emissivity from this process is given by

$$\epsilon_{\alpha}^{\text{ph,coll}} = C_{\text{Ly}\alpha}(T) n_{\text{H I}} n_e, \quad (2)$$

where $C_{\text{Ly}\alpha}(T) \propto T^{-1/2} \exp(-h\nu_{\alpha}/kT)$ is the Ly α collisional excitation coefficient in units of $\text{ph cm}^3 \text{s}^{-1}$. Note that the Ly α recombination and collisional excitation emissivities scale differently with temperature. Moreover, whereas the recombination term is proportional to the H II number density, the collisional excitation term is proportional to the H I number density. The relative importance of the two processes will therefore depend on the local temperature and ionization state of the gas.

3.1.3. Limiting Equilibrium Cases

We assume ionization equilibrium, which is generally valid since the timescale $t_{\text{eq}} \equiv [\Gamma_{\text{H I}} + (\Gamma_{\text{H I},c}(T) + \alpha_{\text{H I}}^{\text{A}}(T))n_e]^{-1}$ to reach equilibrium is small compared to the dynamical timescale in both optically thin and self-shielded gas. The statistical equilibrium equation for hydrogen is

$$\Gamma_{\text{H I}} n_{\text{H I}} + \Gamma_{\text{H I},c}(T) n_e n_{\text{H I}} = \alpha_{\text{H I}}^{\text{A}}(T) n_e n_{\text{H II}}, \quad (3)$$

where $\Gamma_{\text{H I}}$ is the photoionization rate, $\Gamma_{\text{H I},c}(T)$ is the collisional ionization coefficient, and $\alpha_{\text{H I}}^{\text{A}}(T)$ is the case A recombination coefficient.

Physical intuition can be gained by considering the two limiting cases of photoionization equilibrium and of pure collisional ionization equilibrium (CIE; $\Gamma_{\text{H I}} \ll \Gamma_{\text{H I},c}(T)n_e$). As we will show, the two cases are directly relevant to our problem. While most of the cosmic volume is well approximated by the photoionization equilibrium regime, the dense cold gas (including the cold streams of interest) can self-shield from the external ionizing radiation and is then more accurately described by the pure CIE case. Figure 1 shows $\epsilon_{\alpha}^{\text{ph}}/n_{\text{H}}^2$ from both recombination and collisional excitation. The left panel shows the case of gas in ionization equilibrium with a photoionization rate $\Gamma_{\text{H I}} = 10^{-12} \text{ s}^{-1}$ (approximately the magnitude of the




Figure 1. Ly α emissivity per n_{H}^2 . Left: gas in ionization equilibrium with a hydrogen photoionization rate $\Gamma_{\text{H I}} = 10^{-12} \text{ s}^{-1}$. Right: gas in pure collisional ionization equilibrium. The solid curves show the Ly α emissivity from collisional excitation and the dashed curves show the emissivity from recombination. For the photoionization equilibrium case, the collisional excitation curve depends on n_{H} ; we show $\log(n_{\text{H}}/\text{cm}^{-3}) = -6, -4$, and -2 with curves of increasing thickness. The emissivity from collisional excitation scales with $n_{\text{H I}} \propto \Gamma_{\text{H I}}^{-1}$ in the photoionization equilibrium case. The unavoidable effects of collisional ionization are included in the photoionization equilibrium case. These curves assume that all the helium is in the form of He III.

cosmic ionizing background at $z \approx 2-4$, e.g., Faucher-Giguère et al. 2008a, 2008b), and the right panel shows the case of gas in collisional ionization equilibrium. The most important point to note is that the collisional excitation contribution is exponentially sensitive to temperature, at the temperatures $T \sim 10^4$ K characteristic of cold accretion streams, in both cases. As collisional excitation is the dominant Ly α cooling mechanism, it is critical to accurately capture the thermal state of the emitting gas. The curves shown in Figure 1 assume that all the helium is in the form of He III for simplicity. The emissivity from both processes is only weakly sensitive to the helium ionized fractions, except for temperatures $T \lesssim 10^4$ K for the CIE case and at very high densities in the photoionization case, when most of the hydrogen is neutral and helium dominates the free electrons. Since little Ly α emission originates from these regimes, our results are broadly insensitive to the helium ionization state, although we do solve for the correct helium ionization fraction in self-shielded regions in our simulations with the on-the-fly self-shielding approximation (Section 3.2.2).

In our simulations, the Ly α luminosity is evaluated directly from the SPH particles. Specifically, each particle is assigned an Ly α luminosity $L_{\alpha, \text{p}}^{\text{ph}} \equiv V_{\text{p}}(\epsilon_{\alpha}^{\text{ph, rec}} + \epsilon_{\alpha}^{\text{ph, col}})$, where the emissivity terms are evaluated using its density, ionization state, and temperature, and $V_{\text{p}} \equiv M_{\text{p}}/\rho_{\text{p}}$ is its volume, defined as the ratio of its mass to its density. This approach is desirable as it accurately takes into account the clumping of the gas on small scales, relevant to calculate the emission from the density-squared processes, and because it avoids artificial mixing that could occur if hot and cold phases were averaged in a gridding procedure. Such artificial mixing could boost the predicted Ly α luminosity by a large factor owing to the nonlinearity of the emissivity function.

3.2. Self-shielding

Since the emission processes scale with density squared (Equations (1)–(2)), the emissivity peaks in the densest regions, which are the most likely to self-shield. Because our hydrodynamical simulations lack proper ionizing radiative transfer,

they do not correctly capture self-shielding (Section 2.3). As explained in Section 2.4, naively integrating over multiphase SPH particles could also induce large errors in the predicted cooling luminosity. As we will show, *it is necessary to both exclude multiphase particles from the calculation and to model self-shielding* to accurately predict the cooling luminosity. The rest of this section is dedicated to demonstrating the importance of each potential source of error and to developing a consistent approximation to the Ly α cooling luminosity.

We follow the following steps.

1. Naively calculate the Ly α luminosity from a simulation with standard UV background and star formation treatments.
2. Using post-processing ionizing radiative transfer, identify the self-shielded gas in step 1.
3. Using the post-processed output, illustrate how the predicted Ly α luminosity depends on the assumed thermal state of the self-shielded gas.
4. Rerun a hydrodynamical simulation with the same initial conditions, but with the ionizing background turned off in self-shielded regions on the fly as an approximation to the self-consistent effects of self-shielding.
5. By rerunning identical simulations with star formation turned off, separate the effects of incorrectly including multiphase particles from those of ignoring self-shielding.

3.2.1. Post-processing Self-shielding

A technical description of our ionizing radiative transfer code is provided in Appendix B. Briefly, the hydrogen photoionization rate of the cosmic background, $\Gamma_{\text{H I}}^{\text{bkg}}$, is specified and taken as the boundary condition at the faces of the cubical radiative transfer volume. For the radiative transfer calculations, the simulation outputs are interpolated onto a Cartesian grid taking into account the smoothing kernels, with N_{p} grid points along each dimension. We employ the fiducial choice $N_{\text{p}} = 256$ and a radiative transfer volume of $(1 \text{ comoving Mpc } h^{-1})^3$, centered around each halo considered, which convergence tests suggest is sufficient (Section 4.3). Rays normal to each of the six faces

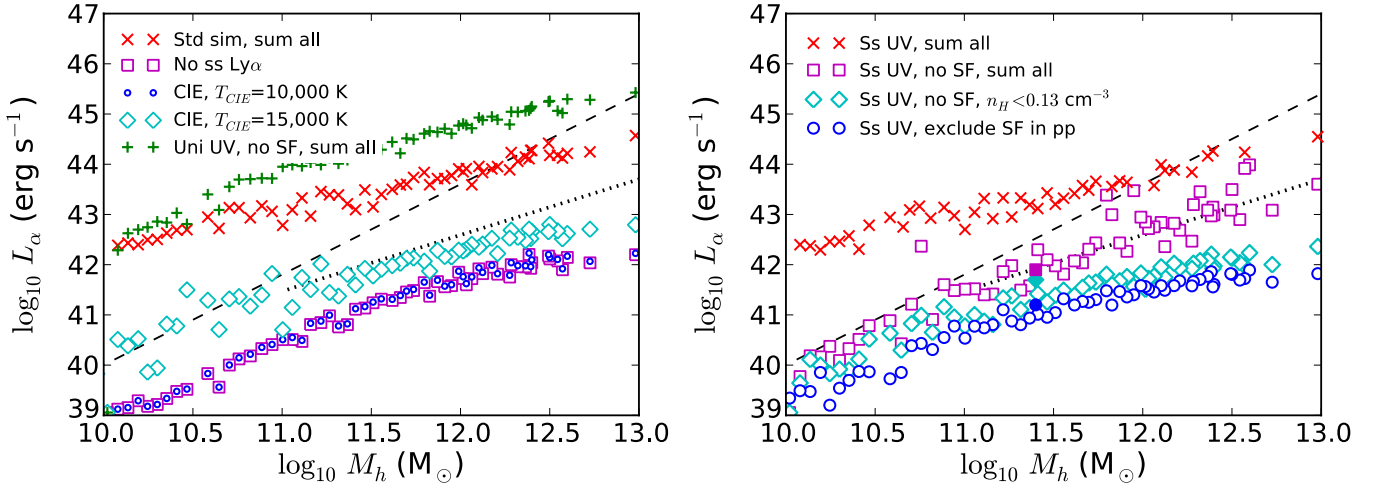


Figure 2. Ly α luminosity within the virial radius as a function of halo mass at $z = 3$ calculated from our cosmological volume simulations, for different physical assumptions. Each point corresponds to a randomly selected halo. The different cases are defined in Table 3, where corresponding remarks are given. Left: standard hydrodynamical simulation with uniform UVB and a multiphase star formation model, in some cases post-processed with the ionizing radiative transfer scheme to identify the self-shielded gas, corresponding to prescriptions 1, 2, 3, and 4. The green +s show the result of a simulation with a uniform UVB but no star formation, where the cooling luminosity is integrated over all the particles (prescription 5). Right: simulations with the ionizing background turned off in regions where $n_{\text{H}} > 0.01 \text{ cm}^{-3}$ to approximate self-shielding, corresponding to prescriptions 6, 7, 8, and 9. For the last three (most consistent) cases, the filled symbols show the results for the A1 halo in our zoom-in simulation with $8\times$ better mass resolution. Caution should be exercised when interpreting the quantitative details of the lowest mass halos shown, as the hydrodynamics may not be fully converged (Section 3.2.2). The dashed lines show an analytic estimate for the maximum average cooling luminosity available from the release of gravitational potential energy (Appendix A); the dotted lines show the more sophisticated analytic model of Dijkstra & Loeb (2009) for their fiducial parameter $f_{\text{grav}} = 0.3$.

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

are then sent inward and the optical depth to ionizing photons is calculated along each ray. Given the attenuated photoionization rate at each point, the ionization equilibrium is updated taking into account photoionization, collisional ionization, and recombination. The procedure is iterated until the ionized fraction has converged in all the cells. In solving for the equilibrium ionization balance, the gas temperatures used are those provided by the hydrodynamical simulation. These should be accurate in the optically thin regions and therefore our scheme should accurately capture the onset of self-shielding. In this post-processing treatment, the temperature structure will however be inaccurate in the self-shielded regions, since the modifications of the heating and cooling functions are not properly modeled in the hydrodynamical calculation. As outlined above, we address this in two ways: first, we explore a range of prescriptions for the self-shielded gas, illustrating the sensitivity of the predictions to these prescriptions; then we subsequently improve the accuracy of our calculations by approximating the self-consistent thermal evolution of the gas with simulations in which the ionizing background is switched off in dense regions. The self-shielded cells are defined as those that see an attenuated ionizing background, $\langle e^{-\tau_i} \rangle < \langle e^{-\tau_i} \rangle_{\text{crit}}$ with $\langle e^{-\tau_i} \rangle_{\text{crit}} \equiv 0.1$, where $\langle e^{-\tau_i} \rangle$ is the angle-averaged attenuation factor in the cell after post-processing. Since the optical depth rapidly increases within a self-shielded region, the results are weakly sensitive to the choice of the self-shielding threshold.

To investigate how the Ly α luminosity depends on assumptions regarding the self-shielded gas (steps 1, 2, and 3 above), we start with hydrodynamical simulations with standard treatments of the UV background and of star formation. Halos covering a broad range of masses are selected from our cosmological volume simulation gdm40n512 at $z = 3$ and the luminosity of each is defined as the sum of the luminosities of the gas particles contained within its virial radius. The same prescriptions are also applied to the A1 halo at $z = 3$, which will be used for

the Ly α radiative transfer calculations (Section 4). We explore three cases that are illustrative of the range of possibilities (for similar prescriptions, see Furlanetto et al. 2005).

Naively sum all the particles. In this simplistic prescription, we simply sum all the SPH particles within the virial radius. This includes dense particles that would in reality self-shield but that are artificially illuminated by a uniform ionizing background, and star-forming particles that carry effective multiphase values for their temperature and ionization state, which will introduce luminosity errors as described in Section 2.4.

No Ly α emission from self-shielded gas. Self-shielded gas may become neutral and cool substantially. If the gas is not photoionized and cools below $T \approx 10^4 \text{ K}$, its Ly α emissivity is severely suppressed (Figure 1). To illustrate how this case might differ from the naive calculation above, we model it by the extreme assumption that self-shielded gas does not produce Ly α photons at all. To transport the Ly α photons, we assume that this gas has a temperature $T = 10,000 \text{ K}$.

Collisional ionization equilibrium. Somewhat intermediate between the two above cases, self-shielded gas could settle to collisional ionization equilibrium (CIE) with a temperature $T_{\text{CIE}} \gtrsim 10,000 \text{ K}$ if gravitational heating is sufficiently efficient. In this case, the neutral hydrogen fraction is given by

$$x_{\text{H I}} = \frac{1}{1 + \Gamma_{\text{H I, c}}(T_{\text{CIE}})/\alpha_{\text{H I}}^{\text{A}}(T_{\text{CIE}})} \quad (4)$$

and the Ly α emissivity $\epsilon_{\alpha}/n_{\text{H}}^2$ is a well-defined function of temperature, as shown in the right panel of Figure 1. We explore how the luminosity depends on the prescribed CIE temperature, for $T_{\text{CIE}} = 10,000 \text{ K}$ and $15,000 \text{ K}$. As we will see in our simulations that approximate the effects of self-shielding, the self-shielded gas cools very effectively, so that higher temperatures are not expected. Furthermore, the proximity of $15,000 \text{ K}$ to the peak of the cooling curve implies that this case is already quite optimistic.

Table 3
Prescriptions for Calculating the Ly α Cooling Luminosity

No.	Description	Notes
1	Standard hydro (uniform UVB and multiphase SF model), sum all particles	Overestimates Ly α luminosity due to UVB and multiphase model
2	Standard hydro, with post-processing ionizing RT, no Ly α from self-shielded gas	Satisfies energetic bound
3	Standard hydro, with post-processing ionizing RT, CIE with $T_{\text{CIE}} = 10,000$ K	\approx same Ly α luminosity as 2
4	Standard hydro, with post-processing ionizing RT, CIE with $T_{\text{CIE}} = 15,000$ K	$\approx 10\times$ more Ly α luminous than 3
5	Hydro with uniform UVB but no SF, sum all particles	Overestimates Ly α luminosity due to UVB only
6	Self-shielding approx. for the UVB, with SF, sum all particles	Overestimates Ly α luminosity due to multiphase model only
7	Self-shielding approx. for the UVB, no SF, sum all particles	Consistent Ly α luminosity estimate
8	Self-shielding approx. for the UVB, no SF, only sum $n_{\text{H}} < 0.13 \text{ cm}^{-3}$	Consistent Ly α luminosity estimate
9	Self-shielding approx. for the UVB, SF excluded in post-processing, sum all particles	Consistent Ly α luminosity estimate

For the latter two cases, in which we assume either no emission from self-shielded gas or CIE, we exclude emission from the star-forming particles that might fall outside of the self-shielded regions, in order to avoid potential confusion with artificially high luminosities from multiphase particles. Table 3 summarizes the different prescriptions explored for calculating the Ly α cooling luminosity; the above cases are labeled 1–4.

