

Interpreting ALMA observations of the ISM during the epoch of reionization

Harley Katz,[★] Taysun Kimm, Debora Sijacki and Martin G. Haehnelt

Institute of Astronomy and Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

Accepted 2017 March 9. Received 2017 March 9; in original form 2016 December 6

ABSTRACT

We present cosmological, radiation-hydrodynamics simulations of galaxy formation during the epoch of reionization in an effort towards modelling the interstellar medium (ISM) and interpreting Atacama Large Millimeter Array (ALMA) observations. Simulations with and without stellar radiation are compared at large (Mpc), intermediate (tens of kpc) and small (sub-kpc) scales. At large scales, the dense regions around galaxies reionize first before ultraviolet (UV) photons penetrate the voids; however, considerable amounts of neutral gas remain present within the haloes. The spatial distribution of neutral gas is highly dynamic and is anticorrelated with the presence of stars older than a few Myr. For our specific feedback implementation, most of the metals remain inside the virial radii of haloes, and they are proportionally distributed over the ionized and neutral media by mass. For our most massive galaxy with $M_h \sim 10^{11} M_\odot$, the majority of the C II and O I masses are associated with cold neutral clumps. N II is more diffuse and arises in warmer gas, while O III arises in hotter gas with a higher ionization parameter, produced by photoheating and supernovae. If smaller pockets of high-metallicity gas exist in the ISM, the emission from these ions may be observable by ALMA, while the low metallicity of the galaxy may cause these systems to fall below the local [C II]–star formation rate relation. The presence of dust can cause spatial offsets between UV/Lyman α and [C II] emissions, as suggested by the recent observations of Maiolino et al. [O III] may be spatially offset from both of these components since it arises from a different part of density–temperature phase space.

Key words: galaxies: evolution – galaxies: formation – galaxies: high-redshift – intergalactic medium – infrared: ISM.

1 INTRODUCTION

In the local Universe, the interstellar medium (ISM) of various different galaxy types has been well studied observationally (Kennicutt et al. 2003, 2011). In low-redshift metal-enriched galaxies, fine structure lines are likely to be the dominant coolant of the ISM at $T < 10^4$ K (Spitzer 1978). At energies above the ionization potential of neutral hydrogen (13.6 eV), forbidden lines such as [N II] and [O III] trace the ionized medium, while at lower energies, other lines such as [C II], [O I] or [C I] can trace the neutral ISM (Carilli & Walter 2013). With the *Spitzer Space Telescope*, local galaxies can be spatially resolved in the near-infrared (near-IR), which has allowed for accurate measurements of various ISM properties such as dust mass, dust-to-gas ratio, dust properties and stellar properties (Kennicutt et al. 2003; Draine et al. 2007). More recently, the *Herschel Space Observatory* (Pilbratt et al. 2010) has been particularly important for characterizing the ISM in local

galaxies due to its high spatial resolution in the far-IR, its extended wavelength coverage and its spectrometer. This allows many of the primary emission lines (i.e. [O I], [O III], [N II], [C I] and [C II]) that originate from the different phases of the ISM to be probed (Kennicutt et al. 2011). Combining these observations with dust radiative transfer (RT) modelling can reveal important insights into dust mass, dust production mechanisms and star formation (e.g. De Looze et al. 2016).

At high redshift and, in particular, during the epoch of reionization, our understanding of the ISM and internal properties of galaxies is much less advanced. While hundreds of galaxies have been detected, only the more global properties such as star formation rate (SFR), ultraviolet (UV) luminosity function and UV continuum slopes have been constrained (Zheng et al. 2012; Coe et al. 2013; Ellis et al. 2013; McLure et al. 2013; Oesch et al. 2013, 2014; Bouwens et al. 2014, 2015). Various campaigns have targeted [C II] at 158 μm at $z > 6$ as this is expected to dominate cooling and to be the strongest emission line (Carilli & Walter 2013). Furthermore, it is well established in low-redshift galaxies that [C II] emission correlates with the SFR (e.g. De Looze et al. 2014). The interpretation

[★]E-mail: hk380@ast.cam.ac.uk

of this emission is complex as [C II] can be excited in photodissociation regions (PDRs), the warm and cold neutral media, partially ionized gas, as well as shocked gas (Madden et al. 1997; Kaufman et al. 1999; Graciá-Carpio et al. 2011; Cormier et al. 2012; Appleton et al. 2013; Pineda, Langer & Goldsmith 2014; Velusamy & Langer 2014). Quasar host galaxies have been targeted successfully (Walter et al. 2009; Wang et al. 2013), and [C II] has even been observed in the most distant spectroscopically confirmed quasar (Venemans et al. 2012). However, the galaxies that host these bright quasars tend to be the most massive galaxies in the Universe (e.g. Sijacki, Springel & Haehnelt 2009; Costa, Sijacki & Haehnelt 2014) and do not represent typical star-forming galaxies during the epoch of reionization. Recent observations with the Atacama Large Millimeter Array (ALMA) are beginning to probe [C II] in galaxies that are more representative of the high-redshift galaxy population and have SFRs of $\sim 10 M_{\odot} \text{ yr}^{-1}$ (Capak et al. 2015; Maiolino et al. 2015; Willott et al. 2015; Knudsen et al. 2016; Pentericci et al. 2016). These observations often find that the high-redshift galaxies fall below the local [C II]–SFR relations, which indicates that the internal properties of these high-redshift galaxies are different from what we observe locally.

Many theoretical studies have attempted to model the far-IR emission that may emanate from star-forming galaxies, both with a semi-analytic approach (Gong et al. 2012; Muñoz & Furlanetto 2014; Popping et al. 2014) and with numerical simulations (Nagamine, Wolfe & Hernquist 2006; Vallini et al. 2013, 2015; Pallottini et al. 2017). Nagamine et al. (2006) post-processed smoothed particle hydrodynamics (SPH) simulations that employed a subgrid model of the ISM and found that the amount of [C II] emission crucially depended on the amount of neutral gas present in the galaxy. Likewise, Vallini et al. (2013) used an SPH code to model the formation of a single high-redshift ($z = 6.6$) galaxy at higher resolution and post-processed the simulation with a Monte Carlo RT scheme to track the UV photons that are primarily responsible for exciting the [C II] emission line. However, the underlying simulation included neither star formation processes, nor radiative cooling nor supernova feedback, which are integral to properly model baryons in galaxies. Nevertheless, the Vallini et al. (2013) simulations have predicted that the [C II] emission is spatially offset from where the UV continuum of the galaxy is expected to dominate. This prediction is consistent with a number of observations of high-redshift galaxies that show an offset between the [C II] and UV and Ly α emission (Gallerani et al. 2012; Capak et al. 2015; Maiolino et al. 2015; Willott et al. 2015). This probably means that the [C II] emission likely originates in dense neutral clumps surrounding the galaxy and that high-redshift molecular clouds are rapidly disrupted by strong stellar feedback (Maiolino et al. 2015). Vallini et al. (2015) have improved on the Vallini et al. (2013) simulations by using a more detailed recipe for the metal distribution and identifying the likely locations of molecular clouds. They concluded that most of the [C II] emission likely originates in PDRs. However, the underlying simulations still lacked some of the relevant physical processes (most importantly radiative cooling).

More recently, Pallottini et al. (2017) used a zoom-in technique to model the formation of a $\sim 10^{11} M_{\odot} h^{-1}$ halo at 30 pc resolution. These simulations included a model for the formation of H₂ (Krumholz, McKee & Tumlinson 2008, 2009; McKee & Krumholz 2010) that allowed the authors to follow the atomic to molecular transition and make inferences on the correlation between [C II] emission and the location of H₂. They found that 95 per cent of the [C II] emission originates from the H₂ disc, consistent with the Vallini et al. (2015) simulation. The Pallottini et al. (2017)

simulations, however, do not include on-the-fly RT and therefore do not properly model the enhancement of the local Lyman–Werner background above the mean metagalactic background, inside of the galaxy.

Other simulations have been performed that include on-the-fly RT and detailed non-equilibrium chemistry models. These tend to be full box simulations with small (< 100 Mpc) volumes or zoom-in simulations of individual haloes or regions (e.g. Wise et al. 2012; Gnedin 2014; Kimm & Cen 2014; Xu et al. 2016). These have generally been performed in the context of reionization. While it is well established that the Universe becomes fully ionized at some time around $z \sim 6$ –7 (Fan, Carilli & Keating 2006; Ouchi et al. 2010; Bolton et al. 2011; Choudhury et al. 2015; Planck Collaboration XIII 2016), much remains unknown about reionization. The sources that provide the majority of UV photons for reionization are still debated (Couchman & Rees 1986; Haardt & Madau 1996; Ricotti 2002; Ricotti & Ostriker 2004; Choudhury, Haehnelt & Regan 2009; Dopita et al. 2011; Haardt & Madau 2012; Chardin et al. 2015; Madau & Haardt 2015). Likewise, the time it took to complete reionization is also still uncertain (Bolton & Haehnelt 2007; Bowman & Rogers 2010; Zahn et al. 2012). Answering these questions is important for understanding galaxy formation as reionization can have important effects on galaxy properties.

Unfortunately, RT, non-equilibrium chemistry simulations are computationally very expensive, and there must be a compromise between resolution and computational volume. Here we use cosmological simulations, including on-the-fly RT and detailed non-equilibrium chemistry, in an attempt towards self-consistently modelling the ISM at high redshift in order to better interpret current observations of ‘normal’ star-forming galaxies. The sacrifice we have to make in order to run a full cosmological box with on-the-fly RT is a moderate spatial resolution of ~ 100 pc. Our work improves on the previous studies that have attempted to model the far-IR emission in high-redshift galaxies in a number of ways. First, we explicitly model the relevant star formation and cooling processes that change the underlying structure of the galaxy. Secondly, we use on-the-fly RT, which is coupled to the molecular hydrogen and captures the spatially inhomogeneous flux in the Habing band, which is crucial for predicting the far-IR emission. Thirdly, we relate the internal properties of the galaxy to the global process of reionization. Finally, we attempt to connect large (few Mpc) and small (100 pc) scales and describe how the included physics affects galaxy formation and our understanding of gas properties in the objects that are currently being targeted by ALMA.

In Section 2, we describe our cosmological, radiation-hydrodynamics simulations and introduce a new variable-speed-of-light approximation (VSLA) that better models photon propagation in both low- and high-density regimes. In Section 3, we analyse these simulations on large, intermediate and small scales with a particular emphasis on the stellar, gas, radiation and molecular properties. In Section 4, we compare our simulations to previous work and describe how UV and Ly α emissions can be offset from [C II] and [O III]. Finally, in Section 6, we present our conclusions.

2 COSMOLOGICAL RT-HYDRO SIMULATIONS

2.1 Initial conditions

We simulate a $10 \text{ Mpc } h^{-1}$ box starting at $z = 150$ with a uniform grid of 256^3 dark matter particles ($m_{\text{dm}} = 6.51 \times 10^6 M_{\odot}$) and a similar number of gas cells. For the initial conditions, we use the software package MUSIC (Hahn & Abel 2011). Lagrangian

Table 1. Energy bins used to track the radiation. Energy bin 1 represents the Habing band. It does not affect the temperature or density state of the gas and is propagated so that we can calculate the C II, O I, O III and N II ionization states in order to predict their masses. The photons in these bins can be absorbed in the simulation depending on the metallicity of the cell.

| Bin | E_{\min} (eV) | E_{\max} (eV) | Main function |
|-----|--------------------|--------------------|---|
| 1 | 5.60 | 11.20 | C II, O I, O III and N II ionization states |
| 2 | 11.20 | 13.60 | H ₂ photodissociation |
| 3 | 13.60 | 15.20 | H I photoionization |
| 4 | 15.20 | 24.59 | H I and H ₂ photoionization |
| 5 | 24.59 | 54.42 | He I, H I and H ₂ photoionization |
| 6 | 54.42 | ∞ | He I, He II, H I and H ₂ photoionization |

perturbations and a local Lagrangian approximation are used on the base grid for dark matter particles and gas cells, respectively, so that the initial conditions are accurate to second order. Cosmological parameters reported by the Planck Collaboration are assumed ($h = 0.6731$, $\Omega_m = 0.315$, $\Omega_\Lambda = 0.685$, $\Omega_b = 0.049$, $\sigma_8 = 0.829$ and $n_s = 0.9655$; Planck Collaboration XIII 2016). CAMB (Lewis, Challinor & Lasenby 2000) was used to generate the transfer function with the relevant cosmology for these initial conditions, and the gas is initially assumed to be neutral with 76 per cent hydrogen and 24 per cent helium by mass.

2.2 Numerical implementation of chemistry, and radiation transport

We use the publicly available adaptive mesh refinement (AMR) code RAMSES (Teyssier 2002), and in particular, the RT version, RAMSES-RT (Rosdahl et al. 2013), to model the gravity, detailed hydrodynamics, non-equilibrium chemistry and RT in the cosmological box. We have made significant modifications to both the non-equilibrium chemistry and RT packages native to RAMSES-RT in order to better track the H₂ abundances as well as the propagation of ionizing radiation in low-density regimes and, in particular, the intergalactic medium (IGM).

2.2.1 H₂ chemistry, radiation and cooling

The standard version of RAMSES-RT is able to accurately follow the non-equilibrium chemistry and cooling for six species: H, H⁺, e[−], He, He⁺ and He⁺⁺. There is an additional patch to the code that was used to follow the formation and destruction of a seventh species, H₂, in the context of lower redshift galaxies than considered here and post-reionization (Tomassetti et al. 2015). This implementation is not coupled to the inhomogeneous 3D radiation field, which requires a complicated treatment of H₂ dissociating radiation in multiple bands. We use this patch as a starting point for our own implementation, which follows very closely the methods of radiation-coupled H₂ chemistry presented in Baczynski, Glover & Klessen (2015). Certain differences in the exact implementation arise due to the fact that Baczynski et al. (2015) use a ray-tracing method to follow the RT, while RAMSES-RT is a moment-based scheme that treats the radiation like a fluid.

We follow radiation in six different energy bins in the range [5.6 eV– ∞] listed in Table 1. The five highest energy bins are used to calculate rates for the following reactions:



and



The subscript for each γ determines which energy bin listed in Table 1 contributes to the reaction. As in Baczynski et al. (2015), we have created an extra bin with a lower limit of 15.2 eV, which is the ionization energy of H₂, and we assume that all H₂⁺ that is created by this reaction is immediately destroyed by dissociative recombination such that



For each reaction listed in equations (1)–(5), one must know the average atomic cross-section in order to calculate the photoionization rates. Furthermore, to self-consistently model the photoheating from these reactions, one must also store the energy-weighted cross-section (see e.g. Rosdahl et al. 2013). Because our photon bins represent averages across multiple frequencies, the cross-sections we use for each of the different species are calculated on-the-fly at each time-step and determined by the luminosity-weighted mean energy of the different spectral energy distributions (SEDs) of all of the sources in our box.¹ For each individual star particle, photons are injected into the host cell with an SED for a given metallicity and age (Bruzual & Charlot 2003). We adopt a Chabrier initial mass function (Chabrier 2003). The luminosity of the source is scaled to the mass of the star particle. In order to calculate the average and energy-weighted cross-section of each species in the energy bins greater than 13.6 eV, we use equations (B7) and (B8) as given in Rosdahl et al. (2013). These two equations integrate over the frequency-dependent cross-section for each species. For H, He and He⁺, we use the frequency-dependent cross-sections as given in Hui & Gnedin (1997). The frequency dependence of the H₂ cross-section at $E > 15.2$ eV is consistent with the piecewise fit to the analytical results of Liu & Shemansky (2012) and is listed in table 1 of Baczynski et al. (2015). The average and energy-weighted cross-sections for each species are updated for each of the relevant bins at every coarse time-step to account for the formation of new stars.