The left panel of Figure 2 shows how the Ly α cooling luminosity varies with halo mass at the fiducial redshift $z = 3$ for these different prescriptions. The dashed curve shows an analytic estimate of the average maximum cooling luminosity achievable from the release of gravitational potential energy as a function of halo mass (see Appendix A; also Goerdt et al. 2010 for similar ideas). Briefly, the Λ CDM cosmology predicts the average mass accretion rate onto dark matter halos as a function of mass and redshift (e.g., Neistein & Dekel 2008; McBride et al. 2009; Fakhouri et al. 2010), as well as the shape of the dark matter halo potential wells (e.g., Hernquist 1990; Navarro et al. 1997). The product of the halo potential well depth with the gas mass accretion rate provides an estimate of the rate at which gravitational potential energy that can be radiated is “injected” into the halo. Figure 2 assumes an optimistically high efficiency factor $f_{\text{eff}} = 1$ (Equation (A3)). As we discuss in Appendix 2, this gravitational power estimate is not strictly an upper bound for the total amount of cooling achievable even in the case of pure accretion, without feedback from stars and AGNs, since the accretion streams are embedded in the cosmic ionizing background, which can transfer further energy to them. This contribution is, however, included in our simulations and circumstantial evidence suggests that it does not dominate. In fact, in our most realistic simulations with a consistent self-shielding approximation—shown in the right panel of Figure 2 and to be discussed below—the gravitational power upper bound is never systematically violated, and a simulation of our A1 halo with the ionizing background completely turned off yields a nearly equal cooling luminosity (within 30%) at $z = 3$ as our prescription 9 when excluding the multiphase star-forming regions in the same manner.

The more sophisticated analytic model of Dijkstra & Loeb (2009), shown by the dashed lines for their fiducial efficiency parameter $f_{\text{grav}} = 0.3$, includes a factor $f_{\text{cold}}(M)$ accounting for the decreasing fraction of cold gas in massive halos (e.g., Kereš et al. 2005, 2009b). This model therefore predicts lower cooling luminosities than the above upper bound, with a shallower mass dependence at large masses, which is in better agreement with the simulation data points. It is important to note here that the cooling luminosities shown in Figure 2 are “theoretical” or “intrinsic,” meaning that they include all the photons emitted within the virial radii of the halos. These luminosities will in general

be higher than the observationally inferred luminosities, which only include the emission above a certain surface brightness threshold determined by the observation. Furthermore, a certain fraction of the emitted photons are in practice absorbed by the intervening intergalactic medium (IGM). The Dijkstra & Loeb (2009) data points plotted here (which were computed for a 50% IGM transmission factor by these authors) have been multiplied by a factor of two for a fair comparison with our simulation data points (which assume 100% transmission). In future work, we will quantify how the predicted theoretical luminosities translate into observational ones.

There are two main points to take away from the left panel of Figure 2. First, the predicted Ly α cooling luminosity is extremely sensitive to the treatment of the dense gas, with the range exceeding three orders of magnitude at $M_h \sim 10^{10} M_{\odot}$. Second, some prescriptions actually lead to unphysical results, as comparison with the analytic upper bound indicates that they emit more Ly α power than is available from the release of gravitational energy by orders of magnitude (in Appendix A, we show that photoionization from the cosmic background cannot physically account for such large luminosities either). This is the case, in particular, for a standard simulation with an optically thin ionizing balance and a multiphase star formation model in which all the gas particles are naively summed over (red \times 's). Unsurprisingly, the cases of no Ly α emission from self-shielded gas (magenta squares) and of CIE with $T_{\text{CIE}} = 10,000$ K (blue circles) yield nearly identical luminosities, since the cooling curve is already strongly suppressed at this temperature. The case of CIE with $T_{\text{CIE}} = 15,000$ (cyan diamonds) yields more optimistic Ly α luminosities, although these push the limit of the power that can be provided by the release of gravitational potential energy alone (the dashed line in Figure 2) at low masses. Prescription number 5 (green +’s), which uses a simulation in which the multiphase star formation model was turned off, but with a uniform ionizing background penetrating deep into the dense gas, illustrates that artificial photoheating alone can boost the cooling luminosities by orders of magnitude.

The critical question is therefore: what is the correct Ly α cooling luminosity of the cold streams? To address this, we develop a simple approximation to the self-consistent evolution of the gas properties with self-shielding. When self-shielding is properly modeled, it will also be possible to show that naively integrating over multiphase particles alone can also artificially boost the cooling luminosity by a large factor.

3.2.2. On-the-Fly Self-shielding

In Figure 3, we plot the hydrogen neutral fraction $x_{\text{H I}} \equiv n_{\text{H I}}/n_{\text{H}}$ as a function of total proper hydrogen number density n_{H} for the gas around the A1 halo at $z = 3$. The panel on




Figure 3. Hydrogen neutral fraction vs. total hydrogen number density in a cube of side length 1 comoving Mpc h^{-1} centered on the A1 system at $z = 3$. Left: values for a standard simulation with a uniform UVB. Right: same quantity after post-processing with the ionizing radiative transfer method. The ionizing radiative transfer shows that the main effect of self-shielding is to create a vertical “plume” of neutral gas above a density $n_H \sim 0.01 \text{ cm}^{-3}$. Figures 4 and 5 illustrate the effects of self-shielding on the dynamics and thermal state of the gas, when it is approximated during the course of the hydrodynamical simulation. The two-dimensional histograms are in arbitrary (but matching) logarithmic units and weighted by n_H^2 to emphasize the regions where the two-body emission processes are most efficient. Gas from multiphase, star-forming particles ($n_H > 0.13 \text{ cm}^{-3}$) is excluded.

the left shows this distribution for the standard simulation with a uniform ionizing background. The panel on the right shows exactly the same quantity after the post-processing ionizing radiative transfer. The effect of self-shielding is clear. Roughly, it generates a vertical “plume” above $n_H \sim 0.01 \text{ cm}^{-3}$, indicating the fact that the gas becomes mostly neutral above this density. This motivates our approximation to the self-consistent evolution of self-shielded gas in the hydrodynamical simulations. Namely, we rerun simulations with exactly the same initial conditions and other physical parameters, but set the ionizing background to zero in regions where the density exceeds the fiducial threshold $n_H = 0.01 \text{ cm}^{-3}$ (for an alternative scheme in which the UVB is turned off where the gas is optically thick to ionizing photons on a scale $\sim 0.1\text{--}1 \text{ kpc}$, see Sommer-Larsen 2006; Laursen & Sommer-Larsen 2007; Laursen et al. 2009b). By turning the ionizing background off in dense regions on the fly, their thermal and dynamical properties are consistently evolved with the modified cooling and heating functions, and the corresponding dynamical response. The “CDB” simulation analyzed by Goerdt et al. (2010) employed an analogous scheme, but with a density threshold $10\times$ higher, $n_H = 0.1 \text{ cm}^{-3}$; in Section 5, we argue that this difference likely explains much of the discrepancy with our results. Some uncertainty is introduced by our choice of a fixed density threshold for self-shielding, and in the future it would be useful to improve the methodology by performing proper ionizing radiative transfer on the fly, which our codes do not allow us to do at present. There are however reasons to believe that this choice is a good one, which we outline next.

Figure 4 summarizes the hydrodynamical properties (total gas distribution, neutral gas distribution, and temperature structure) for the A1 system at $z = 3$ for the different treatments of self-shielding: standard uniform ionizing background, post-processing ionizing radiative transfer, and the on-the-fly self-shielding approximation. When the ionizing radiative transfer is taken into account, the neutral hydrogen column density of the

cold streams can be greatly enhanced, especially in the higher density regions close to the central and satellite galaxies, indicating the fact that they self-shield (some of the cold gas at larger radii however remains optically thin). The ionization structure obtained with the on-the-fly self-shielding approximation is furthermore remarkably similar to the one obtained with the post-processing ray tracing scheme, supporting the validity of using our simple density criterion during the course of the hydrodynamical simulation. The simulation with on-the-fly self-shielding is the most accurate as it consistently captures the thermal evolution and dynamical response of the self-shielded gas. Figure 5 illustrates the effects of self-shielding on the temperature structure of the gas more explicitly: the self-shielded gas with $n > 0.01 \text{ cm}^{-3}$ is generally cooler (with $T \lesssim 10^4 \text{ K}$) when its evolution is consistently modeled. This simply results from the suppression of artificial photoheating and the enhancement of the cooling function in CIE. At these temperatures, the gas radiates very inefficiently in Ly α , which provides further evidence that the simple self-shielding density threshold is not introducing large errors: the Ly α cooling luminosity versus halo mass predicted using prescription 9 (discussed below) is quite close to what is obtained by effectively suppressing the Ly α emission from all the self-shielded gas, as in prescriptions 2 and 3 in which the self-shielded gas is identified using a ray tracing method and does not assume a particular density threshold. We have also run a simulation of the A1 halo with the ionizing background completely turned off, so that all the cooling in this case originates from gravitational energy and requires no self-shielding correction. The cooling luminosity for this simulation equals the one obtained with the simulation with on-the-fly self-shielding within $\sim 30\%$, when the star-forming regions are identically excised. We are therefore confident that our simple density threshold for self-shielding yields relatively accurate results.

The prescriptions for calculating the Ly α cooling luminosity from the simulations with on-the-fly self-shielding

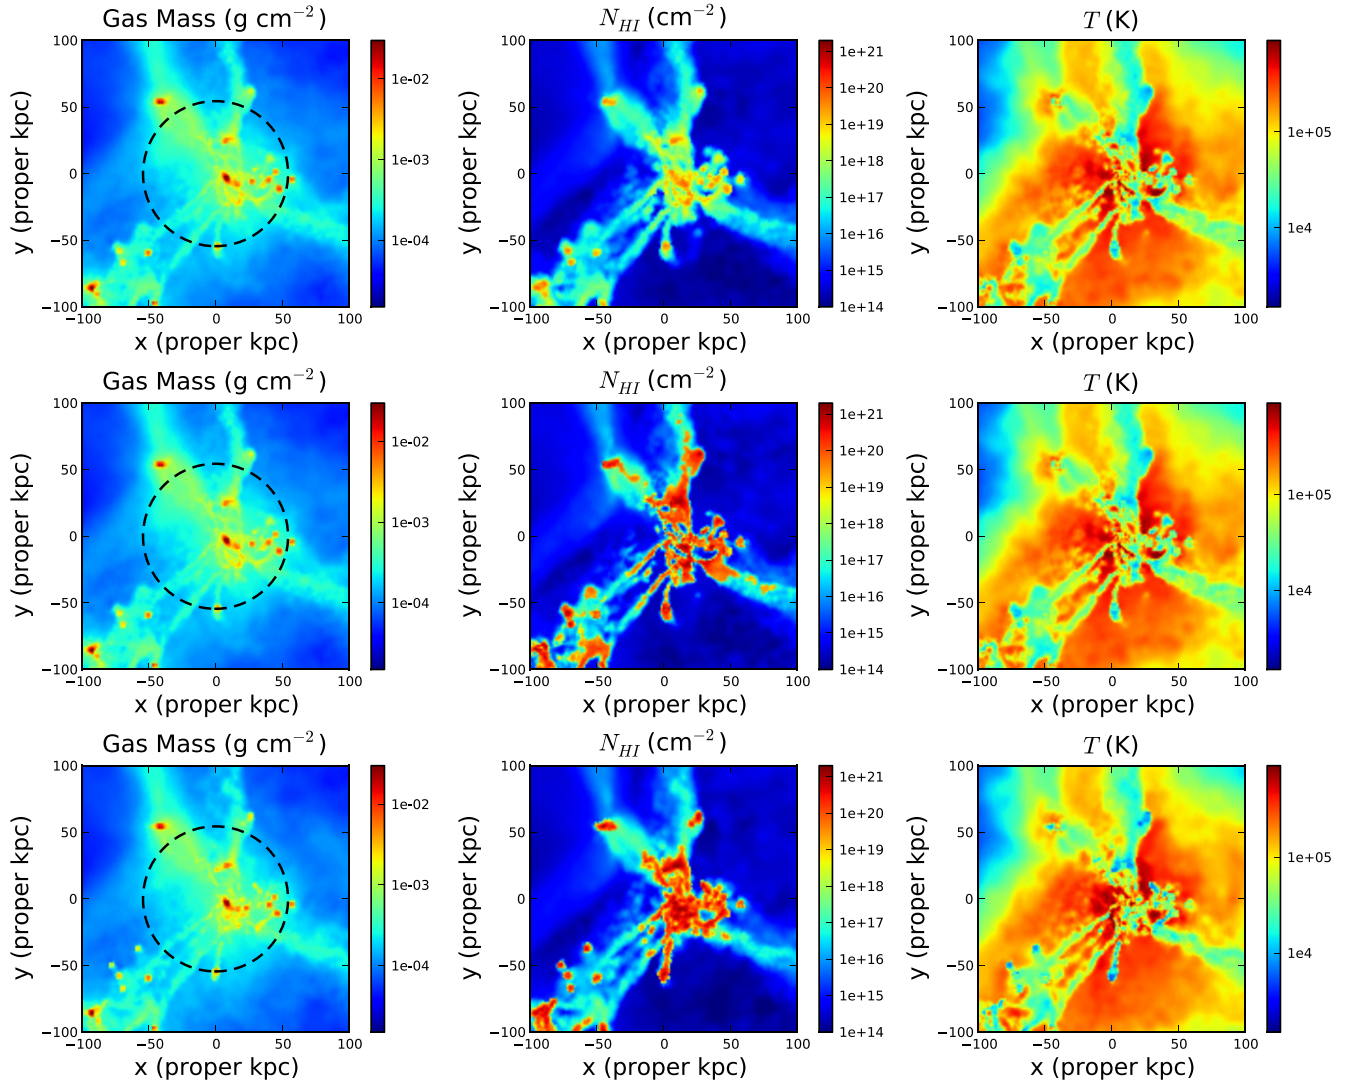


Figure 4. Hydrodynamical properties of the A1 system at $z = 3$ as a function of the self-shielding treatment. The left column shows the total projected gas mass, the central column shows the neutral hydrogen column density, and the right column shows the projected gas temperature. In all cases, the projected depth is 1 comoving Mpc h^{-1} . The projected temperature is weighted by density squared to emphasize the dense filaments and the virial radius of the halo is indicated by the dashed circles. Top: standard simulation output, with a uniform ionizing background. Middle: same, but post-processed with the ionizing radiative transfer scheme to identify the self-shielded regions (Section 3.2.1). Bottom: same initial conditions but with the on-the-fly self-shielding approximation, in which the ionizing background is turned off in regions with $n_{\text{H}} > 0.01 \text{ cm}^{-3}$ as the simulation proceeds to capture the effects on the thermal and dynamical evolution of the gas. Figure 5 illustrates the effects of self-shielding on the temperature structure of the gas more explicitly. The simulations shown include star formation.

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

approximation are labeled 6–9 in Table 3 and the corresponding results are shown in the right panel of Figure 2. For the last three (most consistent) cases, the filled symbols show the results for the A1 halo in our zoom-in simulation with $8\times$ better mass resolution. Prescription 6 (red \times 's), in which self-shielding is modeled but in which we naively sum over multiphase particles, demonstrates how the multiphase particles alone can produce artificially high cooling luminosities; these should therefore always be excluded, or treated separately. As an aside, comparison of prescriptions 1, 5, and 6 indicates that having the high density regions turn into multiphase particles limits the amount of artificial photoheating by effectively shielding the very dense gas. Three prescriptions (7, 8, and 9) correspond to physically plausible cases: one with the on-the-fly self-shielding approximation and star formation, but with multiphase star-forming particles excluded from the $\text{Ly}\alpha$ luminosity sum (9; blue circles); one also with the on-the-fly self-shielding approximation, but with

the star formation model turned off, summed over all the particles (7; magenta squares); and the intermediate case of summing only the particles with $n_{\text{H}} < 0.13 \text{ cm}^{-3}$ at which the gas would have become multiphase if the star formation model had been on (cyan diamonds; 8). All three are physically realistic in the sense that they are uncontaminated by either artificial photoheating or by the sub-resolution multiphase model. The only difference between the three cases is in how the gas with $n_{\text{H}} > 0.13 \text{ cm}^{-3}$, the density at which the multiphase model becomes active if on, is treated. Since the multiphase model was calibrated to match the observed Kennicutt–Schmidt relation (Springel & Hernquist 2003), this threshold density corresponds approximately to the density above which stars should start forming (although the exact value depends on some physical assumptions and may depend on redshift; e.g., Schaye 2004; Wolfe & Chen 2006).