While H₂ photoionization at $E > 15.2$ eV is a continuum process, H₂ dissociation in the Lyman–Werner band is a line-driven process, and therefore we must treat the cross-section of H₂ in this band differently because self-shielding becomes important. Similar to Baczynski et al. (2015), we define an effective cross-section,

$$\sigma_{\text{H}_2, \text{LW}} = \sigma_{\text{H}_2, \text{LW}, \text{thin}} f_{\text{shd}}, \quad (7)$$

where $\sigma_{\text{H}_2, \text{LW}, \text{thin}} = D/F$, $D = 5.18 \times 10^{-11} \text{ s}^{-1}$ is the photodissociation rate in the optically thin limit (Röllig et al. 2007) and $F = 2.1 \times 10^7 \text{ s}^{-1} \text{ cm}^{-2}$ is the photon flux in the Lyman–Werner band in the interstellar radiation field (Draine & Bertoldi 1996). For f_{shd} , we follow Gnedin, Tassis & Kravtsov (2009) and assume

$$f_{\text{shd}} = \frac{1 - \omega_{\text{H}_2}}{(1 + x)^2} + \frac{\omega_{\text{H}_2}}{\sqrt{1 + x}} e^{-0.00085\sqrt{1+x}}, \quad (8)$$

where $\omega_{\text{H}_2} = 0.2$ and $x \equiv N_{\text{H}_2}/(5 \times 10^{14} \text{ cm}^{-2})$. N_{H_2} is the H₂ column density and we make the very simple approximation that

¹ Note that this discussion of cross-section is not applicable to H₂ photoionization by Lyman–Werner band photons, as listed in equation (1).

$N_{\text{H}_2} = n_{\text{H}_2} \Delta x$, where n_{H_2} is the number density of H_2 in the cell and Δx is the physical length of the grid cell in the simulation.

With cross-sections readily available, we can compute the photoionization and photodissociation rates for each of the different species, update their number densities and deplete the number of photons in each bin by the relevant amount, which fully couples the non-equilibrium chemistry to the RT.

In addition to the individual species as sinks for photons, in principle, there may also be dust present that can absorb photons across a wide energy range. We do not self-consistently track the formation of dust in the simulation, and in order to calculate the dust number density in each cell, we assume that the dust traces the metals. Similarly to Gnedin et al. (2009), we assume that the dust-to-gas ratio scales linearly with metallicity and take the dust cross-section to be $\sigma_{\text{dust}} = 4.0 \times 10^{-21} \text{ cm}^2$. This cross-section is much lower than that of the other species relevant to the more energetic radiation bins in our simulation, and the metallicity in our simulations is rather low compared to solar.² We are likely to resolve pockets of higher metallicity that would then have a higher dust content and possibly affect the absorption; therefore, photon absorption by dust is a very marginal effect. The dust number density, however, could become relevant for our lowest energy bin (which is the lower energy range of the Habing band) where none of the species in our simulation can act as absorbers. The flux in this band is relevant for photoelectric heating as well as for making predictions for $[\text{C II}]$, which may be observed by ALMA.

The total rate equation we use to track the abundance of H_2 in our simulation is

$$\frac{dx_{\text{H}_2}}{dt} = -C_{\text{coll}}x_{\text{H}_2} - (k_{\text{UV}} + k_{\text{LW}})x_{\text{H}_2} + (R_{\text{d}} + R_{\text{p}})n_{\text{H I}}, \quad (9)$$

where $n_{\text{H I}}$ is the number density of neutral hydrogen, $C_{\text{coll}} = \sum_{i=\text{e}^-, \text{H}, \text{He}, \text{H}_2} k_{\text{coll}, i} n_i$ and $k_{\text{coll}, i}$ are the collisional dissociation rates for each of the four listed species, as given in Glover & Abel (2008). k_{UV} and k_{LW} are the UV and Lyman–Werner photoionization and photodissociation rates discussed earlier and are calculated from the cross-section and number density of photons. Finally, R_{d} is the formation rate for H_2 on dust, which is relevant for metal-enriched gas, and R_{p} is the formation rate of H_2 via the H^- channel, which becomes relevant in the zero-metallicity limit. For R_{d} , we use the expression from Gnedin et al. (2009) so that

$$R_{\text{d}} = 3.5 \times 10^{-17} Z C_{\text{f}} \text{ cm}^3 \text{ s}^{-1}, \quad (10)$$

where Z is the metallicity of the gas and C_{f} is the clumping factor, which we set to 10, consistent with Gnedin et al. (2009). The expression for the primordial channel of H_2 formation is primarily due to the presence of H^- , and this reaction is taken from Glover et al. (2010). For this we must know the abundance of H^- , which we do not explicitly track in the simulation. In order to calculate this quantity, we assume an equilibrium abundance by solving for the equilibrium rate of the formation of H^- using reactions 1, 2, 5 and 13 in the appendix of Glover et al. (2010). Thus, we use the following equations for the primordial channel of H_2 formation:

$$k_1 n_{\text{H I}} n_{\text{e}} = k_2 n_{\text{H}^-} n_{\text{H I}} + k_5 n_{\text{H}^-} n_{\text{H II}} + k_{13} n_{\text{H}^-} n_{\text{e}} \quad (11)$$

and

$$\mathcal{R}_{\text{p}} \equiv \frac{k_1 k_2}{k_2 + k_5 x_{\text{H II}} + k_{13} x_{\text{e}}}. \quad (12)$$

In addition to being coupled with the radiation, the H_2 in our simulation is also coupled to the thermal state of the gas. In the low-metallicity regime, H_2 is the dominant coolant below the atomic cooling threshold and, in principle, can cool the gas to $\sim 100 \text{ K}$. We use the H_2 cooling rates from Hollenbach & McKee (1979), which is summed with all of the primordial cooling channels already present in *RAMSES-RT*, described in Rosdahl et al. (2013), in order to determine the net cooling rate. At slightly higher metallicities, metal line cooling becomes the dominant cooling channel below $\sim 10^4 \text{ K}$, and for these rates, we interpolate tables computed with *CLOUDY* (Ferland et al. 2013) with a Haardt & Madau (2012) UV background that were made for the Grackle chemistry and cooling library³ (Bryan et al. 2014; Kim et al. 2014). The values for the metal line cooling rates depend on redshift, density and temperature, and we scale the rates with the total metallicity of the cell. Because we do not resolve the smallest progenitors of $z = 6$ galaxies, we assume a metallicity floor with $Z_{\text{min}} = 10^{-3.5} Z_{\odot}$ at $z = 15$ (Wise et al. 2012).

Besides cooling, H_2 can also contribute to the volumetric heating rate of the gas, and we consider here two processes: heating due to UV photoionization for $E_{\gamma} > 15.2 \text{ eV}$ as well as heating from photodissociation and UV pumping in the Lyman–Werner band. The photoheating rate for H_2 due to photoionization from UV photons is treated like the other species where the excess photon energy above the ionization potential contributes to the heating term. In the Lyman–Werner band, we follow the method of Baczynski et al. (2015) to calculate the volumetric heating rate. For each photodissociation, an excess energy of $\sim 0.4 \text{ eV}$ is deposited into the gas as heat (Black & Dalgarno 1977). Not all absorptions of Lyman–Werner photons by H_2 lead to photodissociation. The H_2 can instead become vibrationally excited until it either fluoresces back down to the ground state or is collisionally deexcited, which can transfer heat to the gas. We calculate the heating rate due to this UV pumping following equations (44)–(48) in Baczynski et al. (2015), which combine the UV pumping rate of H_2 from Draine & Bertoldi (1996) with the energy released per UV pumping event from Burton, Hollenbach & Tielens (1990). Having outlined the thermal coupling of H_2 with the gas, we conclude our description of the H_2 implementation in *RAMSES-RT*.

Thus far, we have neglected the inclusion of our first photon group, Bin 1, which represents the lower energy range of the Habing band. We follow the radiation in this energy range for two specific reasons: (1) It becomes relevant to calculate the heating rate for photoelectric heating by dust. (2) In order to calculate the near-IR emission from C II , O I , N II and O III , one must know the energy density of photons in this band. The densities we probe in the simulations presented in this paper (especially the density at which we form stars) do not become high enough for photoelectric heating by dust to become relevant, and we have therefore neglected it in our simulation. The exact implementation and usage of photoelectric heating by dust is described in Kimm et al. (2017). The latter reason becomes of particular importance when we compare our simulations with ALMA observations. The origin of the $[\text{C II}]$ emission within these high-redshift galaxies is unknown and may originate in low-density neutral or ionized gas as well as from photodissociation regions. In order to better understand the physical properties of high-redshift galaxies that exhibit strong $[\text{C II}]$, $[\text{O I}]$, $[\text{N II}]$ or $[\text{O III}]$ emission, it is crucial that the Habing band is tracked self-consistently in the simulation. Note that the radiation in this

² At higher resolutions, when the ISM is better resolved.

³ <https://grackle.readthedocs.io/en/latest/>

bin does not physically affect the state of the gas in the simulation, but the inhomogeneous spatial distribution of this radiation is the quantity that we require to make predictions for ALMA.

2.2.2 The variable-speed-of-light approximation

Our aim is to run a cosmological simulation with on-the-fly RT to model galaxy formation in the high-redshift universe. *RAMSES-RT* uses an explicit solver for radiation transport, and therefore the RT time-step in the simulation is limited by the RT Courant condition such that

$$\Delta t_{\text{RT}} < \frac{\Delta x}{3c_{\text{sim}}}, \quad (13)$$

where c_{sim} is the speed of light used in the simulation. For high-resolution cosmological simulations, using the full speed of light can be prohibitively expensive because it is so much greater than the typical hydrodynamic velocities that govern the time-step for non-RT simulations. One can adopt the reduced-speed-of-light approximation (RSLA, Gnedin & Abel 2001) so that, for example, $c_{\text{sim}} = 0.01c$. This approximation is a good approximation in many regimes [see sections 4.2 and 4.3 of Rosdahl et al. (2013) for a lengthy discussion], as long as the propagation of the ionization front (I-front) is slower than c_{sim} . Rosdahl et al. (2013) clearly demonstrate that to track the I-fronts properly in the IGM, one must use the full speed of light, while in much higher density regions, such as the ISM, the RSLA is a good approximation. Furthermore, Bauer et al. (2015) have shown that adopting $c_{\text{sim}} < c$ can lead to delayed reionization for the same source model.

One possible option, which can also be used concurrently with the RSLA, is to sub-cycle the radiation time-step and use optimized boundary conditions (see Commerçon, Debout & Teyssier 2014) to make this compatible with adaptive time-stepping. The process of sub-cycling can be implemented in one of two ways: (1) A fixed value of the speed of light is chosen (i.e. $c_{\text{sim}} = 0.01c$) and the number of sub-cycles is determined by maximizing the number of RT time-steps that fit into a single hydro time-step (potentially up to a maximum number of sub-cycles, (e.g. Rosdahl et al. 2013)). (2) The number of sub-cycles is kept fixed and the speed of light is changed to perform that number of sub-cycles (e.g. Gnedin & Abel 2001). In the latter case, the speed of light in the simulation can, in principle, become very large for short hydro time-steps. The method of sub-cycling results in a speed-up because it reduces the total number of hydrodynamic time-steps, but the number of RT sub-cycles can still be prohibitively costly if the full speed of light is used. In this work, we apply this method, and for each hydro time-step, we perform a maximum of 10 RT sub-cycles.⁴

Alternatively, the problem of short time-steps has been circumvented by using uniform grids on GPUs (Aubert & Teyssier 2008, 2010). With this architecture, using the full speed of light is no longer prohibitively costly; however, the method is limited to a uniform grid to achieve such a large speed-up [although see Aubert, Deparis & Ocvirk (2015) for improvements allowing for AMR grids on the GPU]. In this paper, we present a new approach that we call the variable-speed-of-light approximation (VSLA) where the speed of light in the simulation changes depending on the level of refinement in the simulation, which thereby

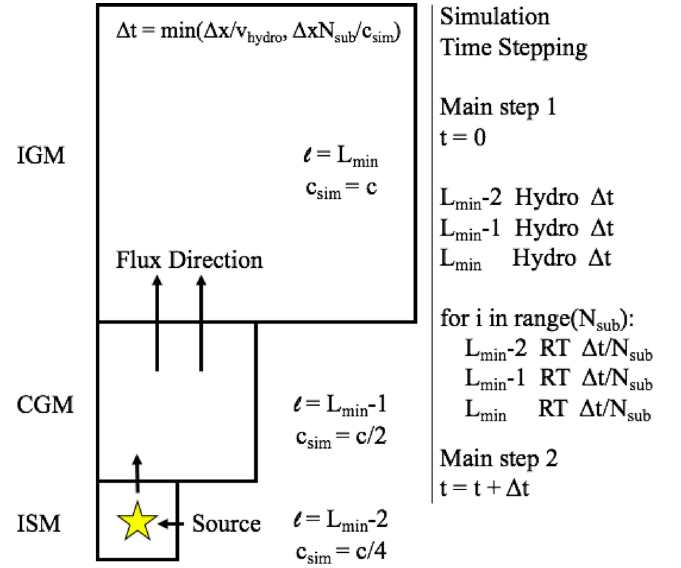


Figure 1. Schematic view of the variable-speed-of-light algorithm (VSLA). The speed of light used in the simulation increases by a factor of 2 at each higher level representing the change in density of the cell. In this three-level example, the most refined level represents the interstellar medium, the middle level represents the circumgalactic medium and the base grid represents the intergalactic medium. The source is located in the highest density region, and as the photons escape from the ISM, they increase in speed while the total flux is conserved. On the right-hand side, we show a pseudo-code of how the algorithm is implemented. The time-step is determined and all levels are synced to the smallest time-step in the simulation. First, the hydrodynamics is evolved for all levels, starting with the most refined. Once the hydro time-step is completed, the RT calculation is completed for the specified number of sub-cycles, looping through all levels at each sub-cycle. At each RT sub-cycle, cooling and photon absorption and emission are calculated.

Table 2. List of simulations. The first two columns denote the name of the simulation and whether it includes on-the-fly radiation from star particles. l_{max} indicates the maximum level of refinement, while Δx_{min} denotes the maximum physical resolution in pc that we maintain throughout the simulation. The box size and dark matter mass are the same for all simulations.

| Name | RT | l_{max} | Box size (cMpc h^{-1}) | M_{DM} (M_{\odot}) | Δx_{min} (pc) |
|--------|-----|------------------|------------------------------|------------------------------------|---------------------------------|
| L12 | No | 12 | 10 | 6.5×10^6 | 500 |
| L12-RT | Yes | 12 | 10 | 6.5×10^6 | 500 |
| L13 | No | 13 | 10 | 6.5×10^6 | 250 |
| L13-RT | Yes | 13 | 10 | 6.5×10^6 | 250 |
| L14 | No | 14 | 10 | 6.5×10^6 | 125 |
| L14-RT | Yes | 14 | 10 | 6.5×10^6 | 125 |

captures the fast- and slow-moving I-fronts at their proper speed. A schematic view of this algorithm is shown in Fig. 1, and a full description of the VSLA is presented in Appendix A. For all simulations in this work, we use the full speed of light on the base grid and divide by a factor of 2 on each subsequently refined level.

2.3 Resolution and refinement

We run a total of six different simulations at three different resolutions with and without on-the-fly RT. The simulations are listed in Table 2. All simulations are run until $z = 6$, except for the L14-RT

⁴ Note that in our case, we do not need to adopt any different boundary conditions because our VSLA routine uses the same hydro time-step, regardless of the level of refinement. Therefore, we can sub-cycle the RT through all levels at once after every coarse hydro time-step.

simulation, which is run until $z = 5.3$. The dark matter particle resolution is the same for all simulations, and we vary only the level the grid can refine to. The grid is allowed to adaptively refine during the course of the simulation when the density of a cell increases by a factor of eight times that of the previous level in either dark matter or baryons. We attempt to enforce a constant physical resolution throughout the simulation to keep the physical size of cells at the highest resolution as close to the values listed in Table 2 as possible. This means introducing further levels of refinement at predefined scalefactors such that when a increases by a factor of 2, the grid refines to a further level.

Throughout our analysis, we will focus mainly on the two runs at the highest resolution with and without stellar RT. We will return to the lower resolution simulations when discussing the properties of the ISM.

2.4 Star formation and feedback

Stars are formed in the simulation based on the Schmidt law (Schmidt 1959) such that $\dot{\rho}_* = \epsilon_{\text{ff}} \rho_{\text{gas}} / t_{\text{ff}}$, where $\dot{\rho}_*$, ϵ_{ff} and t_{ff} are the star formation rate density (SFRD), the efficiency of star formation per free-fall time and the free-fall time of the gas, respectively. We set a star formation density threshold of $n_{\text{H}} = 0.1 \text{ cm}^{-3}$ and require that $\rho_{\text{gas}} > 50 \bar{\rho}_{\text{b}}(z)$ to prevent star formation at very high redshifts. Stars are formed with an efficiency of $\epsilon = 1$ percent (Kennicutt 1998), and the number of star particles formed is drawn from a Poisson distribution, $P(N_*) = (\lambda^{N_*} / N_*!) \exp(-\lambda)$, where $\lambda \equiv \epsilon_{\text{ff}} (\rho \Delta x^3 / m_{*,\text{min}}) (\Delta t / t_{\text{ff}})$. For all simulations, regardless of resolution, we keep a fixed minimum mass for star particles so that $m_{*,\text{min}} = 7.66 \times 10^4 M_{\odot}$. We set a temperature threshold of $2 \times 10^4 \text{ K}$ so that stars cannot form in cells with a temperature higher than this value. Ionization of neutral hydrogen and neutral helium takes place at temperatures less than this value, and therefore the temperature increase in gas cells due to ionizing radiation is alone not enough to prevent star formation. This should allow the simulations that include stellar radiation sources to have similar SFRDs to the simulations without stellar radiation.

Each star particle is assumed to represent a simple stellar population with a Chabrier IMF (Chabrier 2005) with a minimum mass of $0.1 M_{\odot}$ and a maximum mass of $150 M_{\odot}$. After 10 Myr, the massive stars are assumed to explode as supernova (SN), and the star particles lose 31 per cent of their total mass,⁵ which is recycled back into the gas phase, and we assume that 5 per cent of the unenriched mass of the ejecta is composed of metals.⁶ For each SN, 10^{51} erg of thermal energy is injected into the gas. Since our simulations do not resolve the individual phases of the SN, they are likely to suffer from ‘overcooling’, and we therefore employ the ‘delayed-cooling’ model of Teyssier et al. (2013) in order to mitigate this effect. A delayed-cooling parameter is tracked as a passive scalar that decays exponentially as a function of time and cooling is shut off in these cells. We set the delay time-scale to 20 Myr. We have modified the standard feedback routine so that the feedback (i.e. mass, metals, thermal energy, delayed-cooling parameter) is

spread over the nearest 19 cells [see appendix A in Kimm & Cen (2014) for a visualization of the geometry].

3 RESULTS

3.1 Calibrating global properties

In order to fairly compare the different simulations, we must ensure that certain global properties are properly calibrated between the simulations. All six simulations have similar SFRDs as a function of time. This is achieved by keeping the stellar mass and dark matter resolution constant across all simulations while also using the same density threshold ($n_{\text{H}} \geq 0.1 \text{ cm}^{-3}$), temperature threshold ($T < 2 \times 10^4 \text{ K}$) and star formation efficiency. In the left-hand panel of Fig. 2, we compare the SFRDs from each of the six different simulations and see very good convergence. The SFRD in the simulations slightly overpredicts the observed SFRD, although the estimated SFRDs would drop if we consider only haloes with observable SFRs.⁷

In addition to the SFRDs, we attempt to keep the reionization histories constant across the three simulations that include stellar radiation. In the middle panel of Fig. 2, we plot the volume-filling factor of ionized hydrogen to demonstrate that all simulations agree very well. Our box is not completely ionized by $z = 6$, but note that our $10 \text{ Mpc } h^{-1}$ box is not a representative sample of the Universe and our simulations do not resolve some of the small sources with a halo mass of $\sim 10^8$ – $10^9 M_{\odot}$ that are likely very important for reionization (Kimm et al. 2017). Since SN and UV radiations are two of the dominant mechanisms that govern properties of the ISM, the important point is that we have controlled for them by ensuring that all simulations have similar SFRDs and all RT simulations have similar reionization histories.

Furthermore, it is crucial that the physical properties of individual haloes are converged in addition to the more global quantities in the simulations. The most efficient formation channel for H_2 is via the surface of dust grains, and therefore it is necessary that the metal masses of the haloes between the different resolutions are converged (as the dust mass is assumed to scale with metallicity). Note that we achieve the latter for free if we can converge the stellar mass as the total metallicity of a halo can be computed from the integrated SFR.

We use the AHF halo finder (Gill, Knebe & Gibson 2004; Knollmann & Knebe 2009) to locate haloes in our simulation. In the right-hand panel of Fig. 2, we plot the stellar mass and metal mass as a function of halo mass for the 50 most massive haloes in all six simulations at $z = 6$ that show very good agreement. Furthermore, the stellar masses fall within the 2σ contours of the stellar mass–halo mass relation predicted from abundance matching at $z = 6$ from Behroozi et al. (2013), suggesting that our galaxies have produced a reasonable amount of stars. The stellar mass–halo mass relation in our simulation has a slightly flatter slope than what is derived from abundance matching, which could indicate that our lower mass galaxies are forming stars more efficiently than expected from observations.

⁵ This value is calculated assuming a Chabrier IMF (Chabrier 2005) in the mass range 0.1 – $150 M_{\odot}$, where stars with $M > 6 M_{\odot}$ undergo SN, consistent with Vogelsberger et al. (2013) and Crain et al. (2015), assuming that stars with $6 < M(M_{\odot}) < 8$ explode as electron-capture SNe (Chiosi, Bertelli & Bressan 1992).

⁶ We do not include any additional metal enrichment from Type Ia SNe or asymptotic giant branch stars.

⁷ For instance, at $z = 6$ for the L14-RT simulation, taking haloes with $\text{SFRs} > 0.2(0.1) M_{\odot} \text{ yr}^{-1}$ would decrease the SFRD by $\sim 0.8(0.73)$ dex, which is closer to the observations. Note that this calculation is based on fewer than 50 haloes due to our small box. Furthermore, if we consider only those haloes with $M_{\text{halo}} > 10^{10} M_{\odot} h^{-1}$ at $z = 6$ for the SFRD, consistent with what is used for the stellar mass–halo mass relation in the right-hand panel of Fig. 2, we see a decrease in the global SFRD by 0.3 dex.




Figure 2. Left-hand panel: SFRD of all six simulations compared to observations (Zheng et al. 2012; Coe et al. 2013; Ellis et al. 2013; McLure et al. 2013; Oesch et al. 2013, 2014; Bouwens et al. 2014, 2015). Solid and dotted lines represent simulations with and without stellar radiation, respectively, while different colours represent different resolutions, as listed in the legend. Middle panel: volume-filling factor of ionized hydrogen for the simulations that include stellar radiation at three different resolutions. Right-hand panel: stellar (circles and triangles, the higher set of points) and metal (diamonds and squares, the lower set of points) masses in our simulations as a function of halo mass at $z = 6$ for all haloes with $M_{\text{halo}} > 10^{10} M_{\odot} h^{-1}$ at three different resolutions. Circles and diamonds correspond to simulations without RT, while triangles and squares represent simulations with RT. The grey band represents the stellar mass–halo mass relation inferred for $z = 6$ from Behroozi, Wechsler & Conroy (2013). Note the very good convergence in SFRD, ionization history, dark matter mass, stellar mass and metal mass.

Having confirmed that the global properties of our simulation are converged over the different resolutions, we now proceed to examine the physical properties of our simulations on large, intermediate and small scales.

3.2 Large scales: global reionization and thermodynamic properties of the IGM

3.2.1 Ionized bubble topology

Understanding the thermodynamic and temperature states of the IGM on large scales is key for understanding galaxy formation as well as interpreting future measurements of the 21-cm signal. On the largest scales (i.e. Mpc), radiation is the main mechanism that allows galaxies to affect the state of the IGM. In Fig. 3, we show H I maps in thin slices centred around the most massive galaxy as a function of redshift. This galaxy is one of the first haloes in the simulation to form stars and an H II region has formed around the galaxy as early as $z = 13$, which has a diameter across its longest axis of ~ 2 cMpc (the location of this galaxy in the box is circled in blue in the top middle panel of Fig. 3). The H II region is rather elongated and flattened, and it almost appears as if two separate H II regions are touching at the central location. This is due to a dense filament of gas that is feeding the galaxy perpendicular to the long axis of the bubble. The gas in this filament is much denser than the surrounding gas so that it has a much shorter recombination time. This prevents the ionized bubble from expanding efficiently across the vertical direction of the box, leading to this characteristic bipolar shape. The bipolar shape of the H II region is consistent with other studies that have looked at the escape of photons from a galaxy at a node in the cosmic web (Ciardi et al. 2001; Iliev et al. 2006; Abel, Wise & Bryan 2007; Wise & Abel 2008).

As the simulation evolves in time, more galaxies begin to form and emit ionizing radiation. A second galaxy forms to the bottom left-hand side of the most massive object, as can be seen in the bottom middle panel of Fig. 3, and an H II bubble centred around this galaxy appears as early as $z = 12$. By $z = 11$, these bubbles have

begun to merge, and by $z = 9$, a single bubble encapsulates both haloes. Note that the H II bubbles of these haloes are still expanding and have not yet reached the maximum extent of their Strömgen spheres. To get the bubble merging correct at early times, it is important to use the VSLA because it is well established that the RSLA solution will lag behind the true solution until the Strömgen radius is reached (Rosdahl et al. 2013). For our simulations here, we use a relatively small box size that includes many small sources. The I-fronts of these sources travel reasonably slowly at later times, making the RSLA potentially appropriate at later redshifts for hydrogen ionizing radiation bins. However, at early times, using the VSLA or the full speed of light is necessary to capture the early evolution of the Strömgen spheres.

Looking towards lower redshifts, by $z = 6$, nearly all of the distinct bubbles have merged and the topology of the ionized region is simply connected. There are still regions in the $z = 6$ slice that are not yet ionized and this is not surprising as we know from Fig. 2 that our box is not yet completely ionized by $z = 6$. By comparing the distribution of ionized gas at later redshifts, after the ionized bubbles have begun to merge, with the locations of the massive haloes in Fig. 3, we can see that the ionized regions correspond well with the locations of the most massive objects. The largest portion of residual neutral hydrogen at $z = 6$ sits on the left-hand side of the box and has a centre at $y \sim 0$. This corresponds to a void in our simulation. The highest density regions in our simulation, around the massive haloes, ionize first, before the voids. In the top left-hand panel of Fig. 3, we see that this region is completely ionized by $z = 5.5$. The reionization behaviour in our simulation is consistent with the ‘inside-out-middle’ scenario (Gnedin 2000; Finlator et al. 2009), whereby the remaining high density gas in the post-overlap phase (i.e. filaments) are the last regions to become ionized. To demonstrate this, in Fig. 6, we plot the ratio of the mass-weighted volume-filling factor of H II to the volume-weighted filling factor of H II . At high redshift, before bubble overlap, this ratio remains greater than 1 as the high-density regions around galaxies are ionized first. As the UV radiation propagates into the voids, this ratio decreases as more of the volume




Figure 3. Spatial distribution of the H I fraction of the gas as a function of redshift for the L14-RT simulation. The z -coordinate of the slice is centred on the most massive halo in the box. The 10 most massive haloes are circled in the top centre panel, and the blue circle indicates the position of the most massive halo in the simulation. Reionization begins around the most massive halo, and by $z = 5.5$ the Universe is completely ionized.

is ionized. The ratio drops below 1 when the voids are completely ionized and slowly approaches unity as the filaments become more ionized.

3.2.2 $H\text{I}$ photoionization rate

In the left-hand panel of Fig. 4, we show the spatial distribution of the photoionization rate for our highest resolution simulation at $z = 6$. Unsurprisingly, the spatial distribution of the photoionization rate corresponds well to that of the distribution of ionized gas. The gas closest to the most luminous ionizing sources has the highest photoionization rate. The photoionization rate inside of the bubbles is nevertheless relatively homogeneous.

In the middle panel of Fig. 4, we plot a 2D mass-weighted histogram of $\Gamma_{\text{H I}}$ versus n_{H} . The majority of the low-density gas, $n_{\text{H}} < 10^{-4} \text{ cm}^{-3}$, has $\Gamma_{\text{H I}}$ just below 10^{-13} s^{-1} . This is the isotropic

photoionization rate we see permeating throughout the ionized bubbles in the left-hand panel of Fig. 4. There is a plume of gas at these densities with lower $\Gamma_{\text{H I}}$ due to the fact the simulation is not completely ionized by $z = 6$. The residual neutral hydrogen absorbs the photons in these regions, which strongly reduces $\Gamma_{\text{H I}}$.

At higher densities, there is a bifurcation at $n_{\text{H}} \sim 10^{-3} - 10^{-2} \text{ cm}^{-3}$ with one arm pointing to high values and the other towards low values of $\Gamma_{\text{H I}}$. The lower arm at high densities is due to self-shielding in dense regions. This gas can either be pristine material that is fed into the centres of the haloes along dense filaments or metal-enriched gas that has been ejected from the central regions of the galaxy from SN feedback and is now cooling and falling back. For the same temperature, this gas will have much shorter cooling time-scales than the gas in the cold flows due to the presence of the metals. Due to the self-shielding from neutral gas, the radial profiles will not exactly follow r^{-2} . For example, the left-hand panel of Fig. 4 shows




Figure 4. Left-hand panel: spatial distribution of the H I photoionization rate at $z = 6$. The z -coordinate is centred on the most massive halo in the simulation. The photoionization rate is highest around the most massive halo and is somewhat lower in the voids. Middle panel: mass-weighted 2D histogram of H I photoionization rate versus number density of the gas. Note that the hydrogen in the simulation is not completely ionized, and the photoionization rate extends towards low Γ at $n_H \sim 10^{-4} \text{ cm}^{-3}$. At high densities, we see two competing tracks. The gas begins to self-shield at densities slightly lower than $n_H \sim 10^{-2} \text{ cm}^{-3}$, where the photoionization rate branches off towards higher values due to the presence of sources within the cells. Right-hand panel: evolution of the volume-weighted H I photoionization rate in the ionized regions as a function of redshift for the three simulations that include radiation.




Figure 5. Left-hand panel: map of J_{LW} averaged along the z -axis in a small region, 1/10th of simulation box, for the L14-RT simulation. In between haloes, J_{LW} is very homogeneous, indicative that the Universe is optically thin to Lyman–Werner radiation. Middle panel: gas density integrated along the z -axis of the same sub-region for the L14-RT simulation. The gas overdensities correspond well with the sources of Lyman–Werner photons. Right-hand panel: radial profiles of J_{LW} for the most massive halo in the box at $z = 6$ for simulations that include stellar radiation. J_{LW} falls off as r^{-2} as expected for a source in an optically thin region. Note the very good convergence in this behaviour between the three simulations.

a local minimum in $\Gamma_{H I}$ just above the location of the most massive halo in the box. This is due to a dense filament of gas that is feeding the galaxy in this direction, which leads to large column densities along the line of sight and creates a shadow in the photoionization rate.