In principle, the second approach (prescription 7) might seem more accurate since it captures the entire cooling. However, as is




Figure 5. Gas temperature vs. total hydrogen number density in a cube of side length 1 comoving Mpc h^{-1} centered on the A1 system at $z = 3$. Left: values for a standard simulation with a uniform UVB (top panel of Figure 4). These also apply for the simulation post-processed with the ionizing radiative transfer scheme, since it does not update the temperatures (middle panel of Figure 4). Right: same quantity, but for the simulation with the ionizing background turned off in regions with $n_{\text{H}} > 0.01 \text{ cm}^{-2}$ during the course of the hydrodynamical calculation as an approximation to the effects of self-shielding (bottom panel of Figure 4). The self-shielded gas is generally cooler (with $T \lesssim 10^4 \text{ K}$) when its evolution is consistently modeled as a result of the suppression of artificial photoheating and the enhancement of its cooling function. The two-dimensional histograms are in arbitrary (but matching) logarithmic units and weighted by n_{H}^2 to emphasize the regions where the two-body emission processes are most efficient. Gas from multiphase, star-forming particles ($n_{\text{H}} > 0.13 \text{ cm}^{-3}$) is excluded.

apparent in the radiative transfer results of Section 4, the extra cooling luminosity is concentrated around the accreting galaxy. It is unclear whether this central cooling emission would be observable in reality for at least two reasons. First, the density of the medium and the immediate proximity of the galaxy imply that locally produced Ly α photons could be efficiently destroyed by dust (the escape fraction of Ly α photons from Lyman break galaxies (LBGs), for example, covers the entire range $\sim 10^{-3}$ –1; e.g., Kornei et al. 2010). In itself, this is not necessarily an issue for this study in which we focus on a simplified dust-free problem, and would be a well-posed problem for follow-up studies in which dust would be included. Since the extra cooling luminosity occurs in ISM gas, it should however be accompanied by stellar emission which would most likely swamp it locally (see Section 5), and in that case is not really cooling luminosity from the cold streams. Second, turning off the multiphase model removes ISM pressurization, which can lead to catastrophic collapse of the gas-rich disks and deepen the potential wells, allowing extra energy release. Gravitational interactions with dark matter clumps might also artificially transfer energy to unstable gas disks. Since this prescription assumes that all the baryons are in the gaseous component, while in the central galaxies of actual galaxies in massive halos a large fraction would be locked in stars, it is likely an upper limit. Prescriptions 8 and 9 are more conservative as they exclude all cooling emission occurring at densities $n_{\text{H}} > 0.13 \text{ cm}^{-3}$. We expect these three cases to bracket the true cooling luminosity. A better understanding of the energy release at disk interfaces and within galaxies will likely be required to make more definite predictions and should be addressed in future work.

We conclude by noting that some caution is in order when interpreting the quantitative details of the cooling luminosity predictions for the halos at the low end of the mass range in Figure 2, since they contain relatively few SPH particles (~ 200 for $M_{\text{h}} = 10^{10} M_{\odot}$) and may not be well converged. Note,

however, that for the realistic prescriptions 7–9, the slope of the numerically predicted L_{α} – M_{h} relation agrees well with the analytic expectation based on energy conservation (dashed lines in the figure), in this regime where the cold mode dominates and where the scaling should apply.

4. Ly α RADIATIVE TRANSFER

Having described our hydrodynamical simulations and the modeling of Ly α photon production, we proceed to the Ly α radiative transport problem. We use a new three-dimensional Ly α radiative transfer code, αRT , described in more detail in Appendix C. To summarize, the fields defining the physical state of the gas from a simulation are interpolated onto a Cartesian grid placed around a halo of interest, as for the post-processing ionizing radiative transfer (Section 3.2). Monte Carlo Ly α photons are then seeded throughout the gridded volume, with a number proportional to the local Ly α emissivity (Section 3). The multiple resonant scatters of each Monte Carlo photon are simulated until escape. As the photons propagate, the Ly α image and corresponding spectrum in each pixel on the sky are constructed as seen by an observer on Earth, taking into account cosmological surface brightness dimming. We do not however model the effects of IGM filtering in this work (see Section 4.2), so that the results are more properly interpreted as the redshifted emission as it emerges from the galactic halos.

In typical astrophysical situations, the optical depth at the center of the Ly α resonance is very large, easily $\tau_0 > 10^6$ or more (Equation (C1)). As a result, a photon propagates only a short distance before being absorbed by and exciting another neutral hydrogen atom to its $2p$ state. Because of the high Einstein A coefficient for the $2p \rightarrow 1s$ transition, another Ly α photon is quickly re-emitted ($A_{21}^{-1} \sim 10^{-9} \text{ s}$). Since the re-emission is in general in a different direction than the incident photon, the propagation can be effectively pictured as a single

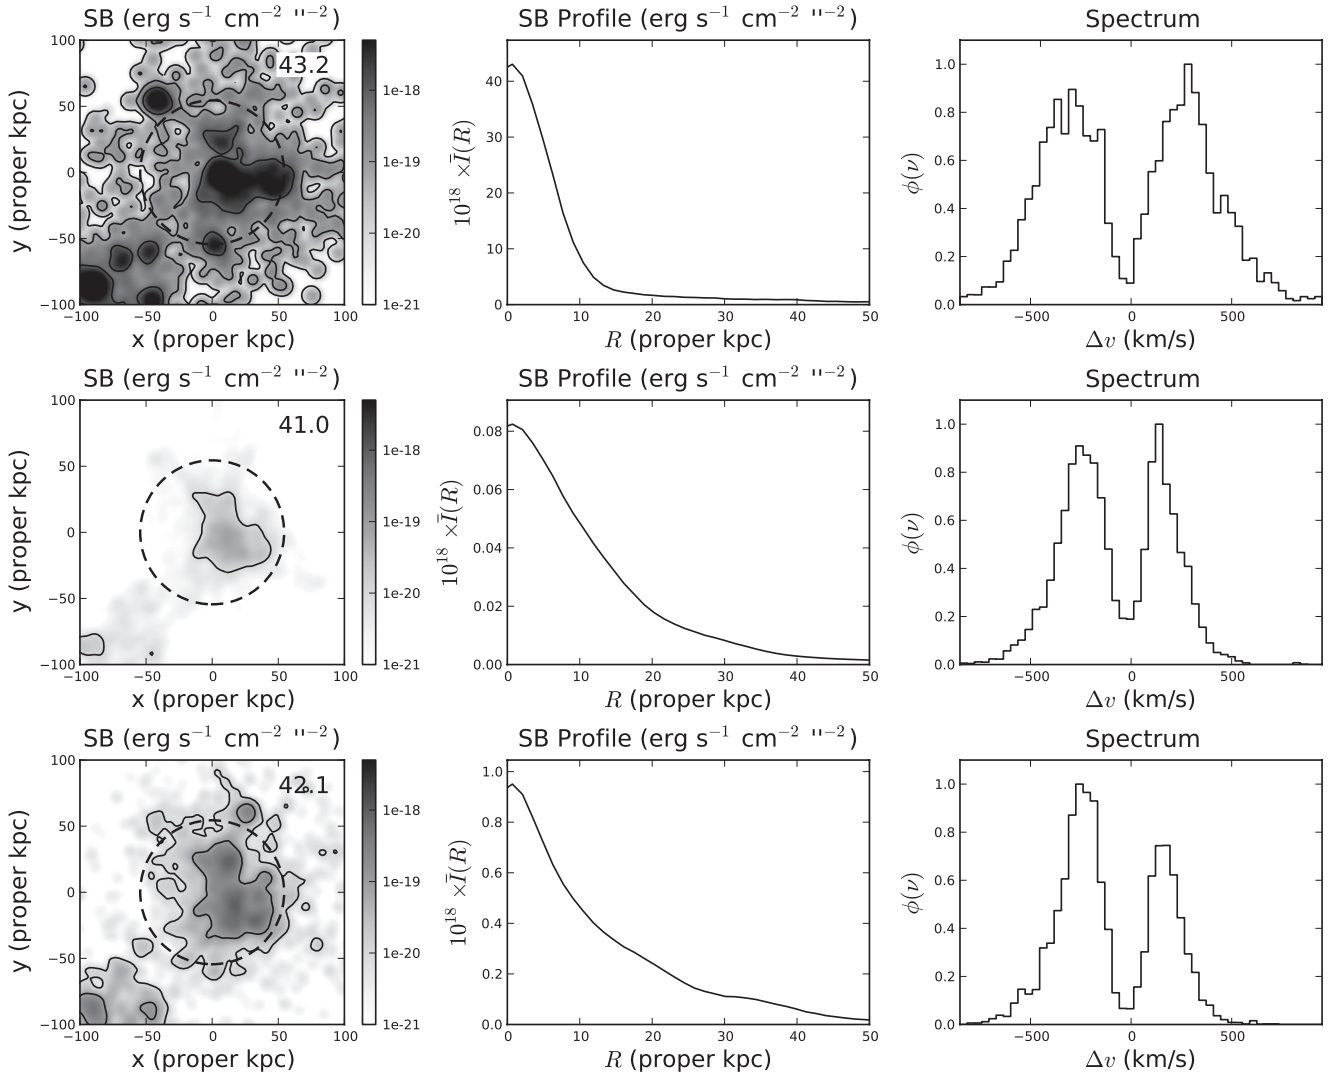


Figure 6. Dependence of the Ly α properties of the A1 system at $z = 3$ for prescriptions 1, 3, and 4 for the state of the dense gas. The left column shows the observed surface brightness distribution (smoothed with a Gaussian of FWHM = $1''$), the middle column shows the corresponding circularly averaged surface brightness profile, and the right column shows the line spectrum integrated within the virial radius of the halo, indicated by the dashed circles. The surface brightness contours correspond to 10^{-18} , 10^{-19} , and 10^{-20} $\text{erg s}^{-1} \text{cm}^{-2} \text{arcsec}^{-2}$. The logarithm of the total apparent line luminosity (in erg s^{-1}) within the virial radius is indicated at the top right corner of each surface brightness panel. Top: (prescription 1) standard hydrodynamical simulation, with uniform ionizing background and a multiphase model for star formation, naively integrated over all the particles. Middle: (prescription 3) same simulation, but post-processed with ionizing radiative transfer (Section 3.2.1). The self-shielded regions are regions assumed to be in CIE with the temperature set to $T_{\text{CIE}} = 10,000$ K. Bottom: (prescription 4) Same but with self-shielded regions assumed to be in CIE with the temperature set to $T_{\text{CIE}} = 15,000$ K.

photon being scattered and undergoing a random walk. Even if the scattering is coherent in the frame of the scattering atom, the motion of the atom combined with the redirection of the photon in general results in a small shift in the frequency as viewed by an external observer. This results in a random walk in frequency space as well. The frequency-space random walk plays a crucial role in shaping the emergent line profile as a photon tends to escape the medium when it finds itself sufficiently far from the line center that the optical depth it sees is reduced to ~ 1 (unless the medium is so optically thick that the photon spatially diffuses out first). The transport of Ly α radiation is thus very different than ordinary lines and a simplified optically thin treatment would lead to fundamental errors both in the theoretical predictions and in interpreting observations.

In this section, we present our basic results on Ly α cooling emission radiative transfer. As the focus of this work is to understand the theoretical and numerical uncertainties, and the

relevant physical effects, we limit ourselves to examining a particular example, our A1 halo at $z = 3$.

4.1. Emission Physics

In Section 3, we demonstrated how the predicted Ly α cooling luminosity depends on the treatment of self-shielding and of sub-resolution physics. We illustrate how the different possible assumptions manifest themselves in other observational properties by performing Ly α radiative transfer calculations for some of the prescriptions studied above (Table 3). The results are shown in Figures 6 (prescriptions 1, 3, and 4) and 7 (prescriptions 5, 7, and 8). These fiducial calculations use $N_{\text{ph}} = 10,000$ Monte Carlo photons. For the radiative transfer plots, we subtracted the peculiar motion of the central galaxy with respect to the simulation box (40 km s^{-1} toward the observer), so that $\Delta v = 0$ corresponds to the galaxy rest frame. As we had found from our study of a sample of halos in Section 3.2, incorrectly

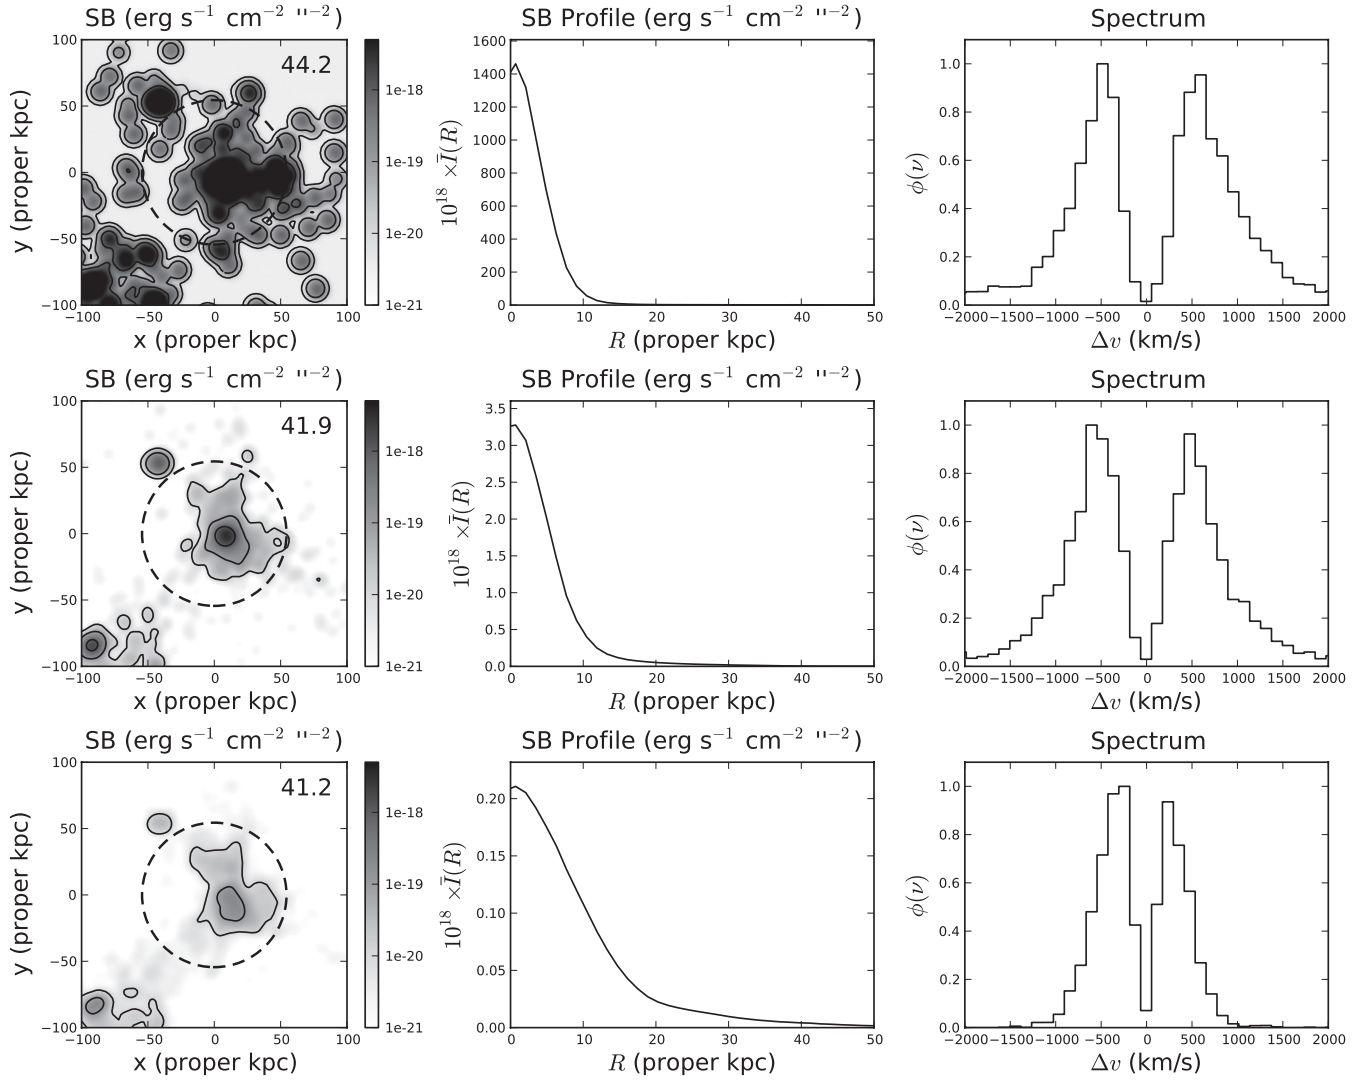


Figure 7. Same as Figure 6, but for prescriptions 5, 7, and 9. Top: (prescription 5) Simulation with a uniform ionizing background but no star formation, integrated over all the particles. Middle: (prescription 7) simulation with the on-the-fly self-shielding approximation and no star formation, integrated over all the particles. Bottom: (prescription 9) simulation with the on-the-fly self-shielding approximation and star formation, but with no Ly α luminosity from star-forming particles. Note the different velocity scale in comparison with Figure 6.

treating self-shielding (either by assuming a uniform ionizing background, or post-processing ionizing radiative transfer with default simulation temperatures) or including gas particles contaminated by the effective sub-resolution multiphase model can lead to order-of-magnitude overestimates of the cooling luminosity. However, it is not only the total luminosity of a system that can be incorrectly predicted, but also its morphology and line spectrum.