The right-hand panel of Fig. 4 shows the volume-weighted evolution of $\Gamma_{H I}$ in ionized regions (defined where $x_{H II} > 0.5$) as a function of redshift for the three simulations that include stellar radiation. At high redshift, the photoionization rate is very large. This is because the ionized regions in the box are close to the photon sources and thus $\Gamma_{H I}$ remains high (see Fig. 3 for the locations of the ionized regions). As the ionized bubbles start to expand, $\Gamma_{H I}$ in the ionized regions begins to decrease as the UV photons reach lower density gas and are spread over a larger volume. The photoionization rate decreases until $z \sim 7$ when the ionized bubbles begin to overlap (see Fig. 3). Once overlap occurs, the mean free path increases and $\Gamma_{H I}$ begins to increase again.

3.2.3 Intensity in the Lyman–Werner band

In addition to H I ionizing radiation, we study the evolution of the Lyman–Werner background, which is important for setting the

intergalactic H_2 abundance and inhibiting gas cooling in low-metallicity systems. While the simulation is not yet completely optically thin to hydrogen-ionizing UV radiation at $z = 6$, it is completely optically thin to photons in the Lyman–Werner band. The only two absorbers of Lyman–Werner photons in the box are H_2 and dust. Nevertheless, as we discuss later, our simulations destroy H_2 more efficiently than one might expect, possibly due to the moderate resolution of our simulations. The metallicity in most of the IGM remains at the level $Z = 10^{-3.5} Z_\odot$ to which it was set at $z = 15$. The left-hand panel of Fig. 5 shows the spatial distribution of the intensity of the radiation in the Lyman–Werner band for an ~ 200 kpc region (1/10th the physical box size) in the simulation at $z = 6$, while the middle panel shows the surface density of the gas. The Lyman–Werner intensity (J_{LW}) is fairly constant in the IGM, and the peaks in J_{LW} are clearly associated with collapsed structures with correspondingly high surface densities. For gas below the cosmic mean baryonic density, $J_{LW} \sim 9.5 \times 10^{-21} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1} \text{ sr}^{-1}$. The amplitude slowly decreases with increasing redshift as it follows the integrated SFR.

The right-hand panel of Fig. 5 shows the radial profile of J_{LW} for the most massive galaxy in our simulation at $z = 6$ for three different resolutions. In the central regions, J_{LW} is elevated a few




Figure 6. Ratio of the mass-weighted volume-filling factor of H II, $Q_{\text{HII},M}$, to the volume-weighted filling factor of H II, $Q_{\text{HII},V}$, as a function of redshift for the L14-RT simulation. This ratio remains above unity as the highest density regions around the galaxies are ionized first. The ratio continues to decrease as the voids become ionized and finally drops below 1 when the only remaining neutral gas resides in the filaments and inside galaxies.

orders of magnitude above the background due to the presence of young stars in the galaxy. For the highest and lowest resolutions, the profile is relatively flat out to a few kpc before decreasing, while for our middle resolution, a flat profile is not seen. The shape of this profile is entirely determined by the distribution of young stars within the galaxy, which is reasonably converged between the three different resolutions; however, the middle-resolution simulation in this snapshot has a slightly lower SFR compared to the other two simulations.

Since the simulation box is optically thin to Lyman–Werner radiation, any version of RSLA where the photon velocity is decreased significantly will lead to long lags in changes of the Lyman–Werner intensity as the fronts are expanding. These fronts, in principle, will move at a velocity close to the full speed of light. A strong burst of star formation in a massive galaxy will raise the Lyman–Werner background locally well above the mean in the IGM. This, then, should propagate to nearby haloes at the full speed of light and affect the formation of H_2 . It is thus important to use the full speed of light or the VSLA in order to transport these feedback effects at the correct speed.

3.2.4 Thermal state with and without RT

Including the radiation in the simulation has significant effects on the temperature of the gas. In Fig. 7, we show the density–temperature phase-space diagrams for the two highest resolution simulations at $z = 6$. It is clear that the temperature–density relation is very different with the inclusion of stellar radiation. Most notably, at low densities [$-5 < \log_{10}(n_{\text{H}}) < -4$], much of the gas in the L14-RT simulation sits at $T \sim 10^4$ K due to the ionization of hydrogen, while for the L14-No RT simulation, the majority of the gas has $T < 10^2$ K. Without ionizing radiation, the gas will simply adiabatically cool as it expands and remain at low temperatures.

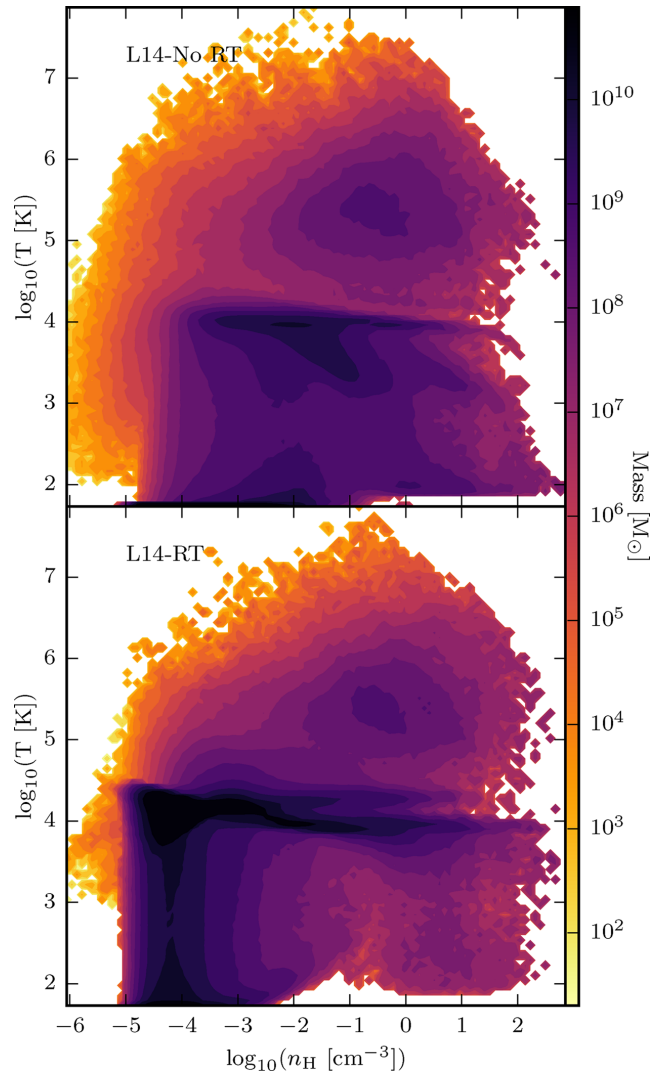


Figure 7. 2D histograms of density versus temperature for our two highest resolution simulations at $z = 6$. The upper panel is for the L14-No RT simulation, while the lower panel represents the L14-RT simulation. Very little low-density ($n_{\text{H}} < 10^{-4} \text{ cm}^{-3}$) gas has temperature $T \geq 10^4$ K in the L14-No RT simulation because all of the gas in the IGM is neutral. At higher densities, we see much more similar behaviour between the two simulations with a plume of gas at $10^5 < T(\text{K}) < 10^6$ due to SN feedback, and gas at $T < 10^4$ K due to radiative cooling.

Note that the evolution of gas in the density–temperature phase space in the L14-RT simulation is rather different from a simulation run with a spatially uniform metagalactic UV background such as those prescribed in Haardt & Madau (2012) and Faucher-Giguère et al. (2009). With a spatially uniform background, the low-density regions will be ionized first, and thus once the UV background turns on, the temperature of the lowest density gas will jump to $\sim 10^4$ K. In our L14-RT simulation, the ionized bubbles begin forming in the highest density regions around the haloes, while the voids are ionized last.

3.3 Intermediate scales – feeding the circumgalactic medium from the cosmic web

Characterizing the ionization state and metallicity of the circumgalactic gas that is feeding the galaxies will help in interpreting the

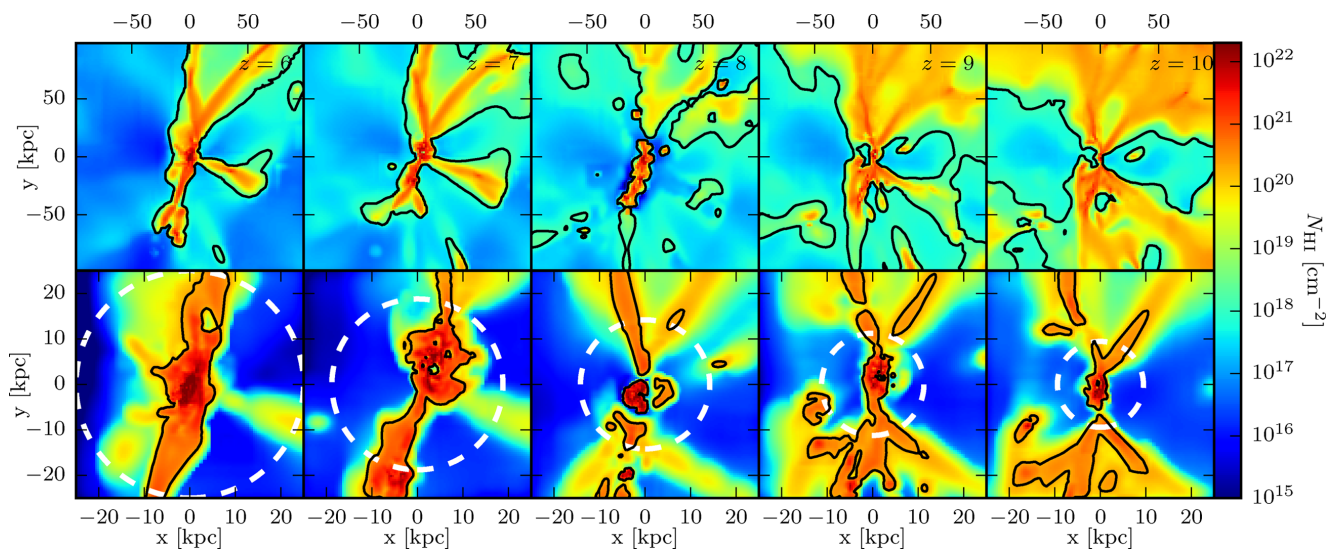


Figure 8. Time series from $z = 10$ to 6 of the column density of neutral hydrogen for the most massive halo in the simulation. Dashed white circles in the bottom panel indicate the extent of the virial radius of the halo. Black contours in the top panel indicate a column density of 10^{17} cm^{-2} to highlight LLSs. In the bottom panel, the black contours mark a column density of $2 \times 10^{20} \text{ cm}^{-2}$ to highlight regions causing DLAs. LLSs tend to map out the filaments between haloes and get wider with increasing redshift, while DLAs are present only within the virial radii of haloes. Both the top and bottom panels are integrated along 50 physical kpc.

systems that appear in quasar absorption spectra. The top panel of Fig. 8 shows a time series of a 200 kpc region of the surface density of neutral hydrogen around the most massive halo from $z = 10$ to 6 . At $z = 6$, the halo is fed by three dense, mostly neutral filaments that move with the evolving galaxy. Neutral gas moving along these filaments penetrates deep into the centre of the halo before merging with the central galaxy. The filament feeding the galaxy from the bottom is the most dynamic as a major merger is occurring along this direction. At $z = 7$, a large overdensity of neutral gas can be seen just outside the virial radius of the central halo. By $z \sim 6.5$, the satellite galaxy has fully entered the virial radius and the merger is complete by $z = 6.0$.

Contours in the top panel of Fig. 8 highlight $N_{\text{H I}} = 10^{17} \text{ cm}^{-2}$, where the gas is expected to self-shield from external ionizing radiation. Lines of sight passing inside of these contours will result in Lyman-limit absorption systems (LLSs). This column density threshold encapsulates most of the filamentary structure as well as the central galaxy, consistent with Fumagalli et al. (2011), who find that the streams feeding galaxies remain optically thick in hydrodynamics simulations post-processed with radiation transfer. The fraction of the surface area inside the virial radius of the central galaxy that is covered by neutral gas at this surface density changes quite substantially between $z = 10$ and 6 . The covering fraction in our simulations at $z = 6$ is higher than the 25 per cent expected from Fumagalli et al. (2011) at $z = 3$ and much closer to the expectations for simulated massive haloes at $z = 4$ from Rahmati et al. (2015). The covering fraction is expected to increase with redshift (e.g. Cen 2012; Rahmati et al. 2015). At $z = 6$, more than 50 per cent of the cross-section enclosed by the virial radius is covered by LLSs, while at $z = 8$, this fraction is significantly smaller. At this redshift, there is a strong deficit in neutral gas, suggesting that a strong burst in star formation has occurred within the past 10 Myr. Fig. 8 shows that the width of the filaments that are above an H I column density of $N_{\text{H I}} = 10^{17} \text{ cm}^{-2}$ strongly increases towards high redshift as the photoionization rate in the IGM is decreasing (see Fig. 2). By $z = 10$,

the majority of the gas in a 200 physical kpc cube surrounding this galaxy would result in an LLS. Note however, that at $z \gtrsim 8$ gas is still sufficiently neutral that LLSs are very extended and not well defined.

Fig. 9 shows the H I column density in the 200 kpc regions around the next five most massive haloes in the simulation. Some of the haloes are only being fed by two filaments, while other haloes are fed by three or more. Some of the filaments are fragmented and others are clear remnants of mergers. The structure strongly varies between the different haloes.

In the bottom rows of Figs 8 and 9, we zoom in closer to the haloes and highlight the $N_{\text{H I}} = 2 \times 10^{20} \text{ cm}^{-2}$ threshold for damped Ly α systems (DLAs). The regions predicted to cause DLAs can be fragmented and generally have very asymmetric shapes with many of these systems coming in along the filaments and potentially extending out to the virial radius, consistent with Fumagalli et al. (2011) and Cen (2012). There are holes in the distribution of neutral gas as a result of photoionization from young stars and SNe blowing out the gas. Star formation occurs preferentially in the highest density regions at the centres of haloes, and this energetic feedback is very efficient at disrupting the H I .

We turn now to the spatial distribution of metals in our simulations, which can be probed by associated metal absorption lines in QSO spectra during the epoch of reionization at $z > 6$ (Ryan-Weber et al. 2009; Becker et al. 2011; Simcoe et al. 2011; D’Odorico et al. 2013). In Fig. 10, we show the density-weighted spatial distribution of the metallicity surrounding the five most massive haloes in the L14-RT and L14-No RT simulations at $z = 6$. The metallicity inside the enriched bubbles is relatively uniform in the inner regions and remains rather low compared to solar. For the most massive halo, the metallicity reaches a maximum of a few $\times 10^{-1} Z_{\odot}$, and is somewhat smaller for the less massive haloes. Interestingly, very little of the metals actually penetrates out of the virial radius and into the IGM, and this is similar to the results of Ma et al. (2016) at $z = 0$ where their haloes of this mass in their simulations




Figure 9. Column density of neutral hydrogen for the next five most massive haloes in the simulation at $z = 6$. Dashed white circles in the bottom panel indicate the extent of the virial radii of the haloes. Black contours in the top panel indicate a column density of 10^{17} cm^{-2} to highlight LLSs. In the bottom panel, the black contours mark a column density of $2 \times 10^{20} \text{ cm}^{-2}$ to highlight regions causing DLAs. The numbers listed in the top row represent $\log_{10}(M_{\text{DM}} [M_{\odot}])$ for that object. The morphologies of DLAs and LLSs are very dependent on the system where they are located. Both the top and bottom panels are integrated along 50 physical kpc.




Figure 10. Spatial distribution of metallicity for the 100 kpc region centred on the five most massive haloes in the L14-No RT (top) and L14-RT (bottom) simulations at $z = 6$. The integration along the line of sight (50 physical kpc) has been weighted by density. The most massive halo is shown in the left-hand panels, and the halo mass decreases in each subsequent panel. The virial radius of each galaxy is indicated as a white circle. The white contours in the bottom row outline neutral gas with a column density $> 10^{17} \text{ cm}^{-2}$. Most of the metals tend to remain inside the virial radius of the haloes. The locations where the filaments have been enriched above the metallicity floor generally correspond to a smaller galaxy that has formed along the filament. However, some filaments remain unenriched.

have retained nearly 100 per cent of their metals. However, much of the circumgalactic medium CGM (in cross-section) is enriched to values slightly above the metallicity floor. The gas feeding the galaxies along the filaments appears to be close to primordial metallicity. Smaller haloes that have formed stars and are accreting along the filament lead to cases where the metallicity of the filament is raised above the metallicity floor. Note again that we do not resolve the small mini-haloes that may pre-enrich the IGM at earlier epochs (Wise et al. 2012) so the exact metallicity of the filaments is not reliably predicted by our simulations. In the simulations with and without stellar radiation, the distribution of metals and stellar

mass are similar, however, which is consistent with the expectations of Okamoto, Gao & Theuns (2008), where it was shown that the galaxy filtering mass is likely smaller than the resolution limit of dark matter haloes in our simulations.

3.4 Galactic scales – from the circumgalactic medium to stars and the ISM

We will now focus our attention on the most massive galaxy in the simulation, which is the best resolved object, and discuss the properties of its ISM.




Figure 11. Distribution of young (white), middle-aged (blue) and old (white contours) stars for the three different simulations that include stellar radiation. The underlying coloured map shows the neutral hydrogen column density. The young stars tend to form in clumps, while middle-aged stars have spread out more. The old stellar population forms a smooth distribution coincident with the dense gas.




Figure 12. H_2 mass (black), the mass of young stars (red) and the mass of middle-aged stars (blue) as a function of redshift in the three different simulations that include stellar RT. The formation of H_2 is synched with the formation of young stars as both tend to form in dense neutral gas. Because of strong radiation feedback, the fluctuations in H_2 mass are anticorrelated with those in the mass of middle-aged stars.

3.4.1 Spatial distribution of young, middle-aged and old stars

At high redshifts, galaxies are predominantly identified through their UV emission (e.g. Bouwens et al. 2015), and here we study the different stellar populations that give rise to the emission at different wavelengths. We identify three specific groups of stars: young stars (age < 10 Myr), middle-aged stars ($10 < \text{age} < 30$ Myr) and old stars (age > 30 Myr). The first group represents the stars that have yet to go SN, the second group represents the stars that have gone SN recently and the remnants of the explosions are keeping the gas at a high temperature since we delay cooling for 20 Myr, and the final group represents the passive stellar population that are remnants of previous episodes of star formation. In Fig. 11, we show the surface density of these three populations at $z = 6$ for the most massive halo for the three different resolution simulations that include stellar radiation. While all three populations are centred around the same location, the young stellar population exhibits a clumpy distribution, the middle-aged population is slightly more spread out, while the spatial distribution of the old population is very smooth.

Once the stars have gone SN, the density of the host gas cells decreases as matter is expelled out of the central regions of the galaxy. This decreases the local potential well, allowing the middle-

aged stars to spread out and respond to the overall gravitational potential well of the galaxy.

3.4.2 Time evolution of star formation

Observations of high-redshift galaxies provide only an instantaneous snapshot of the current state of a galaxy. Understanding these observations therefore requires understanding the duty cycle of star formation in these systems, which puts limits on a galaxy's observability at high redshift. In Fig. 12, we show how the mass in young and middle-aged stellar populations changes as a function of redshift between $z = 8$ and 6 for the three simulations with different resolutions that include stellar radiation. For the two lowest resolutions, there is a clear sinusoidal pattern in the mass of both young and middle-aged stars where the peaks have a separation of $\Delta z \approx 0.25$, corresponding to a time-scale of ~ 30 – 50 Myr in this redshift range. For the young stars, the time-scale between the first star formation and when the last SN remnants allow the gas to cool again is ~ 40 Myr,⁸ which corresponds well with the time-scale of this sinusoidal pattern. The peaks in middle-aged stellar mass are

⁸ 10 Myr for the time-scale between first stars forming plus 20 Myr for delayed cooling plus an additional 10 Myr because young stars will form

somewhat offset from the peaks in young stellar mass. This is simply because the middle-aged stellar mass increases as the young stars age into this category. SN feedback from the young stars will prevent more stars from forming and thus decrease the young stellar mass while the middle-aged stellar mass increases.

In the highest resolution simulation, much of the prominent sinusoidal pattern in the young and middle-aged stellar mass is not evident. This is likely due to the fact that as the resolution increases, the gas can reach higher densities and fragment more. In the two lower resolution simulations, the distribution of the gas is dominated by only a few large clumps. Once SNe explode in these clumps, it is as if the entire galaxy becomes regulated by these episodic star formation events. As more and more clumps form at higher resolution, disrupting them all simultaneously becomes difficult, and thus the pattern becomes washed out. The right-hand panel of Fig. 12 shows that the similar features we see at the two lower resolutions are present at the highest resolution as well. In particular, there is a clear peak in young stellar mass at $z \sim 7.3$. As the young stellar mass decreases, the middle-aged stellar mass increases as expected. There are obvious other features such as the wide peak in young stellar mass spanning $z = 6.8$ – 6.2 , which is similarly followed by a delay in the middle-aged stellar mass.

These patterns can be recognized in the spatial distribution in Fig. 13, where we show surface density maps of the young stellar mass, middle-aged stellar mass and old stellar mass as a function of redshift from $z = 6.8$ to 6.0 . A prime example of this is the large hole in the distribution of young stars at $z = 6.8$ due to an earlier burst of star formation. A trough in the young stellar mass can be seen at this redshift in the right-hand panel of Fig. 12. By $z = 6.6$, this hole is once again filled with young stars, and at this redshift, there is a peak in the young stellar mass in the right-hand panel of Fig. 12. It is evident from this figure that both the mass and distribution of young stars in the galaxy, which are the primary sources of UV and Ly α radiation, are dynamic both in spatial distribution and in mass over short time-scales. Looking at the old stars, we do not see much of a change in the distribution as a function of redshift. In this redshift range, the total stellar mass of the galaxy increases from 4.3×10^8 to $8 \times 10^8 M_\odot$. Given the large dynamic range of the colour bar in Fig. 13, a small change will not be visible in the distribution, and therefore the old stellar mass looks relatively constant.

3.4.3 Spatial variations in the intensity of radiation

Star formation will leave a unique signature on the radiation field within the galaxy as young stars are the primary emitters of ionizing photons, both at $E > 13.6$ eV and in the Lyman–Werner band. In Fig. 13, we characterize the spatially inhomogeneous radiation field that is key for predicting the ionization states of metals whose emission can be observed by ALMA. We show $\Gamma_{\text{H I}}$ and the intensity in the Lyman–Werner band (J_{LW}) as a function of redshift between $z = 6.8$ and 6.0 . Comparing the two quantities, one can see a stark contrast. At $z = 6$, $\Gamma_{\text{H I}}$ has a strong peak in the centre of the halo corresponding to the highest density in young stars. This quickly fades by a few orders of magnitude even at very small radii, although small secondary peaks in $\Gamma_{\text{H I}}$ can be seen throughout the image. These are due to secondary star formation sites that have less mass in young stars compared to the central regions.

By contrast, the intensity of the radiation in the Lyman–Werner band is extremely smooth at $z = 6$. There is a clear central peak in the radiation field, and this intensity falls off as r^{-2} at larger radii (see the right-hand panel of Fig. 5). This difference arises because of the nature of the absorbers in each band. Photons with $E > 13.6$ eV are readily absorbed by neutral hydrogen. The mass- (volume-)weighted mean free path for hydrogen ionizing photons is ~ 500 kpc (~ 11 Mpc) inside the virial radius of the halo, while for the Lyman–Werner band, we find a significantly longer mass- (volume-)weighted mean free path of ~ 700 Mpc (~ 1.2 Gpc).

3.4.4 Time evolution of the thermal and ionization states of the ISM

The radiative and SN feedback significantly affects the ionization state and temperature of the gas, and here we show how these properties correlate with star formation. The top panel of Fig. 13 shows the time evolution of temperature. A large ionization bubble exists around this galaxy (see Fig. 3) so it is unsurprising to see gas at $T \sim 10^4$ K surrounding the galaxy. The temperature changes from 10^3 to 10^6 K, and the peaks in temperature correspond to peaks in middle-aged stellar mass (i.e. compare middle-aged star surface mass density to temperature in the $z = 6.3$ and 6.7 columns of Fig. 13).

The second and third rows of Fig. 13 show the surface density of ionized and neutral gas as a function of redshift. Towards the outskirts of the galaxy, there is significantly more ionized gas compared to neutral gas. This gas is diffuse and the recombination time-scale is large, and thus the gas is maintained in an ionized state (hence why the temperature is $\sim 10^4$ K in these regions). A minimum in the surface density of ionized gas can be seen between $z = 6.1$ and 6.3 , which is associated with the decline in young stellar mass (see Fig. 12). The intensity of ionizing radiation is sensitive to the presence of young stars (see Section 3.4.3) and, in particular, stars with age $\lesssim 3$ Myr as the emission of ionizing photons drops quickly for stars past this age (Bruzual & Charlot 2003).

In the neutral gas, we see a markedly different structure compared to that in the ionized gas. The dynamic range is significantly higher and many more strong spatial features are present. The neutral gas distribution changes rapidly. We see many examples where holes appear in the distribution due to star formation and SN feedback.

In Table 3, we list the physical properties for the 10 most massive haloes in the L14-RT simulation at $z = 6$. In 9 of 10 haloes, the H II mass is dominant over the H I mass, indicating that these haloes are sufficiently ionizing their surroundings. The metal mass is spread fairly proportionally over the neutral and ionized gas, consistent with the mass of these two quantities (i.e. for the most massive halo, 32 per cent of the total metal mass is associated with neutral gas, while 34 per cent of the total gas is in the neutral phase). Most haloes have retained their cosmic baryon fraction of gas, consistent with the expectations from Okamoto et al. (2008) and Gnedin & Kaurov (2014) for the halo masses considered here. The feedback from radiation and photoionization does not seem to be strongly affecting the accretion of gas on to these massive haloes. Most of the gas is being fed through cold, neutral, dense filaments (see Fig. 9), and the outflows are not preventing fresh gas from penetrating down to the centres of the haloes (Kimm et al. 2015).

For observations, it is imperative to know the temperature states of each of the different gas components inside the halo. In Fig. 14, we show the probability distribution function (PDF) of the temperatures for H I, H II and metals inside the virial radius of the most massive halo in the L14-RT simulation at $z = 6$. For H II, we see a

during the 10 Myr before the first stars have undergone SN and thus the final SN explode 10 Myr after the first.




Figure 13. Time series of various physical properties of the most massive galaxy in the L14-RT simulation shown from $z = 6.0$ to 6.8 in intervals of $\Delta z = 0.1$ starting from the left-hand side. Each image is 20 proper kpc in width. From the top to bottom, the quantities are temperature (slice), H II column density, H I column density, metal surface mass density, H_2 surface density, $\Gamma_{\text{H I}}$ (slice), J_{LW} (slice), young stellar mass surface density, middle-age stellar mass surface density and old stellar mass surface density. Properties shown as slices represent a small section of the simulation box centred in the middle of the halo.

double-peaked profile: the lower temperature peak is due to photoheating, while the higher temperature peak is from SN feedback. Hydrogen efficiently recombines only at $T < 10^4$ K, and thus we see H I only at temperatures lower than this value. The metals are very evenly spread across the galaxy, and for this reason, we see that the temperature PDF of this component is a combination of both the H I and H II distributions.

3.4.5 The evolution of H_2 with and without RT

Of particular interest is H_2 because it is well known to correlate with star formation (Kennicutt 1998; Bigiel et al. 2008). Furthermore, in

low-metallicity environments, as is the case in many regions of our simulation box, H_2 becomes the primary coolant below 10^4 K and thus alters the thermodynamic state of the gas.

In Fig. 15, for each of the six simulations, we plot the total mass of H_2 contained within the virial radius at $z = 6$ against the SFR (measured by averaging over the past 10 Myr) for all haloes with $M_{\text{halo}} > 10^{10} M_{\odot} h^{-1}$. Without RT, the H_2 masses form a very tight, positive and steep correlation with the SFRs of the haloes. For the best resolved halo (the three circles at the top right-hand panel of Fig. 15), we see that the instantaneous SFR and H_2 masses are converged within a factor of ~ 2 – 3 at $z = 6$. When RT is introduced into the simulations, the total mass of H_2 decreases

Table 3. Properties of the 10 most massive haloes in the L14-RT simulation at $z = 6$. M_{DM} , M_* and M_{gas} represent the dark matter mass, stellar mass and gas mass in the haloes, respectively. $M_{\text{Z,H I}}$ and $M_{\text{Z,H II}}$ are the metal masses associated with neutral and ionized gas. $M_{\text{H I}}$, $M_{\text{H II}}$ and M_{H_2} are the masses in neutral hydrogen, ionized hydrogen and molecular hydrogen. SFR is the current star formation rate of the halo.

| Halo | M_{DM} $\log_{10}(M_{\odot})$ | M_* $\log_{10}(M_{\odot})$ | M_{gas} $\log_{10}(M_{\odot})$ | $M_{\text{Z,H I}}$ $\log_{10}(M_{\odot})$ | $M_{\text{Z,H II}}$ $\log_{10}(M_{\odot})$ | $M_{\text{H I}}$ $\log_{10}(M_{\odot})$ | $M_{\text{H II}}$ $\log_{10}(M_{\odot})$ | M_{H_2} $\log_{10}(M_{\odot})$ | SFR $M_{\odot} \text{ yr}^{-1}$ |
|------|---|---------------------------------|--|--|---|--|---|--|------------------------------------|
| 1 | 11.17 | 9.08 | 10.43 | 6.73 | 7.06 | 9.85 | 10.13 | 6.78 | 4.91 |
| 2 | 10.83 | 8.51 | 10.08 | 6.25 | 6.45 | 9.56 | 9.74 | 5.76 | 2.12 |
| 3 | 10.66 | 8.20 | 9.91 | 5.90 | 6.20 | 9.34 | 9.61 | 3.38 | 0.72 |
| 4 | 10.64 | 8.37 | 9.92 | 5.97 | 6.42 | 9.29 | 9.64 | 3.54 | 0.69 |
| 5 | 10.59 | 8.26 | 9.86 | 5.78 | 6.34 | 9.20 | 9.59 | 3.98 | 0.47 |
| 6 | 10.55 | 8.19 | 9.81 | 5.73 | 6.24 | 9.14 | 9.55 | 4.31 | 0.68 |
| 7 | 10.51 | 7.96 | 9.76 | 5.64 | 5.96 | 9.17 | 9.46 | 3.40 | 0.53 |
| 8 | 10.48 | 8.24 | 9.72 | 5.71 | 6.26 | 8.96 | 9.49 | 5.28 | 0.93 |
| 9 | 10.48 | 8.11 | 9.71 | 5.70 | 6.10 | 9.10 | 9.42 | 3.69 | 0.63 |
| 10 | 10.41 | 8.18 | 9.73 | 6.05 | 6.06 | 9.32 | 9.29 | 4.53 | 0.43 |




Figure 14. PDFs of the gas temperature of H I, H II and metals inside the virial radius of the most massive halo in the L14-RT simulation at $z = 6$. The majority of high-temperature gas ($T > 2 \times 10^4$ K) is ionized, while most of the gas at $T < 2 \times 10^4$ K is neutral. The metals are proportionately spread between the two components by mass inside of the haloes. Approximately two-thirds of the metal mass is found in the ionized component, while the other one-third remains neutral.

drastically because of the presence of UV radiation, and the scatter in the H_2 -SFR relation significantly increases to about ~ 1 dex. The range of SFRs exhibited by the haloes in the runs with RT is similar to those in the runs without (e.g. Rosdahl et al. 2015).