It is easy to understand how different prescriptions for seeding the Ly α photons can affect the observed morphology and spectrum. In terms of morphology, different prescriptions seed the photons not only in different amounts, but also in different places. For instance, including artificially luminous multiphase gas particles produces a disproportionately large Ly α luminosity concentrated in the densest central parts of the system, approximating a bright point source rather than spatially extended emission from the cold streams. Seeding the photons in different places also has implications for the emergent spectrum: photons seeded deep within dense self-shielded regions have a much larger H I column density to traverse before escape, and so result in more widely separated double peak profiles.

Some specific points are worth noting. As indicated by the surface brightness maps and the circularly averaged surface brightness profiles, the resulting objects are spatially extended by tens of proper kpc, comparable to the virial radius of the halo (see also Dijkstra & Loeb 2009) and to many of the observed Ly α blobs (e.g., Matsuda et al. 2004; Saito et al. 2008). This spatial extent results from a combination of some of the cooling radiation being emitted in the accretion streams far from the galaxy, and of spatial diffusion owing to resonant scattering. However, the spatial extent is strongly dependent on the surface brightness threshold and consequently on the luminosity prescription, so that even an intrinsically diffuse source could appear relatively compact in observations. For instance, for our optimistic prescription 7 for the A1 halo at $z = 3$ shown in the second row of Figure 7, only the central few kpc would stick out above the surface brightness threshold of $\approx 2 \times 10^{18} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ of the narrowband images of Matsuda et al. (2004). In this case, the more diffuse cooling halo would be completely missed, but would show up over a larger area in deep long-slit spectra sensitive to lower surface brightnesses (e.g., Rauch et al. 2008). The predicted line spectra

are double peaked, a common characteristic of Ly α radiative transfer reflecting the fact that the photons can escape the medium either on the blue side or on the red side of the line center, where the opacity is typically too extreme (e.g., Neufeld 1990; Zheng & Miralda-Escudé 2002a; Dijkstra et al. 2006; Verhamme et al. 2006). In most cases (but not all, reflecting the effects of the complex geometry in the different prescriptions), the blue peak is slightly more pronounced than the red peak. This is a signature of systematic infall in the problem at hand, in which the velocity gradients tend to smear the line opacity on the red side of the Ly α line, making it easier for the photons to escape on the blue side. Almost all the star formation powered Ly α emitters (Shapley et al. 2003; Steidel et al. 2010), as well as many Ly α blobs and fainter analogs (Matsuda et al. 2006; Saito et al. 2008; Rauch et al. 2008), instead show dominant red peaks indicative of outflows. We do not see this phenomenon here simply because we have not modeled galactic winds in our simulations to simplify the physical problem. Unlike the thin accretion streams (with small covering factor) that produce only slightly stronger blue peaks, outflows (with order unity covering factor) are expected to boost the red peak more drastically (e.g., Dijkstra et al. 2006; Verhamme et al. 2006). Intergalactic absorption, neglected here, would also tend to preferentially suppress the blue peak. In future work, we will include outflows and IGM filtering, which should provide a better match to the observational data.

In Section 3, we had identified three physically plausible prescriptions for calculating the Ly α cooling luminosity (7, 8, and 9); the radiative transfer results for 7 and 9 are shown in the second and third rows of Figure 7. The surface brightness maps now make it clear that the extra luminosity obtained using prescription 7, i.e., when using the on-the-fly self-shielding approximation but turning off the star formation model and summing over all the particles, is concentrated in the densest central parts of the system. This also manifests itself in the predicted line spectrum, which shows more widely separated peaks as a result of the higher column densities of the gas through which the bulk of the photons must propagate to escape, as discussed in the following section.

4.2. Radiative Transfer Physics

Since radiative transfer effects play a crucial role in shaping the observational properties, we pause to illustrate exactly how important each piece of physics is. To do so, we repeat the same fiducial calculations but sequentially turn off different physical effects. We only repeat the calculations for the plausible prescription 9, which suffices to illustrate the physics. The results are presented in Figure 8.

First, we keep resonant scatters but artificially set the gas velocities to zero. In this case, the most important change is in the line spectrum, which becomes a nearly symmetric double peak. This is easily understood in the context of the plane-parallel analytic solution derived by Neufeld (1990, for an adaptation to spherical geometry, see Dijkstra et al. 2006) for a monochromatic source in an extremely optically thick, static medium. Assuming that the source is located at the center of the slab and that τ_0 is the line center optical depth from the source to the surface, Neufeld (1990) showed that the emergent spectrum is a symmetric double peak profile, with each peak offset by

$$|\Delta v_p| \approx 191 \text{ km s}^{-1} \left(\frac{T}{10^4 \text{ K}} \right)^{1/6} \left(\frac{N_{\text{H I}}}{10^{20} \text{ cm}^{-2}} \right)^{1/3} \quad (5)$$

from the center. In dimensionless units in which $x \equiv (\nu - \nu_0)/\Delta\nu_D$ is the frequency offset in Doppler-broadening units, $|x_p| = 1.06(a\tau_0)^{1/3}$. For convenience, the corresponding offset in observed wavelength units can be obtained from

$$\Delta\lambda_{\text{obs}} \approx 8.1 \text{ Å} \left(\frac{1+z}{4} \right) \left(\frac{\Delta\nu}{500 \text{ km s}^{-1}} \right). \quad (6)$$

In particular, the emergent line profile and corresponding effective line width are in this case determined by the physics of Ly α radiative transfer and of self-shielding (which sets the H I column densities) and have little to do with the global properties of the host halo.

Next, we turn velocities back on but artificially turn off resonant scattering. This is equivalent to assuming that the Ly α line were an ordinary, optically thin line. In this case, there are two important changes. First, the morphology on the sky is slightly sharper since there is no longer a spatial diffusion effect arising from photon random walks. The spatial diffusion effect appears subtle here as the cooling emission is intrinsically diffuse. It is better illustrated in Figure 9, in which we have kept the same physical set-up but seeded the Ly α photons in proportion to the local star formation rate, which produces a much more intrinsically compact source. To facilitate visual comparison of the morphology with the pure cooling cases, we have normalized the star formation Ly α emission to the cooling luminosity of the halo. This halo is, however, forming stars at a rate $\sim 10 M_\odot \text{ yr}^{-1}$, which could result in up to $\sim 10^{43} \text{ erg s}^{-1}$ of star formation powered Ly α emission if the escape fraction were unity (Leitherer et al. 1999); in this case, it would dominate over the cooling luminosity of the halo. Second, the line spectrum is very different and has essentially lost its double peaked nature. Most importantly, the spectrum (particularly its width) is determined by completely different physics, since it is now directly representative of the velocity dispersion within the host halo through the Doppler effect.

An important implication of the radiative transfer effects on the observed Ly α line width concerns the way the masses of extended Ly α sources like the Ly α blobs are estimated. In fact, it is often assumed that the Ly α line width can be associated with random motion within the halo (e.g., Bower et al. 2004; Matsuda et al. 2006). As could be anticipated from more general Ly α radiative transfer studies (e.g., Zheng & Miralda-Escudé 2002a; Dijkstra et al. 2006; Verhamme et al. 2006), our results show that neglecting radiative transfer effects will tend to overestimate the mass, since they alone broaden the Ly α line even for completely static emitting media. However, the situation is in reality more complex since the Ly α lines that one actually observe on Earth are further filtered by the intervening IGM, which can significantly attenuate the Ly α line (e.g., Zheng et al. 2010a, 2010b; Laursen et al. 2010). At $z = 3$, for example, the diffuse IGM transmits only $\approx 67\%$ of photons immediately blueward of Ly α and this fraction decreases rapidly with increasing redshift (e.g., Faucher-Giguère et al. 2008d). Moreover, local matter overdensities and systematic infall around massive halos (e.g., Faucher-Giguère et al. 2008c) can amplify and shift the absorption, whereas galactic winds act to “let through” more radiation by redshifting it away from the Ly α line center (e.g., Santos 2004; Dijkstra et al. 2007; Dijkstra & Wyithe 2010). It is therefore in general difficult to relate the observed Ly α line alone to the properties of the halo producing it, and it is not even clear in any given case whether the observed line width overestimates or underestimates the halo velocity dispersion.

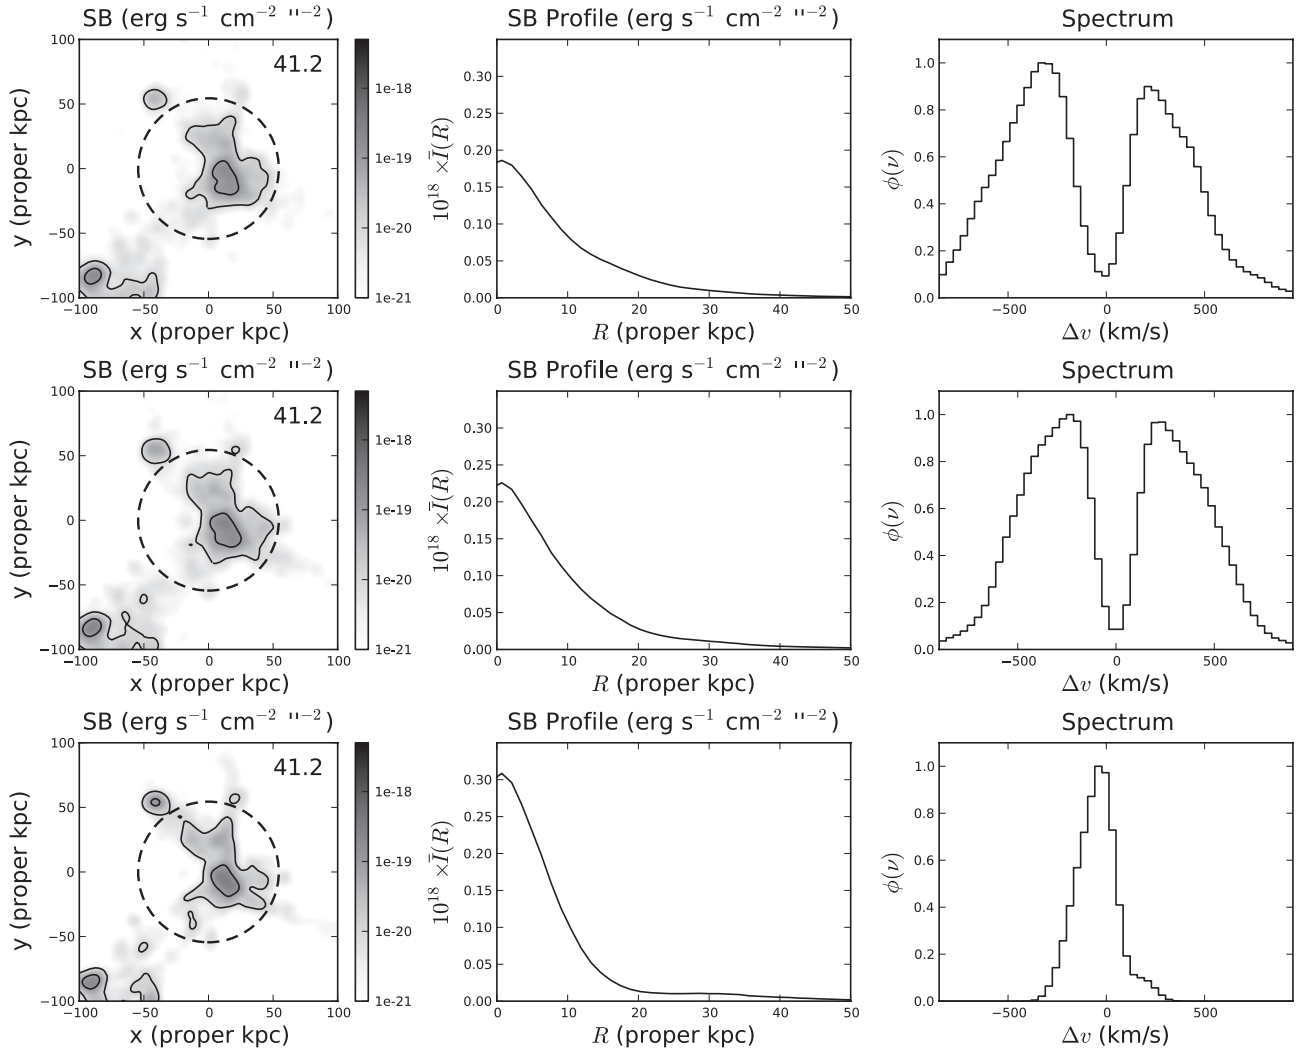


Figure 8. Illustration of the effects of radiative transfer on the morphology and spectrum of the Ly α cooling radiation, for prescription number 9 (on-the-fly self-shielding approximation with the multiphase star formation model, but with the multiphase particles excluded from the luminosity calculation) on the A1 system at $z = 3$. Top: fiducial calculation, with all radiative transfer effects (repeated from Figure 7 but with different axis scales). Middle: bulk velocities artificially set zero. The double peaked spectrum is now symmetric since the velocity flows no longer break the symmetry between the blue and the red peaks. The effective line width is set by the physics of self-shielding (which determines the H I column densities) and of Ly α radiative transfer. In particular, it has little to do with the global properties of the host halo. Bottom: bulk velocities back on, but resonant scatters artificially turned off. The morphology on the sky is now sharper since there is no longer spatial diffusion of the photons. More importantly, the spectrum has lost its double peaked nature and its width is set by completely different physics, being in this case representative of the velocity dispersion of the emitting gas.

More systematic studies quantifying the relation between the observed Ly α line properties and the host halo mass, extending the type of calculations presented in this work, are certainly warranted. We here simply caution against taking halo masses estimated from the Ly α line too literally. This remark applies even if the Ly α blobs are not predominantly powered by cooling radiation since radiative transfer effects will necessarily be at play in any Ly α source.

4.3. Convergence

Most of our radiative transfer calculations employed physical fields stored on a Cartesian grid with 256^3 cells in a volume of $(1 \text{ comoving Mpc } h^{-1})^3$. Since the Ly α photon trajectories are not tied to the grid points (see Appendix C), their increments were however finer than that this by a factor of five. Furthermore, the cell luminosities were always calculated directly from the SPH particles so that they captured the full gas clumping and were not subject to resolution degradation. To verify that our

radiative calculations are robust, we have repeated many of them with 512^3 grid points in the same volume instead. As shown in Figure 10 for prescription 9 applied to the A1 halo at $z = 3$, the calculations in fact appear well converged both in terms of Ly α morphology and spectrum.

5. DISCUSSION AND CONCLUSION

Motivated by observations of a variety of extended Ly α sources at high redshift and by recent studies suggesting that cold mode accretion could account for the majority of them (e.g., Dijkstra & Loeb 2009), we have used hydrodynamical simulations to predict the Ly α cooling emission of forming galaxies. We have introduced a new Ly α radiative transfer code, αRT , and for the first time applied such a code to the particular problem of cooling radiation in cosmological hydrodynamical simulations. In this work, which will serve as the foundation for follow-up studies of the astrophysical implications, we have quantified the theoretical and numerical uncertainties in

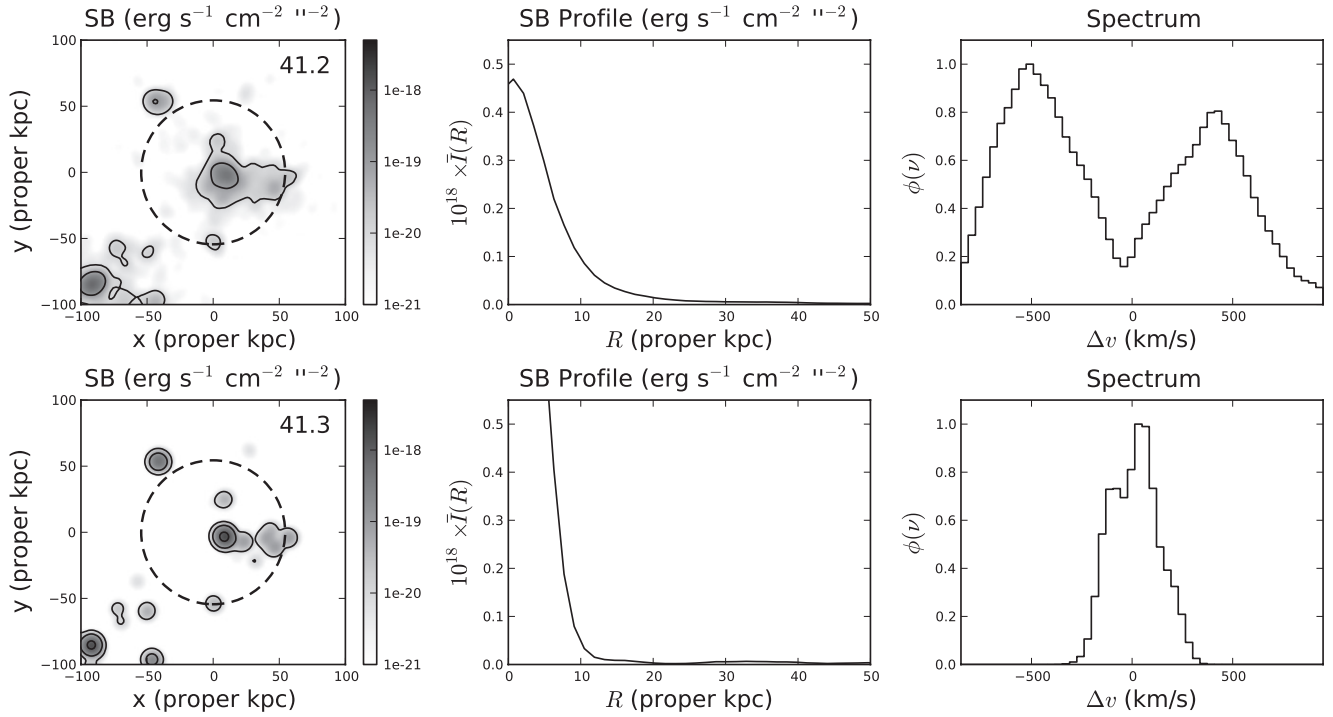


Figure 9. Cleaner illustration of the spatial diffusion effect owing to the Ly α resonant scatters. Top: same physical set up as for the cooling radiation prescription number 9 (Figure 8) but with all the Ly α photons emitted at the locations of the star-forming particles, in number proportional to the star formation rate. To facilitate the visual comparison of the effect on morphology, the total Ly α luminosity has been normalized to the cooling luminosity for prescription 9. Bottom: same calculation but with resonant scattering turned off to show the intrinsic compactness of the star-forming sources.