In Fig. 16, we show surface density maps of H_2 in the most massive halo for each of the six different simulations at $z = 6$. Moving from lower to higher resolution increases the amount of structure in the H_2 regions, regardless of whether or not we include RT. The differences between the H_2 surface density maps for a given halo are likely due to a number of reasons. The H_2 abundance is sensitive to density, temperature, metallicity and the local radiation field. At higher resolutions, we resolve higher densities so, in principle, H_2 should be able to form slightly more efficiently in the most resolved cells of the highest resolution simulations. Furthermore, the star formation algorithm is based on random Poisson sampling, and for these low-resolution simulations, we cannot expect perfect agreement as a function of resolution because of the different densities sampled. We also cannot expect a perfect agreement inside a spe-




Figure 15. H_2 masses versus SFR for all haloes with $M_{\text{halo}} > 10^{10} M_{\odot} h^{-1}$ at $z = 6$. Simulations with and without RT are shown as squares and circles, respectively. Black, red and blue points represent simulations refined to L12, L13 and L14. The enlarged points represent the most massive halo in the box. The scatter in the positive correlation between H_2 mass and SFR increases significantly when stellar radiation is included.

cific halo for the temperature that is once again dependent on the density of the gas and where the stars are located.

There is a much more extended halo of H_2 in the simulations without RT. This is because the haloes in our simulation are optically thin to Lyman-Werner radiation and all low-density H_2 is easily destroyed. This is better seen in Fig. 17, where we plot a 2D histogram of the H_2 fraction versus density for all cells in the two highest resolution simulations at $z = 6$. In the simulations without RT, the H_2 fraction of the IGM is $\sim 10^{-6}$, but as it approaches the more metal enriched and slightly denser CGM, it increases to a value of $\sim 10^{-3}$ as its formation is not stopped since there is no Lyman-Werner background. As the gas reaches higher and higher densities, the H_2 fraction continues to increase. From the bottom panel of Fig. 17, we see a markedly different dependence of the H_2 fraction as a function of density in the simulations that include stellar RT. At volume densities below $n_{\text{H}} \sim 10^{-2} \text{ cm}^{-3}$, the Lyman-Werner background is extremely efficient at destroying any residual H_2 and the molecular fraction remains at $\sim 10^{-10}$. It is only at these




Figure 16. First three columns: column density maps of H_2 in 20 kpc cubes surrounding the most massive galaxy in our simulations at $z = 6$. The locations of the young stars ($t < 10$ Myr) are indicated as white points. The young stars are the most likely sources of the UV and Ly α flux, while the H_2 regions may be the dominant source of [C II] emission. Fourth column: surface mass density of metals. In both simulations, the locations of the young stars are coincident with the locations where H_2 has formed. However, there is significantly less H_2 in the simulations that include stellar radiation due to strong Lyman–Werner feedback.

densities where even marginal self-shielding can begin to take place in the metal-enriched gas. However, only at the highest densities probed by our simulation does the H_2 fraction even approach $\sim 10^{-2}$. Effectively, we find that the amount of H_2 in our simulations with stellar radiation is very small. The consumption time of this gas due to star formation is very short. Our simulation is either too low resolution to model H_2 self-shielding well or we have not increased the sub-grid clumping factor enough in order to form a sufficient quantity of H_2 .

In Fig. 12, we also plot the H_2 mass as a function of redshift for the most massive halo in the three simulations that include stellar radiation. We see that the H_2 mass for a single galaxy can fluctuate by more than ~ 2 dex within a few Myr. The peaks in the H_2 mass correspond well with the peaks in the young stellar mass. Likewise, the peaks in the H_2 mass are associated with the troughs in the middle-aged stellar mass. Even for the highest resolution run where the sinusoidal pattern is smoothed out, there is still a clear offset between the middle-aged stars and the young stars and H_2 . This offset occurs in time as well as spatially. In Fig. 16, one can see that the location of the young stars agrees well with the highest surface density of H_2 because the two will form in cold dense gas clouds. The middle-aged stars are much more spread out (see Fig. 11) and are not associated with these cold dense clumps of gas, and SN feedback has already effectively destroyed the birth clouds.

The mass and spatial distributions of H_2 in the simulated galaxy are very dynamic and change on short time-scales. H_2 is very sensitive to the temperature, density and ionization state of the gas so it is not surprising that we find an anticorrelation between H_2 and middle-aged stars. One major shortcoming of our simulations is the density threshold at which star formation occurs. We have set this density threshold to $n_{\text{H}} = 0.1 \text{ cm}^{-3}$, which happens to coincide with the density above which H_2 will begin to form efficiently. Realistically, star formation occurs at much higher densities, and thus if we increase the density threshold for star formation, there will be

a time offset between when H_2 begins to form and when the stars begin emitting high-energy, UV photons. During this time delay, enough H_2 may form where it can self-shield better and thus our estimate for the H_2 mass will change.

3.4.6 The spatial distribution of C II, N II, O I and O III

Four spectral lines of particular interest for ALMA observations at high redshift are [C II] at $158 \mu\text{m}$, [N II] at $122 \mu\text{m}$, [O I] at $63 \mu\text{m}$ and [O III] at $88 \mu\text{m}$. Because our RT simulations follow the temperature, metallicity and radiation at all frequencies, including and shorter than the Habing band, we can make predictions for the spatial distribution and total mass of C II, N II, O I and O III in our galaxy. We post-process the central region ($r < 20 \text{ ckpc } h^{-1}$) of the most massive galaxy in the L14-RT simulation at $z = 6$ using CLOUDY (Ferland et al. 2013) to calculate the ionization states of each of these four ions in our simulation. We use the total hydrogen density, temperature and metallicity of each of the ~ 20000 cells in the central region of the galaxy as inputs. For the radiation field, we assume a spectral shape within each radiation bin that is similar to the spectrum in the Milky Way ISM (Black 1987). We renormalize this spectrum in each bin such that the intensity within the energy bins is consistent with the intensity measured for each individual cell in our simulation. For energies lower than the Habing band, we normalize the spectrum by the same scalefactor as we measure for the Habing band in the simulation so that the shape is continuous at low energies. We add an additional radiation field to model the contribution of the CMB at $z = 6$.

In Fig. 18, we plot the phase-space diagrams of temperature versus total gas density for all cells in the central region of the most massive galaxy in the L14-RT simulation at $z = 6$. The top left-hand panel of this figure shows the mass-weighted phase-space diagram. We see that most of the gas mass has $T \geq 10^4 \text{ K}$, while a few cells scatter to lower temperatures. The additional five panels show the




Figure 17. Mass-weighted phase-space diagrams of the H_2 fraction as a function of gas mass-density at $z = 6$ for the L14-No RT simulation and the L14-RT simulation. When stellar radiation is included, any residual H_2 at low densities is destroyed by the Lyman–Werner background. Overall, significantly less H_2 forms in the simulations that include stellar radiation.

same phase-space diagram as has been weighted by the ionization parameter [$U = n_\gamma(E > 13.6 \text{ eV})/n_{\text{H}}$], and C II , N II , O I and O III fractions, respectively. While most of the gas has $T \geq 10^4 \text{ K}$, the cells that have high C II fractions have $T \leq 10^4 \text{ K}$. There still exists some C II at warmer temperatures up to $T \sim 10^5 \text{ K}$. By contrast, almost all of the N II resides in gas with $T \sim 10^4 \text{ K}$ and not much N II exists at lower temperatures. O I looks similar to C II in that it is most prevalent at $T \leq 10^4 \text{ K}$. However, we see a much stronger decline in O I at warmer temperatures than we see for C II . O III exists at somewhat higher temperatures [$10^4 < T(\text{K}) < 10^5$], and coincides with ionized hydrogen. We can also see that O III exists where the ionization parameter is high.

In Fig. 19, we plot the surface density of each of the four ions, and the features we identified in the phase-space diagrams are reflected spatially. The surface density of C II is clumpy and these clumps are associated with the cold neutral gas. The features in N II are not as strong as in C II because the neutral gas at $T \sim 10^4 \text{ K}$ has a more extended distribution than the much colder, neutral gas (see Fig. 13). Unsurprisingly, O I looks very similar to C II , but is not quite as extended as C II because it cannot exist at slightly warmer

temperatures. In general, the surface density of O I is higher than for C II mainly because the total abundance of oxygen is greater than that of carbon by a factor of ~ 2 (Asplund et al. 2009; Grevesse et al. 2010). The spatial distribution of O III is different from that of O I and C II because it is the dominant ionization state of oxygen in a much different part of phase space. Inside the regions shown, we find that the total mass of C II , N II , O I and O III is 6.3×10^5 , 5.3×10^4 , 1.3×10^6 and $2.6 \times 10^5 M_\odot$, respectively, compared to a total metal mass in the same region of $1.7 \times 10^7 M_\odot$. We also show in Fig. 19 the spatial distribution of the ionization parameter along the line of sight, weighted by density, and we can see how this anticorrelates with the location of C II and O I . Finally, in the top left-hand panel of Fig. 19, we plot the neutral hydrogen column density that correlates well with C II and O I as expected and anticorrelates with the ionization parameter. The red and blue contours show the normalized distributions of C II and O III column densities, respectively. We can see that O III is more extended and there is an offset between the peaks. Likewise, the yellow contours show the normalized distribution of the column density of young stars. We see two distinct peaks of young stars, one of which is associated with the peak in C II and the other occurs to the bottom left-hand side, off centre from the densest region of the galaxy.

Ideally, we would convert the C II mass into a luminosity and discuss its observability for ALMA. Various works have attempted to do this (e.g. Vallini et al. 2013, 2015, 2017); however, modelling this emission is complicated. The fine structure line is likely excited by collisions with neutral hydrogen, free electrons and protons. Our modelling allows us to make predictions for these quantities, but they are likely to fail in the highest density, metal-enriched regions, because we do not take into account the free electrons that come from dust grains, photoelectric heating and the ionization states of the metals (Draine 1978; Helou et al. 2001). The structure of the ISM at these scales is likely important due to self-shielding across the face of the molecular clouds. For these reasons, we refrain here from estimating the $[\text{C II}]$ and $[\text{O III}]$ emissions of our galaxy.

4 DISCUSSION

4.1 Comparison with previous simulations

Various previous simulations have attempted to model the IR emission coming from galaxies during the epoch of reionization (Nagamine et al. 2006; Vallini et al. 2013, 2015; Pallottini et al. 2017). Our work builds on these previous results, and, in particular, the inclusion of on-the-fly, multifrequency, RT and radiation-coupled H_2 chemistry allows for an improved modelling of star-forming galaxies during the epoch of reionization. In this section, we compare some of our results to these previous simulations.

Vallini et al. (2013, 2015) have run detailed simulations attempting to identify the origin of $[\text{C II}]$ emission in a ‘normal’ high-redshift star-forming galaxy. Vallini et al. (2013) use cosmological SPH simulations to model the formation of a $10^{11} M_\odot$ halo at $z = 6.6$ that has very similar mass to the most massive galaxy in our work. These simulations were then post-processed with UV radiation transfer using a Monte Carlo ray-tracing code. They find that the warm neutral medium resides in overdense clumps that are displaced from the central star-forming regions and that cold gas resides in very dense clumps. Furthermore, they show that at solar metallicity, 95 per cent of $[\text{C II}]$ emission originates from cold gas.

We find a large difference in the morphology of our galaxy compared to Vallini et al. (2013), which forms one large dense clump of stars (see their fig. 1). This likely arises because the simulations




Figure 18. Phase-space diagrams of density versus temperature for all cells within a viral radius of the most massive halo in the L14-RT simulation at $z = 6$. The cells in the top left-hand panel are mass-weighted, while the cells in the other panels are weighted by the ionization parameter (U), and C II, N II, O I and O III fractions, respectively. The top colour bar shows the colour scale of the mass-weighted plot, while the bottom colour bar represents the colour scale for the other five plots. U remains enhanced in the high-temperature gas where we see the highest fractions of O III. C II and O I occupy a very similar part of the phase-space diagram at low temperatures corresponding to the neutral gas.

used in Vallini et al. (2013) include neither radiative cooling nor SNe so a complex morphology cannot arise and radiation post-processing will not affect the structure of the gas. Because the stars are concentrated in one central clump, we find a different temperature distribution within our galaxy, which can significantly affect the ionization state of carbon.

The work of Vallini et al. (2015) significantly improves on Vallini et al. (2013) by using the same underlying SPH simulation but including a recipe for the inhomogeneous distribution of metals, the formation of molecular clouds and the effect of the cosmic microwave background on [C II] emission. Although our simulations likely underpredict the formation of H_2 compared to Vallini et al. (2015), we find that the molecular clouds form in the central and most dense regions of the galaxy, consistent with Vallini et al. (2015). They find that most of the [C II] emission arises from PDRs rather than the cold neutral medium, indicating that the [C II] is likely coming from illuminated molecular clouds. If this is in fact the case, we can use our simulation to potentially understand the morphology of the [C II] emission as well as the UV and Ly α to better understand the observations of Maiolino et al. (2015) (see Section 4.2).

(Pallottini et al. 2017, hereafter P17) performed a similar exercise using a zoom-in simulation and attempted to model the [C II] emission from a single high-redshift $z = 6$ galaxy of similar mass to our work and Vallini et al. (2013, 2015). In contrast to Vallini et al. (2013, 2015), this simulation models radiative cooling, feedback and metal enrichment and forms a disc galaxy with an SFR of $\sim 15 M_\odot \text{ yr}^{-1}$. This allows the authors to explicitly track the

location of metals. They find that 30 per cent of the C II mass has been ejected from the central, radiation pressure-supported disc due to outflows, but 95 per cent of the [C II] luminosity comes from the disc. Furthermore, they show that this galaxy is underluminous compared to the local [C II]–SFR relation.

Comparing our work to P17, the total stellar mass of our most massive galaxy is significantly lower (by an order of magnitude). The different feedback implementations are likely to be the main reason for this difference. The delayed cooling stellar feedback model used in our work efficiently regulates star formation in the galaxy, which allows our stellar mass estimate to be consistent with Behroozi et al. (2013). To support this conclusion, P17 find that the metallicity of the galaxy drops to the artificial metallicity floor at 12 kpc, whereas we find metals at nearly double this radius, indicative that our feedback model is driving stronger outflows. Because the metals in our galaxy are much more spread out, the average metallicity of the gas is lower, which can have drastic effects on the observability of [C II] (Vallini et al. 2013). This would mean that our galaxy would fall even lower on the [C II]–SFR relation compared to P17.