Figure 10. Illustration of the convergence of our radiative transfer results with the number of grid points along each dimension. Left: fiducial cooling luminosity calculation for prescription number 9, with 256³ radiative transfer grid points. Middle: same, but with 512³ radiative transfer grid points in the same volume. Right: corresponding line spectra within the virial radius. Both the morphology and the line spectrum appear well converged with the radiative transfer resolution.

predicting the properties of Ly α cooling radiation. To do so, we have considered the simplest physical problem of galaxy assembly embedded in the cosmic UV background but without feedback processes. We have shown that the Ly α cooling luminosities, morphologies, and spectra of the cold streams are strongly dependent on the treatment of the ionization and thermal state of the emitting gas, and that naive assumptions can produce artificially high cooling luminosities. Previous studies have usually relied on such assumptions, because most existing hydrodynamical simulations do not self-consistently follow the transfer of ionizing radiation (but see Sommer-Larsen 2006; Laursen & Sommer-Larsen 2007; Laursen et al. 2009b), whereas the cold streams in reality self-shield. The predicted cooling luminosity is so sensitive to the thermal state of the gas principally because of the exponential dependence of the collisional exci-

tation coefficient on temperature in the range $T \sim 10^4$ – 10^5 K characteristic of the cold streams. We have also demonstrated that sub-resolution physics models, when not carefully taken into account, can lead to large errors as well.

Having explicitly illustrated the range of uncertainties, we have made a systematic attempt to converge on the correct answer. Our strategy consisted of first post-processing our hydrodynamical simulations with ionizing radiative transfer to identify the self-shielded regions. By inspection, we found that there is a fairly well-defined total hydrogen density $n_H \sim 0.01$ cm⁻³ above which the gas self-shields, at least at the redshifts $z \sim 3$ characteristic of the existing observations. We have then rerun simulations with exactly the same initial conditions and other physical parameters, but with the ionizing background turned off in regions exceeding the fiducial density

threshold $n_H = 0.01 \text{ cm}^{-3}$. By turning the ionizing background off in dense regions on the fly, this dense gas is consistently evolved with modified heating and cooling functions, as well as the corresponding dynamical response. These simulations provide us with our best estimates for the actual Ly α cooling luminosity of the cold streams, in good agreement with energetic considerations.

For our Ly α radiative transfer calculations, we focused on a particular halo (A1) of total mass $M_h = 2.5 \times 10^{11} M_\odot$ at $z = 3$. We have shown that it is not only the integrated Ly α luminosity of the system that depends strongly on assumptions regarding the thermal state of the gas, but also its apparent morphology and spectrum. To isolate the role of different physical effects in shaping the observational properties, we have sequentially turned off separate pieces of physics. Both the Ly α resonant scatters and the bulk velocity structure of the system are critical in shaping the emergent spectrum. The resultant effective line width in general differs strongly from the optically thin expectation owing to these radiative transfer effects, and we therefore caution against measuring the masses of systems based on the Ly α line alone. The intrinsically extended nature of the cooling emission and the spatial diffusion owing to resonant scattering combine to produce objects that are spatially extended on the sky, but with an apparent size that depends sensitively on the observational surface brightness threshold. As a result, only the central few kpc of the A1 halo Ly α cooling emission at $z = 3$ sticks out above the surface brightness threshold $\sim 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ arcsec}^{-2}$ typically achieved to date (Figure 7). If the characteristic size of the Ly α emission scales as $R \propto M_h^{1/3}$, then the surface brightness should scale as $SB \propto L_\alpha/R^2 \propto L_\alpha/M_h^{2/3}$. For our consistent prescriptions 7–9, Figure 2 shows that the high-mass luminosity slope varies from about 3/4 to 1 in the most optimistic case, i.e., $SB \propto M_h^{-0.33 \dots -0.08}$, so that even the higher mass systems are unlikely to show up as sources with extents $\sim 100 \text{ kpc}$ in existing observations from cooling luminosity alone. Fainter sources, however, may well be detectable in deeper observations, such as the 100 hr long-slit spectrum reported by Rauch et al. (2008).

Our analysis of a sample of halos from a cosmological volume at $z \approx 3$ confirms that (in agreement with previous studies; Haiman et al. 2000; Fardal et al. 2001; Dijkstra & Loeb 2009; Goerdt et al. 2010) cooling emission alone can *in principle* produce luminosities $L_\alpha \sim 10^{43} \text{--} 10^{44} \text{ erg s}^{-1}$ sufficient to account for luminous Ly α blobs. This however requires quite optimistic assumptions under which most of the energy is released close to, or within, the accreting disks, at densities sufficient to form stars according to the observed Kennicutt–Schmidt relation. In those optimistic predictions based on prescription 7, star formation was artificially turned off (yielding purely gaseous disks), so that it is unlikely that this entire cooling emission can be realized in reality without being overwhelmed by stellar emission.⁹ Our findings in this respect are at odds with the recent simulations of Goerdt et al. (2010), who argue that pure cooling radiation can explain the observed giant Ly α blobs. Our investigation of the sources of error in the Ly α luminosity predicted from simulations (Section 3) allows us to understand the differences. First, we have demonstrated that the predicted Ly α luminosity is very sensitive to the treatment of self-shielded gas and that at $z \approx 3$, the characteristic density

above which gas self-shields from the UVB is $n_H \approx 0.01 \text{ cm}^{-3}$. Goerdt et al. (2010) did not explicitly perform ionizing radiative transfer and instead assumed a higher self-shielding threshold of $n_H = 0.1 \text{ cm}^{-3}$. This higher density threshold translates into an overestimate of the Ly α luminosity owing to artificial photoheating of gas in the density range $0.01 \leq n_H \leq 0.1 \text{ cm}^{-3}$. Examination of the luminosity-weighted $T - n_H$ histogram in their Figure 7 in fact indicates that the bulk of the Ly α luminosity they predict originates from this density regime (including in clumps that would correspond to satellite galaxies, even when the central galaxy is excised), whereas Ly α emission is strongly suppressed at these densities in our approximation to the consistent thermal evolution of self-shielded gas, owing to rapid cooling below 10^4 K (Figure 5). While there is some uncertainty in the precise value of the self-shielding threshold, our radiative transfer calculations show that it is clearly below 0.1 cm^{-3} at the redshifts under consideration (Figure 3), and the value $n_H \approx 0.01 \text{ cm}^{-3}$ is further supported by recent radiative transfer calculations by other groups (Nagamine et al. 2010; Aubert & Teyssier 2010). The cleanest contrast between our results and those of Goerdt et al. (2010) is provided by their predictions of the Ly α luminosity originating from the “streams” alone, in which they excluded all emission from within 20% of the virial radius. Those can be compared, at best, with our pessimistic prescriptions 8 and 9, in which only gas with $n_H > 0.13 \text{ cm}^{-3}$ (above the sub-resolution star formation threshold) is excluded; our predictions are in those cases a factor $\gtrsim 10$ lower than theirs. It would not be fair to compare our most optimistic prescription 7 with the Goerdt et al. (2010) results with the central 20% of the virial radius excluded, since most of the extra emission under this prescription occurs very close to, or inside, the central galaxy. Finally, the simulations of Goerdt et al. (2010) included supernova feedback, which provides an additional source of energy and could also contribute significantly to boosting the cooling luminosity, even in the absence of ionizing photons from associated local sources (e.g., Taniguchi & Shioya 2000; Taniguchi et al. 2001; Ohya et al. 2003; Mori et al. 2004).

It is further notable that many Ly α blobs show spectral signatures of outflows (dominant red peak), at odds with the expectation of a more prominent blue peak in the case of pure cooling of infalling gas.¹⁰ Nonetheless, the release of gravitational potential energy through cooling radiation at observationally interesting levels is an inevitable prediction of galaxy formation in Λ CDM and the growing body of theoretical work supporting the cold mode scenario implies that a significant fraction would come out in Ly α . Such sources are therefore poised to be routinely detected in future surveys, and likely account for at least a subpopulation of the more modest extended Ly α sources, such as the ones detected by Rauch et al. (2008). In fact, at least three of the Rauch et al. (2008) sources (#15, 36, and 37) show spectral signatures suggestive of infall (with a more prominent blue peak). Interestingly, *Galaxy Evolution Explorer* non-detections indicate that the bright Ly α blobs are much rarer at $z = 0.8$ than at $z \gtrsim 2$ (Keel et al. 2009). This redshift evolution is reminiscent of the gradual disappearance of the cold mode at low redshift predicted by theoretical studies (e.g., Birnboim & Dekel 2003; Dekel & Birnboim 2006; Kereš et al. 2005, 2009b). It is important to note that the cold mode could play an important role in the existence

⁹ In their simulations of LBGs at $z = 3.6$, Laursen et al. 2009a find that only $\sim 10\%$ of the total Ly α emission comes from cooling, although we show in Appendix A that this ratio should depend on halo mass and redshift.

¹⁰ There are exceptions, for example certain regions of the LAB2 blob at $z = 3.09$ (Wilman et al. 2005).

and properties of the Ly α blobs even if they are not energetically dominated by cooling emission. In fact, the presence of cold neutral gas in galactic halos enhances the conversion stellar or AGN power into Ly α photons, and scattering off such gas may be necessary to produce the morphological and spectral properties observed.

Although we have made important progress in accurately predicting the Ly α properties of cold accretion, much work remains to be done. For instance, we have focused on the simplest physical problem of accretion embedded in the cosmic UV background and neglected feedback processes. In the actual universe, galactic winds are observed to be ubiquitous at high redshift (e.g., Shapley et al. 2003; Steidel et al. 2010) and simulations suggest that they are in fact needed to reproduce the observed baryonic mass function of galaxies (e.g., Kereš et al. 2009a; Oppenheimer et al. 2010). Furthermore, AGNs can inject large amounts of energy in the surrounding gas (Springel et al. 2005; Hopkins & Hernquist 2006; DeBuhr et al. 2010), which could radiate even after the nucleus has shut off. Such processes will modify the kinematic and thermal properties of the circumgalactic medium, and therefore its Ly α signatures. Moreover, they provide additional sources of energy that could enhance the total emission. The ionizing radiation produced by embedded star formation or AGN (e.g., Kollmeier et al. 2010), the effects of metals on the cooling function (e.g., Cantalupo 2010), and the destruction of Ly α photons by dust (as shown by Laursen et al. 2009b) are also likely to be important to reproduce the observed sources, but have not been modeled here for simplicity. Since cold accretion likely fuels star formation in halo centers, at rates similar to the cold gas accretion rates (e.g., Kereš et al. 2005, 2009b; C.-A. Faucher-Giguère et al. 2010, in preparation), gas accretion and star formation are intrinsically linked and their roles in Ly α emission are difficult to decouple. Follow-up studies will build upon the technical foundation established in this work and include some of these effects. Ultimately, cosmological statistics like the luminosity and correlation functions of the Ly α sources will be predicted. It will also be interesting to consider complementary observational probes of the cold streams, including their Ly α polarization (e.g., Rybicki & Loeb 1999; Dijkstra & Loeb 2008), their absorption signatures, and other emission lines (including from metals).

On a more basic level, our understanding of the physics of the cold mode is incomplete. In particular, it is not exactly clear how the cold streams release their gravitational energy. At least some simulations indicate that they maintain a roughly constant velocity, rather than accelerate, as they fall into galactic halos (Kereš et al. 2005). This implies that the streams would continuously radiate the gravitational work done on them, perhaps by undergoing a series of weak shocks. However, these weak shocks have not been explicitly identified in existing simulations. Birnboim & Dekel (2003) instead envision a scenario in which the cold streams free fall into the halos and release little energy before hitting the central disk in a strong shock, which may be partially supported by our results which suggest that a large fraction of the energy could be released near the halo center. Infalling cold gas might also release some energy through interactions with halo sub-structure and with the lower density, hot halo gas. These interactions could involve hydrodynamic instabilities below the resolution of current generation cosmological hydrodynamic simulations, both SPH and AMR. While we have taken the pragmatic point of view of predicting the observational signatures implied by our current galaxy formation models, it is conceivable that the

existing simulations are subject to numerical limitations. Work is underway to elucidate some of these hydrodynamical issues using the new shock-capturing, moving mesh code AREPO (Springel 2010). When the capability becomes available, it would also be useful to perform truly self-consistent ionizing radiative transfer on the fly in order to definitively capture the thermal evolution of the gas.

We are grateful to Volker Springel for making the expanded GADGET code available to us, and to Yuval Birnboim, Juna Kollmeier, Adam Lidz, Avi Loeb, Matt McQuinn, Eliot Quataert, Chuck Steidel, and David Weinberg for useful discussions. We also thank the referee, Jesper Sommer-Larsen, for a detailed and constructive review. C.-A.F.-G. is supported by a fellowship from the Miller Institute for Basic Research in Science, and received further support from the Harvard Merit Fellowship and FQRNT during the course of this work. D.K. is supported by NASA through Hubble Fellowship grant HST-HF-51276.01-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under contract NAS 5-26555. Partial funding was also provided by NSF grants ACI 96-19019, AST 00-71019, AST 02-06299, AST 03-07690, AST 05-06556, AST 05-06556, AST 09-07969, and PHY 08-55425, and NASA ATP grants NAG5-12140, NAG5-13292, NAG5-13381, and NNG-05GJ40G. Further support was provided by the David and Lucile Packard Foundation, the Alfred P. Sloan Foundation, the John D. and Catherine T. MacArthur Foundation, and Harvard University funds. The computations in this paper were run on the Odyssey cluster supported by the FAS Sciences Division Research Computing Group at Harvard University.

APPENDIX A

ANALYTIC CONSIDERATIONS

In this section, we outline analytic arguments that motivate the expectation of significant Ly α emission from galactic gas accretion, at a level comparable to observed Ly α blobs. Similar arguments were made by Haiman et al. (2000), Dijkstra & Loeb (2009), and Goerdt et al. (2010); their essence is summarized here as a basis to understand our numerical results. We also provide an analytic estimate of the amount of power that could be contributed by photoionization from the cosmic background.

We consider gas accreting from the IGM to the bottom of a dark matter halo potential well and ask how much energy is available to be radiated. Averaged over the accretion timescale, the gravitational power available can be expressed as $\langle \dot{E}_{\text{grav}} \rangle \approx f_{\text{eff}} \dot{M}_{\text{gas}} |\Delta \Phi_{\text{grav}}(r_{\text{min}})|$, where \dot{M}_{gas} is the gas mass accretion rate, $\Delta \Phi_{\text{grav}}(r_{\text{min}})$ is the potential difference from the IGM to the radius r_{min} from the bottom of the potential well at which the gas is assumed to settle, and $f_{\text{eff}} < 1$ is an efficiency factor. The efficiency factor accounts the fraction of the gravitational energy that remains in bulk kinetic or thermal form. In addition, only a certain fraction of the energy that is radiated comes out in the Ly α line.

Using the fitting formula derived by Neistein & Dekel (2008) for the average halo mass accretion rate, we can write (for the *WMAP5* cosmology)

$$\dot{M}_{\text{gas}} \approx 210 M_{\odot} \text{ yr}^{-1} \left(\frac{M}{10^{12} M_{\odot}} \right)^{1.4} \left(\frac{1+z}{4} \right)^{2.5} \left(\frac{f_{\text{gas}}}{0.165} \right), \quad (\text{A1})$$

where f_{gas} is the fraction of the accreted mass that is gaseous (see also McBride et al. 2009; Fakhouri et al. 2010). For a Navarro et al. (1997) halo profile,

$$|\Delta\Phi_{\text{grav}}(r_{\text{min}})| \approx \frac{GM_{200c}}{r_{\text{min}}} \frac{\ln[1 + (r_{\text{min}}/r_{200c})c]}{\ln(1+c) - c/(1+c)} \approx \frac{GM_{200c}}{r_{200c}} \frac{c}{\ln(1+c) - c/(1+c)}, \quad (\text{A2})$$

where M_{200c} is the halo mass defined as that exceeding 200 times the critical density, r_{200c} is the corresponding radius, and c is the halo concentration parameter. In the last step, we have assumed $cr_{\text{min}}/r_{200c} \ll 1$ to simplify the expression. In general, $M_{200c} \sim M$ ($\equiv M_{\text{FoF}} \approx M_{180b}$; see Section 2.2), but since r_{180b} may exceed r_{200c} by as much of 75% on cluster scales (White 2002), M_{200c} can be lower than M by a factor of as much as 5 at $z \sim 0$. For the present crude estimate, we will ignore the distinction between M and M_{200c} , which is much smaller at the high redshifts of interest. Combining Equations (A1) and (A2), we find

$$\langle \dot{E}_{\text{grav}} \rangle \approx 3.8 \times 10^{43} \text{ erg s}^{-1} f_{\text{eff}} \left(\frac{M}{10^{12} M_{\odot}} \right)^{1.8} \left(\frac{1+z}{4} \right)^{3.5} \times \left(\frac{f_{\text{gas}}}{0.165} \right) \left(\frac{c/[\ln(1+c) - c/(1+c)]}{5.2} \right). \quad (\text{A3})$$

This expression is valid at $z \gtrsim 1$ (where the cosmological constant can be neglected) and fiducially assumes $c = 5$, approximately the median concentration of all massive halos at $z \sim 2-3$ (e.g., Zhao et al. 2003).

As an estimate of the sensitivity of this analytic prediction to the shape of the dark matter halos, the calculation can be repeated for a Hernquist (1990) profile,

$$\Phi_{\text{grav}}(r) = -\frac{GM_{200c}}{r+a}, \quad (\text{A4})$$

using the relation $a = (r_{200c}/c)\sqrt{2[\ln(1+c) - c/(1+c)]}$ (e.g., Springel et al. 2005). We then obtain the ratio

$$\frac{\langle \dot{E}_{\text{grav}} \rangle^{\text{Hernquist}}}{\langle \dot{E}_{\text{grav}} \rangle^{\text{NFW}}} = \sqrt{\frac{\ln(1+c) - c/(1+c)}{2}}. \quad (\text{A5})$$

For $c = 3, 5$, and 10 , this ratio takes the values $0.56, 0.69$, and 0.86 , respectively. Since these differences, at the tens of percent level, are subdominant compared to the other sources of uncertainty, we will only plot the prediction for the Navarro–Frenk–White (NFW) model in this work.

In the absence of additional sources heating (see below, where we consider the importance of the ionizing background), the actual Ly α cooling luminosity emergent from a given halo will in general be somewhat lower than predicted by Equation (A3), since some of the gravitational power will be converted to kinetic and thermal forms, rather than radiated, and the radiated fraction will not come out entirely in Ly α . For halos dominated by the cold mode, however, the f_{eff} is expected to be significant, perhaps up to $\sim 50\%$ if half of the gravitational energy is radiated and most of this radiation comes out in Ly α . At higher masses, the efficiency will be suppressed as the hot mode becomes more important and a higher fraction of the baryons are accreted in the form of stars.

The analytic model of Dijkstra & Loeb (2009) is more sophisticated than the simple energetic argument above, as it accounts

for the fact that the fraction of cold gas in halos diminishes with increasing mass. Since only the cold gas radiates efficiently in Ly α , this yields a shallower slope for the $L_{\alpha}-M_h$ relation at high masses, which is in fact in better agreement with the simulation results (cf. Figure 2). Equation (A3) is thus really an upper bound that we do not expect to be saturated in this regime.

In addition to gravitational potential energy, there is an additional source of power that can be converted into Ly α photons even in the case of pure accretion, without feedback from stars or AGNs: the cosmic ionizing background. An upper bound for how much power the ionizing background can contribute to Ly α cooling within a dark matter halo can be estimated as the total inward flux of energy from ionizing photons across a sphere of a virial radius. A more realistic estimate is obtained by multiplying this quantity by a factor f_{cov} accounting for the fact that only a small fraction of these ionizing photons are actually absorbed within the virial shell:

$$\dot{E}_{\text{ion}} = 4\pi r_{\text{vir}}^2 F_{\text{ion}} f_{\text{cov}}, \quad (\text{A6})$$

where $F_{\text{ion}} = \pi \int_{\nu_{\text{H I}}}^{\infty} d\nu J_{\nu}$ and J_{ν} is the specific intensity of the ionizing background, assumed homogeneous and isotropic, which is a good assumption just above the hydrogen ionization edge at $\nu_{\text{H I}}$ (e.g., Faucher-Giguère et al. 2009). Since the background spectrum has a strong absorption edge at $4\nu_{\text{H I}}$ owing to intergalactic He II absorption, we can optimistically take $J_{\nu} \approx J_{\nu_{\text{H I}}}$ for $\nu \in [\nu_{\text{H I}}, 4\nu_{\text{H I}}]$, and $J_{\nu} = 0$ beyond $4\nu_{\text{H I}}$. Then

$$\dot{E}_{\text{ion}} \approx 12\pi^2 r_{\text{vir}}^2 \nu_{\text{H I}} J_{\nu_{\text{H I}}} f_{\text{cov}}. \quad (\text{A7})$$

For the hydrogen photoionization rate $\Gamma = 0.6 \times 10^{-12} \text{ s}^{-1}$ measured from the Ly α forest at $z = 3$ (Faucher-Giguère et al. 2008b), this spectrum implies $J_{\nu_{\text{H I}}} = 1.5 \times 10^{-22} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}$. Expressing in terms of halo mass, this gives

$$\dot{E}_{\text{ion}} \approx 1.9 \times 10^{41} \text{ erg s}^{-1} \left(\frac{M_h}{10^{12} M_{\odot}} \right)^{2/3} \left(\frac{4}{1+z} \right)^2 \times \left(\frac{J_{\nu_{\text{H I}}}}{1.5 \times 10^{-22} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}} \right) \left(\frac{f_{\text{cov}}}{0.05} \right). \quad (\text{A8})$$

Comparison of Equations (A3) and (A8) suggests that photoionization from the cosmic background can account for only about a percent of the maximum power available from the release of gravitational potential energy at the fiducial halo mass $M_h = 10^{12} M_{\odot}$. Owing to the different mass dependences, and because our numerical results indicate that the gravitational power upper bound can be far from saturated in some plausible cases (Section 3.2), we however cannot exclude that the ionizing background contributes a significant fraction of the Ly α cooling emission in some regimes. This contribution is included in our simulations, though, and while it is difficult to disentangle from gravitational energy, circumstantial evidence suggests that it does not dominate. Indeed, in our most realistic simulations with a consistent self-shielding approximation the gravitational power upper bound is never systematically violated (see the right panel of Figure 2), and a simulation of our A1 halo with the ionizing background completely turned off yields a nearly equal cooling luminosity (within 30%) at $z = 3$ as our prescription 9 when excluding the multiphase star-forming regions in the same manner.

Although we intentionally ignore Ly α photons produced by star formation throughout most of this work, it is interesting to

analytically estimate the photoionization stellar Ly α emission relative to pure cooling under simplified assumptions. Numerical simulations indicate that, in the absence of strong feedback, the star formation rate in high-redshift halos closely follows the cold gas accretion rate (e.g., Kereš et al. 2005, 2009b; C.-A. Faucher-Giguère 2010, in preparation), i.e., $\text{SFR} \approx \dot{M}_{\text{cold}}$, except for modest. Stellar Ly α emission results from the conversion of the ionizing radiation from young stars via absorption and recombination by the surrounding medium, and therefore scales with the star formation rate, with a proportionality constant that depends on the initial mass function. For standard assumptions, $L_{\alpha}^{\text{SF}} = 10^{42} \text{ erg s}^{-1} [\text{SFR}/(M_{\odot} \text{ yr}^{-1})]$ (e.g., Leitherer et al. 1999). If we assume that $\dot{M}_{\text{cold}} \approx \dot{M}_{\text{gas}}$,¹¹ then we can again make use of the average fitting formula in Equation (A1) to derive the average stellar Ly α luminosity as a function of halo mass and redshift:

$$\langle L_{\alpha}^{\text{SF}} \rangle \approx 2.1 \times 10^{44} \text{ erg s}^{-1} f_{\alpha, \text{esc}} \left(\frac{M}{10^{12} M_{\odot}} \right)^{1.4} \times \left(\frac{1+z}{4} \right)^{2.5} \left(\frac{f_{\text{gas}}}{0.165} \right), \quad (\text{A9})$$

where the $f_{\alpha, \text{esc}}$ factor quantifies the fraction of Ly α photons that avoid destruction by dust. We can then combine this with Equation (A3) to find the characteristic ratio of stellar to cooling Ly α :

$$\frac{\langle L_{\alpha}^{\text{SF}} \rangle}{\langle \dot{E}_{\text{grav}} \rangle} \approx 5.53 \left(\frac{f_{\alpha, \text{esc}}}{f_{\text{eff}}} \right) \left(\frac{M}{10^{12} M_{\odot}} \right)^{-0.4} \times \left(\frac{1+z}{4} \right)^{-1} \frac{5.2}{c/[\ln(1+c) - c/(1+c)]}. \quad (\text{A10})$$

The main point to take away is that it is a priori difficult to predict whether stellar Ly α emission should dominate over cooling Ly α , since the ratio depends on a number of uncertain parameters, including $f_{\alpha, \text{esc}}$ and f_{eff} . Furthermore, the ratio depends on both halo mass and redshift, so that theoretical predictions based on model galaxies at fixed M and z (e.g., Laursen et al. 2009a, 2009b) cannot be universally applied. The actual ratio of stellar ionization to cooling Ly α is likely to be somewhat suppressed relative to the simple scaling in Equation (A10) because galactic winds can remove a large fraction of the gas from the star-forming reservoir (e.g., Murray et al. 2005; Oppenheimer et al. 2010), therefore suppressing the Ly α emission from star formation ionization, without significantly affecting the accretion of cold material and hence the cooling Ly α .

APPENDIX B

IONIZING RADIATIVE TRANSFER

The opacity in the Ly α line (Equation (C1)) and the Ly α emissivity (Equations (1) and (2)) depends on the ionization state of the hydrogen atoms. The SPH code GADGET calculates the ionization of gas elements assuming photoionization equilibrium with a uniform UV background (e.g., Haardt & Madau 1996; Faucher-Giguère et al. 2009) (see Section 2.3). This optically thin treatment misses the effects of self-shielding

in dense regions (Section 2.3). To correct for this, we use a post-processing method to solve the radiative transfer problem on the gas distribution provided by GADGET to obtain more realistic ionization fractions and to identify the regions that self-shield from external radiation. The basic assumption is that the gas dynamics does not significantly depend on the radiative transfer effects missed in the SPH calculation. However, re-simulations are also used to approximate the dynamical and thermodynamical response to self-shielding (Section 3.2.2).

We use a ray tracing algorithm on the same grid as the Ly α radiative transfer (Appendix C). Specifically, we send a normal ray from each cell on the surface of the radiative transfer volume along the inward direction ($x+$, $x-$, $y+$, $y-$, $z+$, or $z-$) and iteratively solve for the equilibrium ionization structure. For this calculation, we take the gas temperatures provided by the SPH calculation and assume that all the helium is in the form of He III. These are good approximations for the gas outside of self-shielded regions and they therefore correctly capture the onset of self-shielding. In principle, the distribution of helium ionization affects the predicted Ly α luminosity through the abundance of free electrons. The exact ionization state of helium can however matter significantly only inside self-shielded gas, since the free electron number density is dominated by hydrogen elsewhere. Assuming that all the helium is in He III overestimates the number of free electrons and hence the recombination and collisional excitation rates. However, the overestimation is only significant for temperatures $T \lesssim 10^4 \text{ K}$, at which the self-shielded gas is found to contribute negligibly to the total Ly α emission (Section 3.2). During the re-simulations, turning off the ionizing background in dense gas results in most of the helium being in the form of He I in self-shielded gas, as should be the case since He I traces H I for realistic ionizing spectra. Errors associated with the helium ionization state are therefore expected to be small in all cases. The rays are assumed to originate from a diffuse ionizing background with a hydrogen photoionization rate $\Gamma_{\text{H I}}^{\text{bkg}}$. For simplicity, we only update the hydrogen ionization fractions and only keep track of the photoionization rate $\Gamma_{\text{H I}, i}$ along each ray, rather than the full ionizing spectrum. This is a fair approximation because of the small thickness of the self-shielding layer, in which the gas transitions from almost fully ionized to almost completely neutral. Indeed, the thickness of this layer is approximately equal to the mean free path of ionizing photons within it:

$$\Delta l_{\text{ss}} \sim \Delta l_{\text{mfp}} \sim \frac{1}{n_{\text{H I}} \sigma_{\text{H I}}(\nu_{\text{H I}})} \approx 5 \text{ pc} \left(\frac{n_{\text{H I}}}{0.01 \text{ cm}^{-3}} \right)^{-1}, \quad (\text{B1})$$

where $\sigma_{\text{H I}}$ is the hydrogen photoionization cross section and $\nu_{\text{H I}}$ is its ionization frequency, corresponding to 13.6 eV. This mean free path is generally much smaller than the spatial resolution of both our SPH simulations (Table 2) and our radiative transfer grids. One circumstance in which the self-shielding layer could be significantly thicker would be the case in which the illumination is dominated by a nearby quasar (e.g., Cantalupo et al. 2008; Kollmeier et al. 2010); we do not attempt to model this here, although it would be interesting to explore in the future.

Along each ray i , the optical depth from the background is calculated assuming that all the ionizing photons effectively have a frequency $\nu_{\text{H I}}$, $\tau_i = \int d\ln_{\text{H I}} \sigma_{\text{H I}}(\nu_{\text{H I}})$, and the photoionization rate is correspondingly attenuated as $\Gamma_{\text{H I}, i} = \Gamma_{\text{H I}}^{\text{bkg}} e^{-\tau_i}$. The equilibrium solution is obtained iteratively as follows. Trial values for the ionized fraction of hydrogen in each cell are first arbitrarily chosen. The optical depth from the background to each cell,

¹¹ This assumption is increasingly violated at high halo masses, where hot halo gas becomes more prevalent, but since the cooling Ly α emission also scales with \dot{M}_{cold} instead of \dot{M}_{gas} to first order, the errors made should approximately cancel when taking the ratio of stellar to cooling Ly α below.

along each normal direction, is then computed. The effective photoionization rate within the cell is taken to be the average of the attenuated value from each direction, $\Gamma_{\text{H I}} \equiv (1/6) \sum_i \Gamma_{\text{H I},i}$. Using this photoionization rate, an updated hydrogen fraction is calculated using the balance Equation (3). For this step, the electron number density n_e is taken equal to its value after the previous update. The procedure is repeated until the ionization fraction has converged in each cell. The convergence criterion is set to one part in 100.

In reality, the gas temperatures provided by the basic hydrodynamical simulations are in error within self-shielded regions. We thus store the mean attenuation from the background, $\langle e^{-\tau_i} \rangle$, in each cell as an indicator of self-shielding. Different prescriptions for the self-shielded gas and re-simulation procedures are explored in Section 3.2.

APPENDIX C

Ly α RADIATIVE TRANSFER: α RT

In this Appendix, we describe the new three-dimensional Ly α radiative transfer code, called “ α RT,” developed for this project.