Interestingly, the H_2 mass measured in P17 is comparable to what we find in our simulations that do not include stellar radiation. P17 uses a different model for H_2 formation that is not explicitly coupled to the radiation field, and the equations are independent of the flux in the Lyman–Werner band (see their equation 2 d). At our resolution, the amount of H_2 that forms in the galaxy is very sensitive to the local enhancement of the Lyman–Werner flux above the background (see Fig. 12). Our simulations likely underpredict

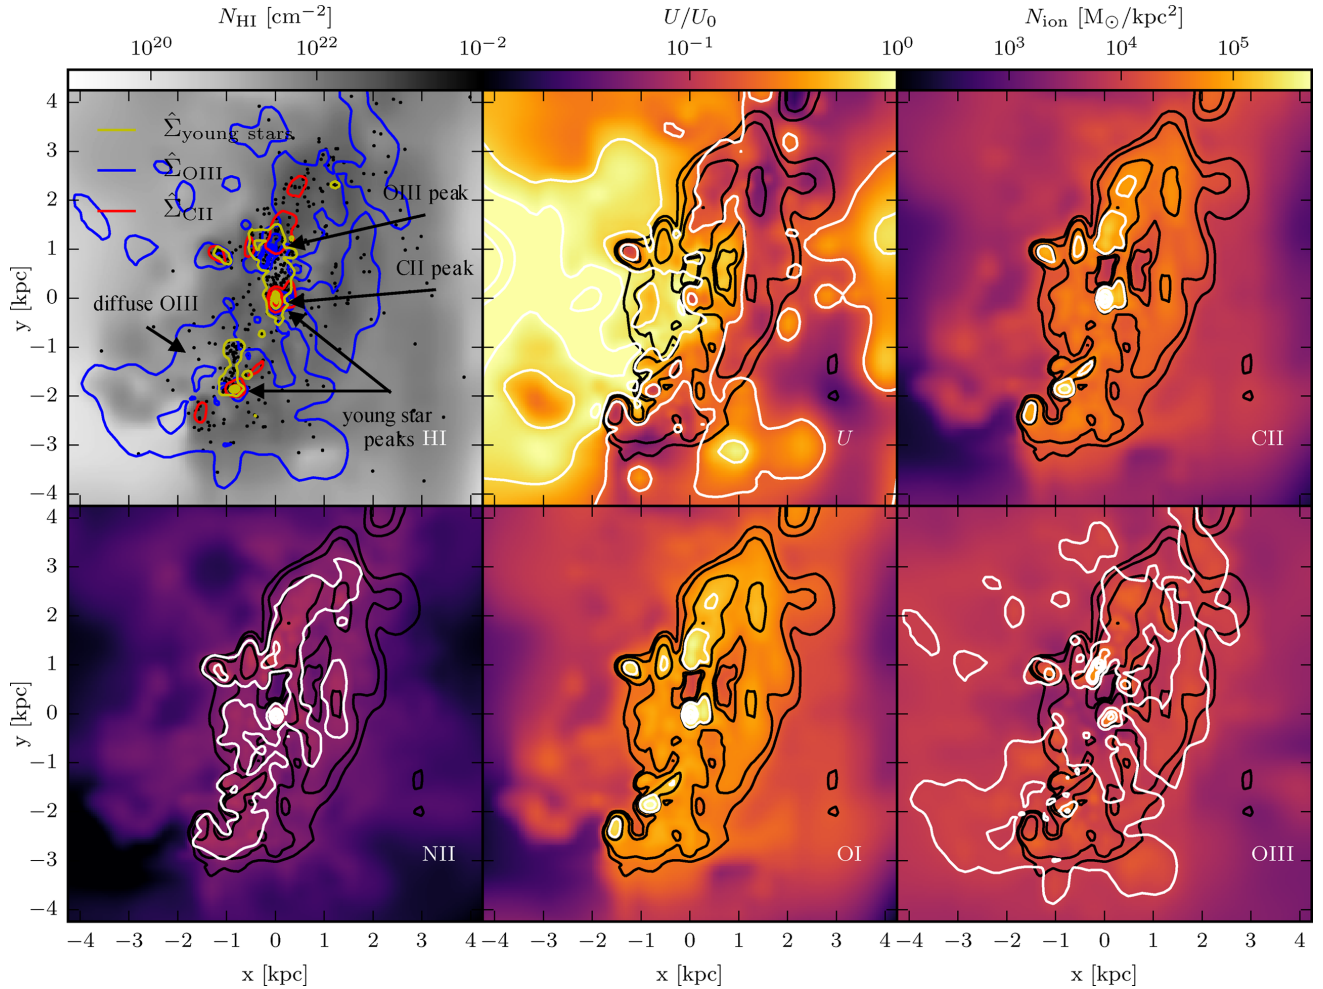


Figure 19. Surface density plots of H I (top left-hand panel), U (top middle panel), C II (top right-hand panel), N II (bottom left-hand panel), O I (bottom middle panel) and O III (bottom right-hand panel) around the most massive halo in the L14-RT simulation at $z = 6$. In the top left-hand panel, the column density contours of O III, C II and young stars are shown in blue, red and yellow, respectively. The contours have been normalized to the peak column density in order to best show the locations of the peak surface densities of each ion. The projected locations of young stars are shown as black points. The O III peak is spatially offset from the C II peak and young star peaks, while the C II peak is colocated with one of the young star peaks. In the other five panels, the black contours represent H I column densities of 2×10^{20} , 3×10^{20} and $7 \times 10^{20} \text{ cm}^{-2}$. The white contours show the normalized surface densities representing 10 per cent, 30 per cent, 50 per cent, 80 per cent and 90 per cent of the maximum value of the quantity plotted in that image. The maximum surface densities of C II, N II, O I and O III are 7.0×10^5 , 2.5×10^4 , 1.6×10^6 and $9.3 \times 10^4 \text{ M}_\odot \text{ kpc}^{-2}$, respectively. The colour scale for U has been normalized to have a maximum of 1, while the top right-hand colour bar applies to the C II, O I, N II and O III maps.

the H_2 mass in these haloes due to low resolution, which hampers self-shielding. Because the H_2 mass in our simulation corresponds well with the high-density peaks, we predict that [C II] luminosity will also derive from the H_2 regions, consistent with P17 and Vallini et al. (2015).

4.2 Interpreting ALMA observations

Various observations targeting the [C II] emission line during the epoch of reionization have been successful (Capak et al. 2015; Maiolino et al. 2015; Willott et al. 2015; Knudsen et al. 2016; Pentericci et al. 2016). These works have revealed a number of interesting properties of star-forming high-redshift galaxies. In particular, they find that the far-IR emission of these galaxies is weaker than similar galaxies at lower redshift. This could indicate a decrease in dust mass at higher redshifts and/or low metallicities. The galaxies tend to fall low on the [C II]–SFR relation derived from local galaxies.

Maiolino et al. (2015) observed three spectroscopically confirmed Lyman break galaxies at $6.8 < z \leq 7.1$ with ALMA that have SFRs of $5\text{--}15 \text{ M}_\odot \text{ yr}^{-1}$. Interestingly, Maiolino et al. (2015) detect no [C II] emission at the location of the Y-band counterpart that probes the rest-frame Ly α and UV emissions and likely the location of young stars. Where [C II] is detected, it is spatially offset by a few kpc from the central core of the galaxy.

Our cosmological RT simulations support the interpretation of the offset between the UV/Ly α and [C II] emissions proposed by Maiolino et al. (2015). The UV and Ly α emissions come from young stars that are predicted to form in tight clumps from our simulation. The [C II] emission is likely to originate in cold, neutral gas, or in PDRs close to young stars. Our simulation includes modelling of the inhomogeneous radiation field of the galaxy that leads to a complex morphology in the H I distribution and temperature. We probably underpredict the total mass in H_2 , although the locations where H_2 forms are likely to be spatially consistent with the cold neutral medium. The final two pieces of the puzzle are the metal

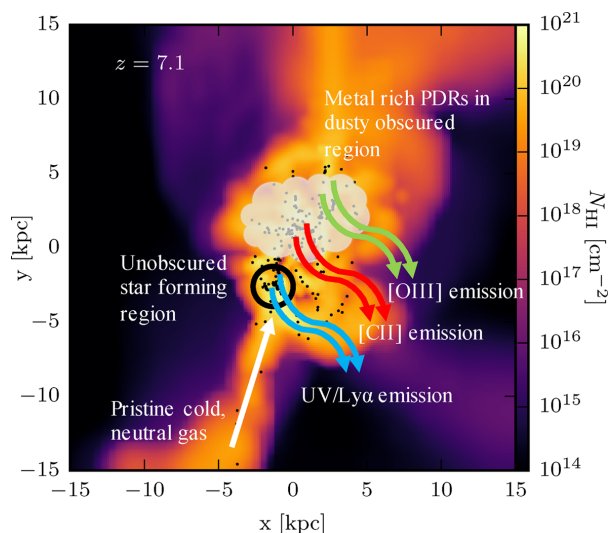


Figure 20. Schematic view for how an offset can arise between UV/Ly α , [C II] and [O III] emissions. The underlying image shows the surface density of neutral gas for the most massive halo in the L14-RT simulation at $z = 7.1$, the redshift of the Maiolino et al. (2015) observation. The locations of young stars that emit in the UV/Ly α are shown as black points. The white arrow shows the direction by which cold, neutral, low-metallicity gas is fed on to the galaxy. The intersection point of the filament and the galaxy creates a star-forming region circled in black, which we assume is relatively unobscured. Towards the central regions of the galaxy, we have placed a dust cloud where we expect the metallicity and dust content to be slightly higher. We propose that here the UV/Ly α emission is obscured, but the [C II] emission coming from PDRs and [O III] emission from higher temperature, ionized gas can escape. The expected UV/Ly α emission is shown as the blue arrows, while the [C II] and [O III] emissions are shown as the red and green arrows, respectively.

and dust distributions. Our simulations do not capture small-scale inhomogeneities in the metal distribution, and we do not explicitly follow dust so the distribution of both quantities is uncertain. Dust is particularly important because of its ability to obscure star-forming regions.

In Fig. 20, we show how spatial offsets between UV/Ly α and [C II] emissions can arise depending on where the dust is located. We show how this may work by annotating a neutral hydrogen surface density map of the most massive galaxy in our simulation at $z = 7.1$, the redshift of the Maiolino et al. (2015) observation. At the bottom of the image, a cold neutral filament is feeding pristine or low-metallicity gas on to the central regions of the galaxy. At the intersection of the galaxy and the filament, a large star-forming clump has formed, which we envision to have reasonably low metallicity and thus little [C II] emission. Because of the low metallicity and dust content, the UV and Ly α emissions escape without being reprocessed and are visible. Towards the central regions of the galaxy, one may expect that the metallicity and dust content are higher, and thus star-forming clouds in this region are much more obscured and no UV or Ly α emission escapes this region. However, in these dense clouds, PDRs produce [C II] emission that can pass through the dust; hence, an offset is created between the UV/Ly α and [C II] emissions. To calculate the dust attenuation, what is required is the dust optical depth, τ_d . Our simulations do not self-consistently trace dust. If we make again the simplistic assumption that the dust optical depth scales linearly with neutral hydrogen column density and metallicity (Gnedin et al. 2009), we find that the maximum τ_d is a factor of about three times higher in the central star-forming

regions compared to the outer regions of our simulated galaxies due to metallicity and density gradients in the galaxy. The difference in surface mass density of young stars between the inner and outer regions is a factor of ~ 1.5 . Note that the attenuation scales as $e^{-\tau_d}$ and that we expect more attenuation in the central regions. These values should, however, be considered with some caution as a more sophisticated model for dust physics will be needed to get more accurate estimates of the attenuation. Furthermore, resolution effects in our simulations limit the maximum densities in the simulations and therefore put an upper limit on τ_d .

If [N II] and [O I] could also be observed, we predict that these will be spatially coincident with the [C II], although the [N II] might be slightly more extended. By contrast, [O III] arises from higher temperature gas at a higher ionization parameter and thus would be offset from [C II]. Therefore, we also might expect an offset between [O III] and [C II].

5 CAVEATS

As with all numerical simulations, there are a number of caveats, many of which we have already discussed in the text. At large scales, reionization is sensitive to the resolution and star particle mass of the simulation (Aubert & Teyssier 2010) and the galaxies resolved. Here we do not resolve the atomic cooling threshold haloes that may help drive reionization (Kimm et al. 2017). Most of the metals in our simulations remain inside the virial radius of the haloes and do not efficiently enrich the IGM. This is rather sensitive to the feedback method implemented in the simulation (Keating et al. 2016), and thus our conclusions may be subject to our chosen feedback model.

The density threshold of star formation in our simulations is much lower (0.1 cm^{-3}) than the typical densities observed in local molecular clouds ($> 100 \text{ cm}^{-3}$). Because the galaxies in our simulations have low metallicity, H_2 forms very efficiently only at densities higher than the density threshold where we assume stars form. For this reason, the Lyman–Werner radiation is extremely efficient at destroying the H_2 as the resolution of our simulations is too low to model self-shielding properly. The formation of H_2 depends on the clumping factor chosen. We have experimented with increasing the clumping factor and, unsurprisingly, significantly more H_2 forms. Higher resolution simulations will be needed to properly model the formation of H_2 . Nevertheless, the conclusion that H_2 is located in dense, metal-enriched gas will continue to hold.

Likewise, because of the still rather moderate resolution of our simulations, the metals efficiently spread throughout the central regions of the haloes. A higher resolution simulation might resolve a more patchy metal distribution in the ISM. If small patches of gas have significantly higher metallicity, this will affect where the [C II] emission originates from and may change the conclusion with regard to its observability (Vallini et al. 2013, 2015).

Finally, our simulations make no attempt to model the spatial distribution of dust that plays a role in various astrophysical processes such as the cooling and fragmentation of gas, the formation of H_2 , absorption of UV photons and the obscuration of UV/Ly α emission.

6 CONCLUSIONS

We have presented six cosmological simulations, three of which include on-the-fly multifrequency RT and three without, at various resolutions, which demonstrate how inhomogeneous stellar radiation affects galaxy formation at large (Mpc), intermediate (tens of kpc) and small (sub-kpc) scales. The simulations employ a new VSLA

that allows I-fronts to travel at the appropriate speed in both low- and high-density regions before Strömgren spheres have reached their maximum extent. The simulations include a non-equilibrium chemistry model for the formation and destruction of H_2 that is coupled to the spatially inhomogeneous radiation from star particles. Our findings can be summarized as follows:

(i) At large scales, reionization proceeds inside-out. The highest density regions surrounding haloes are ionized first. However, a significant amount of neutral gas remains inside the haloes with a typical ratio of masses of ionized to neutral gas of 2:1.

(ii) The simulation becomes optically thin to radiation in the Lyman–Werner band at a significantly higher redshift than it does for photons that ionize hydrogen and helium. Inside of haloes, the flux in the Lyman–Werner band follows a characteristic r^{-2} profile, while hydrogen ionizing radiation is dominated by the individual star-forming clumps as it is readily absorbed internally by cold neutral gas.

(iii) At intermediate scales, the filaments remain mostly self-shielding with H I column densities above the Lyman limit threshold ($N_{\text{H}} > 10^{17} \text{ cm}^{-2}$). The spatial distribution of the DLAs is very dynamic and strongly anticorrelated with the presence of young stars.

(iv) The distribution of young stars inside of haloes is strongly clustered, and the locations of these stars are closely associated with their birth clouds. As stars age, they redistribute throughout the galaxy, which creates a smooth, old stellar component. The UV/Ly α radiation will be dominated by these clumps, and its observability will be subject to the presence of dust.

(v) We find rapid changes in the spatial distribution of ionized and neutral gas inside of the haloes due to young stars ionizing their surroundings and SN feedback blowing holes in the neutral medium.

(vi) The metal mass in our simulations is proportionally distributed between the ionized and neutral gases inside of the haloes. The 10 most massive haloes in our simulation with masses $10^{10.4} < M_{\text{DM}} < 10^{11.2}$ all exhibit the expected cosmic baryon fraction within their virial radii.

(vii) For our specific feedback implementation, most of the metals are confined inside the virial radius of haloes. Many small haloes form along the filaments that increase the metallicity of this gas slightly above the metallicity floor of our simulations; however, much of the gas at these densities has a metallicity that remains close to primordial.

(viii) H_2 masses in the haloes are time variable and extremely sensitive to the presence of young stars. Without stellar radiation, a strong correlation arises between the SFR and the H_2 mass, while with RT, this correlation weakens and the scatter significantly increases. Our simulations, however, likely underpredict the total mass of H_2 in the haloes because they have not reached H_2 densities where it can effectively self-shield.

(ix) The C II in the galaxy is associated with the locations of the cold neutral clumps of gas. O I follows a similar trend; however, it has a less extended distribution compared to C II . By contrast, the distribution of N II is more diffuse and associated with much warmer gas, and O III exists in hotter, ionized gas.