C.1. Basic Physics and Algorithm

The basic algorithm is the standard Monte Carlo method, which follows the scatters of a prescribed number, N_{ph} , of Ly α photons and statistically estimates the emergent morphology and line profile from them (e.g., Zheng & Miralda-Escudé 2002a; Cantalupo et al. 2005; Dijkstra et al. 2006; Verhamme et al. 2006; Tasitsiomi 2006a; Laursen & Sommer-Larsen 2007; Laursen et al. 2009a). The principal improvements over many of these codes are that it is fully three dimensional and its implementation is parallel to take advantage of distributed computing resources. The code therefore scales well to large problems with complex geometries and so can be run on the outputs of large-scale hydrodynamical simulations. It produces Ly α images, spectra, as well as integral field images (spectra as a function of position on the sky). As it was independently developed, its results provide useful checks of existing numerical calculations. It has been demonstrated (Tasitsiomi 2006a; Laursen et al. 2009a) that the Monte Carlo algorithm can be combined with mesh refinement techniques relatively straightforwardly, which is particularly valuable for problems of large dynamic range, for instance when attempting to capture the details of scattering within the clumpy ISM of galaxies at the same time as the large-scale cosmological radiative transport. We plan to build on the infrastructure developed in this work and implement adaptive refinement in future versions of the code. Here, we briefly outline the basic physics and methodology of the current implementation. Our notation follows that of Dijkstra et al. (2006).

There are three relevant reference frames: the frame of the observer, the fluid frame, and the frame of a particular atom, which may have thermal motion with respect to the fluid frame. The observer is assumed to be located at $z = +\infty$ and to be stationary with respect to the simulation volume, apart from a possible Hubble flow that redshifts all the photons from the simulation volume by the same factor. We assume that the simulation volume is sufficiently small that its volume is essentially at a single redshift. A quantity Q in the observer’s frame is denoted Q' in the fluid frame and \tilde{Q} in the frame of the atom.

The Monte Carlo photons are injected with a probability per unit volume proportional to the local Ly α emissivity (Section 3).

The photons are assumed to have energy exactly at the line center in the frame moving with the fluid and the emission is taken to be isotropic. The natural and thermal broadening of the emission line have a negligible impact on the results, as the resonant scatters rapidly erase their memory. The Ly α photons are then each transported by simulating their resonant scatters.

The Ly α optical depth through an H I column density $N_{\text{H I}}$ for a photon at the line center is given by

$$\tau_0 = N_{\text{H I}} \sigma_{\text{Ly}\alpha}(\nu_0) = N_{\text{H I}} f_{12} \frac{\sqrt{\pi} e^2}{m_e c \Delta\nu_D} \approx 8.3 \times 10^6 \left(\frac{N_{\text{H I}}}{2 \times 10^{20} \text{ cm}^{-2}} \right) \left(\frac{T}{2 \times 10^4 \text{ K}} \right)^{-1/2}, \quad (\text{C1})$$

where $\sigma_{\text{Ly}\alpha}$ is the Ly α cross section, ν_0 is the Ly α line center, f_{12} is the oscillator strength of the transition, e is the charge of the electron, m_e is its mass, c is the speed of light, and $\Delta\nu_D = v_{\text{th}} \nu_0 / c$, with $v_{\text{th}} = (2k_B T / m_p)^{1/2}$. It is convenient to locally define a dimensionless frequency $x \equiv (\nu - \nu_0) / \Delta\nu_D$. The optical depth for photons of arbitrary energy is then

$$\tau_x = \tau_0 H(a, x) = \tau_0 \frac{a}{\pi} \int_{-\infty}^{\infty} \frac{e^{-y^2} dy}{(y - x)^2 + a^2}, \quad (\text{C2})$$

where $a = A_{21} / (4\pi \Delta\nu_D)$ is the Voigt parameter and A_{21} is the Ly α Einstein A coefficient.

In the absence of perturbations, the absorption of a Ly α photon is quickly followed (after $A_{21}^{-1} \sim 10^{-9}$ s) by the re-emission of another Ly α photon of the same energy. Owing to the motion of the atom, the scatter is however not coherent in the observer’s frame. Defining $\mathbf{v}_a \equiv \mathbf{v}_{\text{bulk}} + \mathbf{v}_{\text{th}}$ to be the net velocity of the scattering atom, the frequency of a photon of incident frequency x_i after scattering is

$$x_o = x_i + \frac{\mathbf{v}_a \cdot (\mathbf{k}_i - \mathbf{k}_o)}{v_{\text{th}}} + g(\mathbf{k}_i \cdot \mathbf{k}_o - 1) + \mathcal{O}\left(\left(\frac{v_{\text{th}}}{c}\right)^2\right), \quad (\text{C3})$$

where \mathbf{k}_i and \mathbf{k}_o are unit vectors in the directions of the incident and outgoing photons, respectively (e.g., Dijkstra et al. 2006). The Hubble expansion is modeled by including a component $\mathbf{v}_H = H(z)\mathbf{r}$, where \mathbf{r} is the displacement vector from the center of the system and $H(z)$ is the Hubble parameter, in the bulk velocity term. The g “recoil” term accounts for the average transfer of momentum from the photon to the atom during the scatter and can be written as $g = h\Delta\nu_D / 2k_B T$ (Field 1959). The effect of recoil is usually small (Adams 1971) and is ignored in our calculations.

Because the optical depth near the Ly α line center is very large in astrophysical conditions, a Ly α photon typically travels only a short distance before being scattered. The numerous scatters result in a random walk in space. At the same time, the scatters cause the photon to oscillate in frequency and it usually escapes the medium during an excursion far into a damping wing of the line, where the optical depth is greatly reduced (e.g., Zanstra 1949; Unno 1952; Field 1959; Osterbrock 1962; Adams 1972; Harrington 1973; Neufeld 1990; Gould & Weinberg 1996). As a result, the emergent Ly α line profile is heavily affected by the medium through which it propagates and therefore provides a probe of the latter.

To simulate the scatters of the Monte Carlo photons, we at each step randomly pick the optical depth τ before the next scatter. By definition, τ has a PDF $P(\tau) \propto e^{-\tau}$. Given the frequency and propagation direction of the photon, we integrate

through the medium until an optical depth τ has been reached; this defines the position of the next scatter. The frequency seen by a fluid element along the way in general differs from that seen by the observer and is given by the Doppler shift formula $x'_i = x_i - \mathbf{k}_i \cdot \mathbf{v}_{\text{bulk}}/v_{\text{th}}$. The frequency of the outgoing photon is then calculated using Equation (C3), which requires picking \mathbf{k}_o and \mathbf{v}_{th} .

The direction of the outgoing photon is picked from a dipolar phase function, $P(\mathbf{k}_i \cdot \mathbf{k}_o) \propto 1 + (\mathbf{k}_i \cdot \mathbf{k}_o)^2$. Owing to quantum mechanical effects, this Rayleigh phase function is strictly valid only for scatters with $\tilde{x}_i \gtrsim 0.2$ (see Dijkstra & Loeb 2008, and references therein), but the multiple scatters act to quickly randomize propagation directions so that the precise phase function is unimportant. An isotropic phase function, $P(\mathbf{k}_i \cdot \mathbf{k}_o) \propto \text{const}$, gives identical results in practically all astrophysical conditions, aside for polarization, which we neglect here.

The first part of the scattering atom velocity is simply the local bulk velocity of the fluid, \mathbf{v}_{bulk} . To this, a thermal velocity \mathbf{v}_{th} must be added. In a frame moving with the fluid, but oriented such that the z' axis is parallel to \mathbf{k}_i , the thermal velocity can be expressed as $v_{\text{th}}(u'_{\perp 1}, u'_{\perp 2}, u'_{\parallel})$, where u'_{\parallel} is the normalized component parallel to the direction of propagation, and $u'_{\perp 1,2}$ are perpendicular components. The probability of scattering off an atom with a parallel component u'_{\parallel} is the product of the one-dimensional thermal distribution and the absorption cross section,

$$P(u'_{\parallel}) = \frac{a}{\pi} \frac{e^{-u'^2_{\parallel}}}{(x'_i - u'_{\parallel})^2 + a^2} H^{-1}(a, x'_i), \quad (\text{C4})$$

while $u'_{\perp 1,2}$ are simply thermally distributed since the probability of absorption is independent of these normal components.

The procedure is repeated for each scatter of each Monte Carlo until it escapes the computational volume. In this work, we assume that all photons escape the medium but the algorithm can be straightforwardly extended to model photon destruction (e.g., by dust) by specifying a destruction probability at each integration step or scattering event.

C.2. Numerical Aspects

Having described the essential aspects of the algorithm, we now elaborate on numerical aspects of its implementation relevant to its accuracy and performance.

The generation of random numbers is a key element of the Monte Carlo method. In most instances, these are straightforwardly generated using either a transformation method or a rejection method (e.g., Press et al. 1992). An important exception is the generation of u'_{\parallel} , whose PDF (Equation (C4)) is not analytically integrable and for which there is not an obvious efficient bounding function to use in the rejection method. Since a variate from this distribution is generated at each scatter, it is important to have a fast implementation. We generalize the rejection method described in the appendix of Zheng & Miralda-Escudé (2002b) to use three critical velocities (in their notation, u_1 and u_2 in addition to u_0) in order to minimize the number of rejections.

Although the distance between neighboring grid points is $\Delta l \equiv L/N_p$, photon trajectories are traced exactly in the sense that the photon can be located at any coordinate (x, y, z) in the box along its ray. This can greatly improve the accuracy of the calculations possible with limited computational resources. For instance, a large and/or very optically thick medium may

be well described by a relatively coarse grid if it is fairly homogeneous. However, because of the short mean free path of the Ly α photons, the radiative transfer would be poorly captured if the position of the photon were restricted to widely separated grid points. The integration along rays is still done in finite steps of length Δl_{int} , but this parameter can be specified arbitrarily and in particular can be $\ll \Delta l$. Since the photon position is in general between the grid points on which physical properties are stored, these must be interpolated onto the ray. Our code can use either nearest grid point (NGP) or cloud-in-cell (CIC) interpolation (e.g., Hockney & Eastwood 1988). While CIC is more accurate, NGP produces similar results but requires a factor of several fewer operations per integration step. We use NGP interpolation in the calculations presented in this work.

Finally, we use an “accelerated scheme” (e.g., Ahn et al. 2002) that greatly speeds up radiative transfer calculations by skipping scatters in the core of the Ly α line, during which the photon is nearly stationary in space. Specifically, we define a critical frequency x_{crit} and say that a photon with $|x'| < x_{\text{crit}}$ is in the core of the line. Photons in the line core are then forced into the wings by allowing them to scatter only off high-velocity atoms. This is achieved by truncating the thermal Gaussian distribution from which $u_{\perp 1}$ and $u_{\perp 2}$ are drawn to zero for $|u| < x_{\text{crit}}$. Skipping core scatters effectively sets the mean free path of core photons to zero. In the limit $x_{\text{crit}} \rightarrow 0$, the algorithm becomes exact and we can vary x_{crit} to verify the convergence of our calculations. In practice, we find that $x_{\text{crit}} = 3$ generally speeds up the calculations (by as much as orders of magnitude in the very optically thick regime) while producing faithful results. We use this value as default in this work. Following Tasitsiomi (2006a, 2006b) and Laursen et al. (2009a), we further speed up the code by making use of the fact that in extremely optically thick radiative transfer cells (which, in the discretized numerical scheme, are internally static and uniform), the solution to the radiative transfer problem is well approximated by a re-scaled version of the analytic solution obtained by Neufeld (1990). Specifically, the explicit scatters in cells with $a\tau_0 \gtrsim 10^3$ can be skipped by choosing the emerging frequency according to the known analytic solution.

C.3. Visualization

The principal application of our code is to produce images and spectra that can be compared to observations. To do so, the code keeps track of a three-dimensional “integral field” array, of dimensions $N_g^2 \times N_v$. The first two dimensions correspond to the projected position on the sky as viewed by the observer, or pixels. The third dimension divides each pixel into N_v frequency bins between specified boundaries. From the full integral field array, with a spectrum at each pixel, an Ly α image can then be constructed by summing over frequency bins within each pixel. A spectrum toward any particular direction, or along an arbitrarily placed slit, can be produced by considering only the relevant pixels.

To construct the integral field array, a number proportional to the probability that the outgoing photon escapes directly toward the observer is added to the corresponding pixel and frequency bin at each scatter. Since the observer is assumed to be located at $z = \infty$, we take this numerical increment to be

$$\frac{3}{16\pi} [1 + (\mathbf{k}_i \cdot \mathbf{z})^2] e^{-\tau_z}. \quad (\text{C5})$$

The factor in the square brackets accounts for the probability that the photon is re-emitted toward the observer and $e^{-\tau_z}$ is




Figure 11. Redistribution function test: PDF of the outgoing dimensionless frequency x as a function of the incident x_0 , for $x_0 = 0, 1, 2, 3, 4$, and 5 from left to right in each panel. The histograms show the PDFs estimated by our Monte Carlo code and the smooth curves show the corresponding analytic solutions. The left panel shows the case of an isotropic redistribution function and the right panel shows the case of a dipole redistribution function.

the probability of escape in a single flight, where τ_z is the optical depth to $z = \infty$, given the frequency that the photon would have had if it had been scattered in that direction. This procedure (see also Zheng & Miralda-Escudé 2002a; Cantalupo et al. 2005) is necessary because only an infinitesimal fraction of the Monte Carlo photons actually emerge exactly in the $z = \infty$ direction and is termed the “next-event estimator” (Dupree & Fraley 2002). To convert the dimensionless counter to observed physical intensity, we multiply by $L_\alpha / [N_{\text{ph}}(\Delta l)^2(1+z)^4]$, where L_α is the total Ly α luminosity and the z factor accounts for cosmological surface brightness dimming. We have verified that this next-event estimator conserves the number of photons in isotropic cases, i.e., that the luminosity inferred from the resultant image is equal to the luminosity of the central source, from which only a small fraction of the emitted photons directly reach the observer.

C.4. Test Problems

In order to be confident in the results obtained from our radiative transfer code, it is important to test it on problems with known solutions. Figures 11–13 show the results of tests of the αRT code for the following problems:

1. *Redistribution function.* The redistribution function, $R(x'_o, x'_i)$, is the PDF of the outgoing photon dimensionless frequency, x'_o , as a function of the incoming photon frequency, x'_i , in the fluid frame. Although not a solution to the radiative transfer problem *per se*, the Monte Carlo algorithm effectively generates random variates from this distribution at each scatter when it picks \mathbf{v}_{th} and \mathbf{k}_o , and applies Equation (C3). A useful test of this key portion of the algorithm is thus to compare the Monte-Carlo-generated redistribution function with analytically derived expressions (Unno 1952; Hummer 1962; Lee 1974). Figure 11 compares the results of our code with the analytic solutions presented by Lee (1974) (for a naturally broadened intrinsic line) for the cases of isotropic and dipole phase functions. To reproduce the redistribution function for core scatters (small x'_i), the accelerated scheme must be turned off ($x_{\text{crit}} = 0$).
2. *Static sphere.* An analytic solution to the Ly α radiative transfer problem exists for the case of a monochromatic point source at the center of an extremely optically thick, uniform, and static sphere of line-center optical depth τ_0




Figure 12. Static sphere test: emergent line profile for a monochromatic source at the center of a uniform static sphere for optical depths at the line center $\tau_0 = 10^5, 10^6$, and 10^8 . The histograms show the numerical solutions obtained with our code and the smooth curves show the corresponding analytical solutions following Dijkstra et al. (2006) (Equation (C6)). The numerical solution was computed for a temperature $T = 10$ K.

from the source to the surface:

$$J(x) = \frac{\pi^{3/2}}{\sqrt{6}a\tau_0} \left\{ \frac{x^2}{1 + \cosh[\sqrt{2\pi^3/27}(|x|^3/a\tau_0)]} \right\} \quad (\text{C6})$$

(Dijkstra et al. 2006). This solution is a generalization of the plane-parallel results obtained by Harrington (1973) and Neufeld (1990). It differs by a factor of 2π from the expression in Dijkstra et al. (2006) because it is normalized to integrate to unity here. Figure 12 compares the results obtained with our code for $\tau_0 = 10^5, 10^6$, and 10^7 and a gas temperature $T = 10$ K with this solution.

3. *Collapsing/outflowing sphere.* Velocity fields are crucial to the Ly α radiative transfer physics, as Doppler shifts modify the perceived optical depth and can open “paths of least resistance” through which photons can escape the medium. It is therefore important to test the code in dynamic situations. Unfortunately, few analytic results with velocity fields are available. We therefore instead compare our numerical solutions against those obtained with other codes. Specifically, we consider the following




Figure 13. Collapsing and outflowing sphere tests. The left panel shows the numerical results for a collapsing uniform sphere of radius R , neutral hydrogen column density $N_{\text{H I}} = 2 \times 10^{20} \text{ cm}^{-2}$ from the center to the surface, and temperature $T = 20,000 \text{ K}$. In each case, the sphere is collapsing with a velocity $\mathbf{v} = (-r/R)v_{\text{max}}\hat{\mathbf{r}}$, with $v_{\text{max}} = 20$ (dotted), 200 (dashed), and 2000 km s^{-1} (solid). The right panel shows the same results for an expanding sphere with the same parameters, but reversed velocity field. The results can be compared with Figure 7 of Verhamme et al. (2006) for the expanding case. The collapsing case is symmetric under the replacement $x \rightarrow -x$.

cases computed by Verhamme et al. (2006; see also Zheng & Miralda-Escudé 2002a; Dijkstra et al. 2006).

1. A collapsing uniform sphere of radius R with monochromatic central source: $N_{\text{H I}} = 2 \times 10^{20} \text{ cm}^{-2}$ from the center to the surface, $T = 20,000 \text{ K}$, $\mathbf{v} = -(r/R)v_{\text{max}}\hat{\mathbf{r}}$ with $v_{\text{max}} = 20, 200$, and 2,000 km s^{-1} .
2. Same, but expanding with $\mathbf{v} = (r/R)v_{\text{max}}\hat{\mathbf{r}}$ for $v_{\text{max}} = 20, 200$, and 2000 km s^{-1} .

The numerical solutions obtained with our code shown in Figure 13 agree well with Figure 7 of Verhamme et al. (2006) for the expanding case. The collapsing case is symmetric under the replacement $x \rightarrow -x$.

REFERENCES

- Adams, T. F. 1971, *ApJ*, **168**, 575
Adams, T. F. 1972, *ApJ*, **174**, 439
Ahn, S.-H., Lee, H.-W., & Lee, H. M. 2002, *ApJ*, **567**, 922
Aubert, D., & Teyssier, R. 2010, ArXiv e-prints
Barnes, J., & Hut, P. 1986, *Nature*, **324**, 446
Binney, J. 1977, *ApJ*, **215**, 483
Birnbom, Y., & Dekel, A. 2003, *MNRAS*, **345**, 349
Birnbom, Y., Dekel, A., & Neistein, E. 2007, *MNRAS*, **380**, 339
Bower, R. G., et al. 2004, *MNRAS*, **351**, 63
Brooks, A. M., Governato, F., Quinn, T., Brook, C. B., & Wadsley, J. 2009, *ApJ*, **694**, 396
Burlles, S., Nollett, K. M., & Turner, M. S. 2001, *ApJ*, **552**, L1
Cantalupo, S. 2010, *MNRAS*, **403**, L16
Cantalupo, S., Porciani, C., & Lilly, S. J. 2008, *ApJ*, **672**, 48
Cantalupo, S., Porciani, C., Lilly, S. J., & Miniati, F. 2005, *ApJ*, **628**, 61
Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *ApJ*, **292**, 371
Dayal, P., Ferrara, A., & Saro, A. 2010, *MNRAS*, **402**, 1449
DeBuhr, J., Quataert, E., Ma, C., & Hopkins, P. 2010, *MNRAS*, **406**, L55
Dekel, A., & Birnbom, Y. 2006, *MNRAS*, **368**, 2
Dekel, A., et al. 2009, *Nature*, **457**, 451
Dey, A., et al. 2005, *ApJ*, **629**, 654
Dijkstra, M., Haiman, Z., & Spaans, M. 2006, *ApJ*, **649**, 14
Dijkstra, M., Lidz, A., & Wyithe, J. S. B. 2007, *MNRAS*, **377**, 1175
Dijkstra, M., & Loeb, A. 2008, *MNRAS*, **386**, 492
Dijkstra, M., & Loeb, A. 2009, *MNRAS*, **400**, 1109
Dijkstra, M., & Wyithe, S. 2010, *MNRAS*, **408**, 352
Di Matteo, T., Springel, V., & Hernquist, L. 2005, *Nature*, **433**, 604
Dupree, S. A., & Fraley, S. K. 2002, A Monte Carlo Primer: A Practical Approach to Radiation Transport (Dordrecht: Kluwer)
Efstathiou, G. 1992, *MNRAS*, **256**, 43P
Fabian, A. C., & Nulsen, P. E. J. 1977, *MNRAS*, **180**, 479
Fakhouri, O., Ma, C., & Boylan-Kolchin, M. 2010, *MNRAS*, **406**, 2267
Fardal, M. A., Katz, N., Gardner, J. P., Hernquist, L., Weinberg, D. H., & Davé, R. 2001, *ApJ*, **562**, 605
Faucher-Giguère, C.-A., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008a, *ApJ*, **682**, L9
Faucher-Giguère, C.-A., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008b, *ApJ*, **688**, 85
Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2008c, *ApJ*, **673**, 39
Faucher-Giguère, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2009, *ApJ*, **703**, 1416
Faucher-Giguère, C.-A., Prochaska, J. X., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008d, *ApJ*, **681**, 831
Field, G. B. 1959, *ApJ*, **129**, 551
Francis, P. J., Palunas, P., Teplitz, H. I., Williger, G. M., & Woodgate, B. E. 2004, *ApJ*, **614**, 75
Francis, P. J., et al. 1996, *ApJ*, **457**, 490
Francis, P. J., et al. 2001, *ApJ*, **554**, 1001
Furlanetto, S. R., Schaye, J., Springel, V., & Hernquist, L. 2005, *ApJ*, **622**, 7
Geach, J. E., Smail, I., Chapman, S. C., Alexander, D. M., Blain, A. W., Stott, J. P., & Ivison, R. J. 2007, *ApJ*, **655**, L9
Geach, J. E., et al. 2005, *MNRAS*, **363**, 1398
Geach, J. E., et al. 2009, *ApJ*, **700**, 1
Gingold, R. A., & Monaghan, J. J. 1977, *MNRAS*, **181**, 375
Goerdt, T., Dekel, A., Sternberg, A., Ceverino, D., Teyssier, R., & Primack, J. R. 2010, *MNRAS*, **407**, 613
Gould, A., & Weinberg, D. H. 1996, *ApJ*, **468**, 462
Haardt, F., & Madau, P. 1996, *ApJ*, **461**, 20
Haiman, Z., & Rees, M. J. 2001, *ApJ*, **556**, 87
Haiman, Z., Spaans, M., & Quataert, E. 2000, *ApJ*, **537**, L5
Hansen, M., & Oh, S. P. 2006, *MNRAS*, **367**, 979
Harrington, J. P. 1973, *MNRAS*, **162**, 43
Hernquist, L. 1987, *ApJS*, **64**, 715
Hernquist, L. 1990, *ApJ*, **356**, 359
Hockney, R. W., & Eastwood, J. W. 1988, Computer Simulation Using Particles (Bristol: Hilger)
Hopkins, P. F., & Hernquist, L. 2006, *ApJS*, **166**, 1
Hu, E. M. 1992, *ApJ*, **391**, 608
Hui, L., & Gnedin, N. Y. 1997, *MNRAS*, **292**, 27
Hummer, D. G. 1962, *MNRAS*, **125**, 21
Katz, N., & Gunn, J. E. 1991, *ApJ*, **377**, 365
Katz, N., Keres, D., Dave, R., & Weinberg, D. H. 2003, in The IGM/Galaxy Connection. The Distribution of Baryons at $z = 0$, ed. J. L. Rosenberg & M. E. Putman (Astrophysics and Space Science Library, Vol. 281; Dordrecht: Kluwer), 185
Katz, N., Quinn, T., Bertschinger, E., & Gelb, J. M. 1994, *MNRAS*, **270**, L71
Katz, N., Weinberg, D. H., & Hernquist, L. 1996, *ApJS*, **105**, 19
Katz, N., & White, S. D. M. 1993, *ApJ*, **412**, 455
Keel, W. C., Cohen, S. H., Windhorst, R. A., & Waddington, I. 1999, *AJ*, **118**, 2547
Keel, W. C., White, R. E., Chapman, S., & Windhorst, R. A. 2009, *AJ*, **138**, 986
Kennicutt, R. C., Jr. 1998, *ApJ*, **498**, 541

- Kereš, D., Katz, N., Davé, R., Fardal, M., & Weinberg, D. H. 2009a, *MNRAS*, **396**, 2332
- Kereš, D., Katz, N., Fardal, M., Davé, R., & Weinberg, D. H. 2009b, *MNRAS*, **395**, 160
- Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, *MNRAS*, **363**, 2
- Kollmeier, J. A., Zheng, Z., Davé, R., Gould, A., Katz, N., Miralda-Escudé, J., & Weinberg, D. H. 2010, *ApJ*, **708**, 1048
- Komatsu, E., et al. 2009, *ApJS*, **180**, 330
- Kornei, K. A., Shapley, A. E., Erb, D. K., Steidel, C. C., Reddy, N. A., Pettini, M., & Bogosavljević, M. 2010, *ApJ*, **711**, 693
- Kunth, D., Mas-Hesse, J. M., Terlevich, E., Terlevich, R., Lequeux, J., & Fall, S. M. 1998, *A&A*, **334**, 11
- Laursen, P., Razoumov, A. O., & Sommer-Larsen, J. 2009a, *ApJ*, **696**, 853
- Laursen, P., & Sommer-Larsen, J. 2007, *ApJ*, **657**, L69
- Laursen, P., Sommer-Larsen, J., & Andersen, A. C. 2009b, *ApJ*, **704**, 1640
- Laursen, P., Sommer-Larsen, J., & Razoumov, A. O. 2010, arXiv:1009.1384
- Lee, J.-S. 1974, *ApJ*, **192**, 465
- Leitherer, C., et al. 1999, *ApJS*, **123**, 3
- Lucy, L. B. 1977, *AJ*, **82**, 1013
- Mas-Hesse, J. M., Kunth, D., Tenorio-Tagle, G., Leitherer, C., Terlevich, R. J., & Terlevich, E. 2003, *ApJ*, **598**, 858
- Matsuda, Y., Yamada, T., Hayashino, T., Yamauchi, R., & Nakamura, Y. 2006, *ApJ*, **640**, L123
- Matsuda, Y., et al. 2004, *AJ*, **128**, 569
- McBride, J., Fakhouri, O., & Ma, C.-P. 2009, *MNRAS*, **398**, 1858
- Moller, P., & Warren, S. J. 1998, *MNRAS*, **299**, 661
- Mori, M., Umemura, M., & Ferrara, A. 2004, *ApJ*, **613**, L97
- Murray, N., Quataert, E., & Thompson, T. A. 2005, *ApJ*, **618**, 569
- Nagamine, K., Choi, J., & Yajima, H. 2010, arXiv:1006.5345
- Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, *ApJ*, **490**, 493
- Neistein, E., & Dekel, A. 2008, *MNRAS*, **383**, 615
- Neufeld, D. A. 1990, *ApJ*, **350**, 216
- Neufeld, D. A. 1991, *ApJ*, **370**, L85
- Nilsson, K. K., Fynbo, J. P. U., Møller, P., Sommer-Larsen, J., & Ledoux, C. 2006, *A&A*, **452**, L23
- Ocvirk, P., Pichon, C., & Teysier, R. 2008, *MNRAS*, **390**, 1326
- Ohya, Y., et al. 2003, *ApJ*, **591**, L9
- Oppenheimer, B. D., Davé, R., Kereš, D., Fardal, M., Katz, N., Kollmeier, J. A., & Weinberg, D. H. 2010, *MNRAS*, **406**, 2325
- Osterbrock, D. E. 1962, *ApJ*, **135**, 195
- Osterbrock, D. E., & Ferland, G. J. 2006, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei* (2nd ed.; Sausalito, CA: Univ. Science Books)
- Ouchi, M., et al. 2009, *ApJ*, **696**, 1164
- Palunas, P., Teplitz, H. I., Francis, P. J., Williger, G. M., & Woodgate, B. E. 2004, *ApJ*, **602**, 545
- Prescott, M. K. M., Dey, A., & Jannuzi, B. T. 2009, *ApJ*, **702**, 554
- Prescott, M. K. M., Kashikawa, N., Dey, A., & Matsuda, Y. 2008, *ApJ*, **678**, L77
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in C. The Art of Scientific Computing* (2nd ed.; Cambridge: Cambridge Univ. Press)
- Rauch, M., et al. 2008, *ApJ*, **681**, 856
- Rees, M. J., & Ostriker, J. P. 1977, *MNRAS*, **179**, 541
- Rybicki, G. B., & Loeb, A. 1999, *ApJ*, **520**, L79
- Saito, T., Shimasaku, K., Okamura, S., Ouchi, M., Akiyama, M., & Yoshida, M. 2006, *ApJ*, **648**, 54
- Saito, T., Shimasaku, K., Okamura, S., Ouchi, M., Akiyama, M., Yoshida, M., & Ueda, Y. 2008, *ApJ*, **675**, 1076
- Santos, M. R. 2004, *MNRAS*, **349**, 1137
- Schaye, J. 2004, *ApJ*, **609**, 667
- Shapley, A. E., Steidel, C. C., Pettini, M., & Adelberger, K. L. 2003, *ApJ*, **588**, 65
- Silk, J. 1977, *ApJ*, **211**, 638
- Smith, D. J. B., & Jarvis, M. J. 2007, *MNRAS*, **378**, L49
- Sommer-Larsen, J. 2006, *ApJ*, **644**, L1
- Springel, V. 2005, *MNRAS*, **364**, 1105
- Springel, V. 2010, *MNRAS*, **401**, 791
- Springel, V., Di Matteo, T., & Hernquist, L. 2005, *MNRAS*, **361**, 776
- Springel, V., & Hernquist, L. 2002, *MNRAS*, **333**, 649
- Springel, V., & Hernquist, L. 2003, *MNRAS*, **339**, 289
- Steidel, C. C., Adelberger, K. L., Shapley, A. E., Pettini, M., Dickinson, M., & Giavalisco, M. 2000, *ApJ*, **532**, 170
- Steidel, C. C., Erb, D. K., Shapley, A. E., Pettini, M., Reddy, N. A., Bogosavljević, M., Rudie, G. C., & Rakic, O. 2010, *ApJ*, **717**, 289
- Taniguchi, Y., & Shioya, Y. 2000, *ApJ*, **532**, L13
- Taniguchi, Y., Shioya, Y., & Kakazu, Y. 2001, *ApJ*, **562**, L15
- Tasitsiomi, A. 2006a, *ApJ*, **645**, 792
- Tasitsiomi, A. 2006b, *ApJ*, **648**, 762
- Unno, W. 1952, *PASJ*, **4**, 100
- Verhamme, A., Schaerer, D., & Maselli, A. 2006, *A&A*, **460**, 397
- Webb, T. M. A., Yamada, T., Huang, J.-S., Ashby, M. L. N., Matsuda, Y., Egami, E., Gonzalez, M., & Hayashimo, T. 2009, *ApJ*, **692**, 1561
- Weidinger, M., Møller, P., & Fynbo, J. P. U. 2004, *Nature*, **430**, 999
- Weidinger, M., Møller, P., Fynbo, J. P. U., & Thomsen, B. 2005, *A&A*, **436**, 825
- Weinberg, D. H., Hernquist, L., & Katz, N. 1997, *ApJ*, **477**, 8
- White, M. 2002, *ApJS*, **143**, 241
- White, S. D. M., & Rees, M. J. 1978, *MNRAS*, **183**, 341
- Wiersma, R. P. C., Schaye, J., & Smith, B. D. 2009, *MNRAS*, **393**, 99
- Wilman, R. J., Gerssen, J., Bower, R. G., Morris, S. L., Bacon, R., de Zeeuw, P. T., & Davies, R. L. 2005, *Nature*, **436**, 227
- Wolfe, A. M., & Chen, H. 2006, *ApJ*, **652**, 981
- Yang, Y., Zabludoff, A., Eisenstein, D., & Davé, R. 2010, *ApJ*, **719**, 1654
- Yang, Y., Zabludoff, A., Tremonti, C., Eisenstein, D., & Davé, R. 2009, *ApJ*, **693**, 1579
- Yang, Y., Zabludoff, A. I., Davé, R., Eisenstein, D. J., Pinto, P. A., Katz, N., Weinberg, D. H., & Barton, E. J. 2006, *ApJ*, **640**, 539
- Zanstra, H. 1949, *Bull. Astron. Inst. Neth.*, **11**, 1
- Zhao, D. H., Jing, Y. P., Mo, H. J., & Börner, G. 2003, *ApJ*, **597**, L9
- Zheng, Z., Cen, R., Trac, H., & Miralda-Escudé, J. 2010a, *ApJ*, **716**, 574
- Zheng, Z., Cen, R., Trac, H., & Miralda-Escudé, J. 2010b, arXiv:1003.4990
- Zheng, Z., & Miralda-Escudé, J. 2002a, *ApJ*, **578**, 33
- Zheng, Z., & Miralda-Escudé, J. 2002b, *ApJ*, **568**, L71