(x) Depending on the location of dust in the galaxy, spatial offsets are expected between UV/Ly α and $[\text{C II}]$ emissions, which would explain the recent observations of Maiolino et al. (2015). A spatial offset also naturally arises between $[\text{C II}]$ and $[\text{O III}]$ because of the different parts of phase space where each of the ionization states is

dominant. It should thus be expected to regularly find spatial offsets between UV/Ly α , $[\text{C II}]$ and $[\text{O III}]$ emissions in the same object.

With our simulations, we have shown that the inclusion of local radiation from young stars is crucial to realistically model the ISM in high-redshift galaxies and, in particular, the spatial distribution of the different ionization states of oxygen and carbon that can be observed by ALMA. To make meaningful comparisons with observations, future simulations will require a further improved modelling of dust in order to better model the formation of H_2 as well as better understand the obscuration that might cause a spatial offset between $[\text{C II}]$ and UV emissions. Higher resolution is also required to better model the distribution of metals in the haloes and to better resolve the physics governing the $[\text{O III}]$ and $[\text{C II}]$ emissions in a multiphase ISM.

ACKNOWLEDGEMENTS

We thank the referee for their review of this manuscript and very constructive comments. We thank Roberto Maiolino, Stefano Carniani and Joki Rosdahl for comments on the manuscript. HK is grateful to Laura Keating and Fernanda Ostrovski for useful conversations and the revision of the manuscript. Furthermore, we thank the MIAPP for the reionization workshop in Garching, where some of this work was completed.

This work made considerable use of the open source analysis software `PYBODY` (Pontzen et al. 2013). HK thanks Foundation Boustany, the Cambridge Overseas Trust and an Isaac Newton Studentship. Support by ERC Advanced Grant 320596 ‘The Emergence of Structure during the Epoch of reionization’ is gratefully acknowledged. DS acknowledges support by STFC and ERC Starting Grant 638707 ‘Black holes and their host galaxies: coevolution across cosmic time’.

This work was performed using the DiRAC/Darwin Supercomputer hosted by the University of Cambridge High Performance Computing Service (<http://www.hpc.cam.ac.uk/>), provided by Dell Inc. using the Strategic Research Infrastructure Funding from the Higher Education Funding Council for England and funding from the Science and Technology Facilities Council.

This work used the DiRAC Complexity system, operated by the University of Leicester IT Services, which forms part of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment is funded by BIS National E-Infrastructure capital grant ST/K000373/1 and STFC DiRAC Operations grant ST/K0003259/1. DiRAC is part of the National E-Infrastructure.

Furthermore, this work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by the BIS National E-Infrastructure capital grant ST/K00042X/1, STFC capital grant ST/K00087X/1, DiRAC operations grant ST/K003267/1 and Durham University. Dirac is part of the National E-Infrastructure.

REFERENCES

- Abel T., Wise J. H., Bryan G. L., 2007, *ApJ*, 659, L87
- Appleton P. N. et al., 2013, *ApJ*, 777, 66
- Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *ARA&A*, 47, 481
- Aubert D., Teyssier R., 2008, *MNRAS*, 387, 295
- Aubert D., Teyssier R., 2010, *ApJ*, 724, 244
- Aubert D., Deparis N., Ocvirk P., 2015, *MNRAS*, 454, 1012
- Baczynski C., Glover S. C. O., Klessen R. S., 2015, *MNRAS*, 454, 380
- Bauer A., Springel V., Vogelsberger M., Genel S., Torrey P., Sijacki D., Nelson D., Hernquist L., 2015, *MNRAS*, 453, 3593

- Becker G. D., Sargent W. L. W., Rauch M., Calverley A. P., 2011, *ApJ*, 735, 93
- Behroozi P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 770, 57
- Bigiel F., Leroy A., Walter F., Brinks E., de Blok W. J. G., Madore B., Thornley M. D., 2008, *AJ*, 136, 2846
- Black J. H., 1987, in Hollenbach D. J., Thronson H. A., Jr, eds, *Astrophysics and Space Science Library Vol. 134, Interstellar Processes*. D. Reidel Publishing Co., Dordrecht, p. 731
- Black J. H., Dalgarno A., 1977, *ApJS*, 34, 405
- Bolton J. S., Haehnelt M. G., 2007, *MNRAS*, 382, 325
- Bolton J. S., Haehnelt M. G., Warren S. J., Hewett P. C., Mortlock D. J., Venemans B. P., McMahon R. G., Simpson C., 2011, *MNRAS*, 416, L70
- Bouwens R. J. et al., 2014, *ApJ*, 795, 126
- Bouwens R. J. et al., 2015, *ApJ*, 803, 34
- Bowman J. D., Rogers A. E. E., 2010, *Nature*, 468, 796
- Bruzual G., Charlot S., 2003, *MNRAS*, 344, 1000
- Bryan G. L. et al., 2014, *ApJS*, 211, 19
- Burton M. G., Hollenbach D. J., Tielens A. G. G. M., 1990, *ApJ*, 365, 620
- Capak P. L. et al., 2015, *Nature*, 522, 455
- Carilli C. L., Walter F., 2013, *ARA&A*, 51, 105
- Cen R., 2012, *ApJ*, 748, 121
- Chabrier G., 2003, *PASP*, 115, 763
- Chabrier G., 2005, in Corbelli E., Palla F., Zinnecker H., eds, *Astrophysics and Space Science Library Vol. 327, The Initial Mass Function 50 Years Later*. Springer, Dordrecht, p. 41
- Chardin J., Haehnelt M. G., Aubert D., Puchwein E., 2015, *MNRAS*, 453, 2943
- Chiosi C., Bertelli G., Bressan A., 1992, *ARA&A*, 30, 235
- Choudhury T. R., Haehnelt M. G., Regan J., 2009, *MNRAS*, 394, 960
- Choudhury T. R., Puchwein E., Haehnelt M. G., Bolton J. S., 2015, *MNRAS*, 452, 261
- Ciardi B., Ferrara A., Marri S., Raimondo G., 2001, *MNRAS*, 324, 381
- Coe D. et al., 2013, *ApJ*, 762, 32
- Commerçon B., Debout V., Teyssier R., 2014, *A&A*, 563, A11
- Cormier D. et al., 2012, *A&A*, 548, A20
- Costa T., Sijacki D., Haehnelt M. G., 2014, *MNRAS*, 444, 2355
- Couchman H. M. P., Rees M. J., 1986, *MNRAS*, 221, 53
- Crain R. A. et al., 2015, *MNRAS*, 450, 1937
- D'Odorico V. et al., 2013, *MNRAS*, 435, 1198
- De Looze I. et al., 2014, *A&A*, 568, A62
- De Looze I. et al., 2016, *MNRAS*, 459, 3900
- Dopita M. A., Krauss L. M., Sutherland R. S., Kobayashi C., Lineweaver C. H., 2011, *Ap&SS*, 335, 345
- Draine B. T., 1978, *ApJS*, 36, 595
- Draine B. T., Bertoldi F., 1996, *ApJ*, 468, 269
- Draine B. T. et al., 2007, *ApJ*, 663, 866
- Ellis R. S. et al., 2013, *ApJ*, 763, L7
- Fan X., Carilli C. L., Keating B., 2006, *ARA&A*, 44, 415
- Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, 703, 1416
- Ferland G. J. et al., 2013, *Rev. Mex. Astron. Astrofis.*, 49, 137
- Finlator K., Özel F., Davé R., Oppenheimer B. D., 2009, *MNRAS*, 400, 1049
- Fumagalli M., Prochaska J. X., Kasen D., Dekel A., Ceverino D., Primack J. R., 2011, *MNRAS*, 418, 1796
- Gallerani S. et al., 2012, *A&A*, 543, A114
- Gill S. P. D., Knebe A., Gibson B. K., 2004, *MNRAS*, 351, 399
- Glover S. C. O., Abel T., 2008, *MNRAS*, 388, 1627
- Glover S. C. O., Federrath C., Mac Low M.-M., Klessen R. S., 2010, *MNRAS*, 404, 2
- Gnedin N. Y., 2000, *ApJ*, 535, 530
- Gnedin N. Y., 2014, *ApJ*, 793, 29
- Gnedin N. Y., Abel T., 2001, *New Astron.*, 6, 437
- Gnedin N. Y., Kaurov A. A., 2014, *ApJ*, 793, 30
- Gnedin N. Y., Tassis K., Kravtsov A. V., 2009, *ApJ*, 697, 55
- Gong Y., Cooray A., Silva M., Santos M. G., Bock J., Bradford C. M., Zemcov M., 2012, *ApJ*, 745, 49
- Graciá-Carpio J. et al., 2011, *ApJ*, 728, L7
- Grevesse N., Asplund M., Sauval A. J., Scott P., 2010, *Ap&SS*, 328, 179
- Haardt F., Madau P., 1996, *ApJ*, 461, 20
- Haardt F., Madau P., 2012, *ApJ*, 746, 125
- Hahn O., Abel T., 2011, *MNRAS*, 415, 2101
- Helou G., Malhotra S., Hollenbach D. J., Dale D. A., Contursi A., 2001, *ApJ*, 548, L73
- Hollenbach D., McKee C. F., 1979, *ApJS*, 41, 555
- Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
- Iliev I. T., Mellema G., Pen U.-L., Merz H., Shapiro P. R., Alvarez M. A., 2006, *MNRAS*, 369, 1625
- Iliev I. T. et al., 2009, *MNRAS*, 400, 1283
- Kaufman M. J., Wolfire M. G., Hollenbach D. J., Luhman M. L., 1999, *ApJ*, 527, 795
- Keating L. C., Puchwein E., Haehnelt M. G., Bird S., Bolton J. S., 2016, *MNRAS*, 461, 606
- Kennicutt R. C., Jr, 1998, *ApJ*, 498, 541
- Kennicutt R. C., Jr, et al., 2003, *PASP*, 115, 928
- Kennicutt R. C. et al., 2011, *PASP*, 123, 1347
- Kim J.-h. et al., 2014, *ApJS*, 210, 14
- Kimm T., Cen R., 2014, *ApJ*, 788, 121
- Kimm T., Cen R., Devriendt J., Dubois Y., Slyz A., 2015, *MNRAS*, 451, 2900
- Kimm T., Katz H., Haehnelt M., Rosdahl J., Devriendt J., Slyz A., 2017, *MNRAS*, 466, 4826
- Knollmann S. R., Knebe A., 2009, *ApJS*, 182, 608
- Knudsen K. K., Richard J., Kneib J.-P., Jauzac M., Clément B., Drouart G., Egami E., Lindroos L., 2016, *MNRAS*, 462, L6
- Krumholz M. R., McKee C. F., Tumlinson J., 2008, *ApJ*, 689, 865
- Krumholz M. R., McKee C. F., Tumlinson J., 2009, *ApJ*, 693, 216
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
- Liu X., Shemansky D. E., 2012, *J. Phys. B: At. Mol. Phys.*, 45, 095203
- Ma X., Hopkins P. F., Faucher-Giguère C.-A., Zolman N., Muratov A. L., Kereš D., Quataert E., 2016, *MNRAS*, 456, 2140
- McKee C. F., Krumholz M. R., 2010, *ApJ*, 709, 308
- McLure R. J. et al., 2013, *MNRAS*, 432, 2696
- Madau P., Haardt F., 2015, *ApJ*, 813, L8
- Madden S. C., Poglitsch A., Geis N., Stacey G. J., Townes C. H., 1997, *ApJ*, 483, 200
- Maiolino R. et al., 2015, *MNRAS*, 452, 54
- Muñoz J. A., Furlanetto S. R., 2014, *MNRAS*, 438, 2483
- Nagamine K., Wolfe A. M., Hernquist L., 2006, *ApJ*, 647, 60
- Oesch P. A. et al., 2013, *ApJ*, 773, 75
- Oesch P. A. et al., 2014, *ApJ*, 786, 108
- Okamoto T., Gao L., Theuns T., 2008, *MNRAS*, 390, 920
- Ouchi M. et al., 2010, *ApJ*, 723, 869
- Pallottini A., Ferrara A., Gallerani S., Vallini L., Maiolino R., Salvadori S., 2017, *MNRAS*, 465, 2540 (P17)
- Pentericci L. et al., 2016, *ApJL*, 829, L11
- Pilbratt G. L. et al., 2010, *A&A*, 518, L1
- Pineda J. L., Langer W. D., Goldsmith P. F., 2014, *A&A*, 570, A121
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Pontzen A., Roškar R., Stinson G. S., Woods R., Reed D. M., Coles J., Quinn T. R., 2013, *Astrophysics Source Code Library*, ascl:1305.002
- Popping G., Pérez-Beaupuits J. P., Spaans M., Trager S. C., Somerville R. S., 2014, *MNRAS*, 444, 1301
- Rahmati A., Schaye J., Bower R. G., Crain R. A., Furlong M., Schaller M., Theuns T., 2015, *MNRAS*, 452, 2034
- Ricotti M., 2002, *MNRAS*, 336, L33
- Ricotti M., Ostriker J. P., 2004, *MNRAS*, 350, 539
- Röllig M. et al., 2007, *A&A*, 467, 187
- Rosdahl J., Blaizot J., Aubert D., Stranex T., Teyssier R., 2013, *MNRAS*, 436, 2188
- Rosdahl J., Schaye J., Teyssier R., Agertz O., 2015, *MNRAS*, 451, 34
- Ryan-Weber E. V., Pettini M., Madau P., Zych B. J., 2009, *MNRAS*, 395, 1476
- Schmidt M., 1959, *ApJ*, 129, 243

- Shapiro P. R., Giroux M. L., 1987, *ApJ*, 321, L107
 Sijacki D., Springel V., Haehnelt M. G., 2009, *MNRAS*, 400, 100
 Simcoe R. A. et al., 2011, *ApJ*, 743, 21
 Spitzer L., 1978, *Physical Processes in the Interstellar Medium*. Wiley-Interscience, New York
 Teyssier R., 2002, *A&A*, 385, 337
 Teyssier R., Pontzen A., Dubois Y., Read J. I., 2013, *MNRAS*, 429, 3068
 Tomassetti M., Porciani C., Romano-Díaz E., Ludlow A. D., 2015, *MNRAS*, 446, 3330
 Vallini L., Gallerani S., Ferrara A., Baek S., 2013, *MNRAS*, 433, 1567
 Vallini L., Gallerani S., Ferrara A., Pallottini A., Yue B., 2015, *ApJ*, 813, 36
 Vallini L., Ferrara A., Pallottini A., Gallerani S., 2017, *MNRAS*, 467, 1300
 Velusamy T., Langer W. D., 2014, *A&A*, 572, A45
 Venemans B. P. et al., 2012, *ApJ*, 751, L25
 Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
 Walter F., Riechers D., Cox P., Neri R., Carilli C., Bertoldi F., Weiss A., Maiolino R., 2009, *Nature*, 457, 699
 Wang R. et al., 2013, *ApJ*, 773, 44
 Willott C. J., Carilli C. L., Wagg J., Wang R., 2015, *ApJ*, 807, 180
 Wise J. H., Abel T., 2008, *ApJ*, 684, 1
 Wise J. H., Turk M. J., Norman M. L., Abel T., 2012, *ApJ*, 745, 50
 Xu H., Wise J. H., Norman M. L., Ahn K., O'Shea B. W., 2016, *ApJ*, 833, 84
 Zahn O. et al., 2012, *ApJ*, 756, 65
 Zheng W. et al., 2012, *Nature*, 489, 406

APPENDIX A: THE VARIABLE-SPEED-OF-LIGHT APPROXIMATION

The motivation for using a variable speed of light is that in high-density regimes (ISM) and for low-luminosity sources, a reduced speed of light generally gives correct results. However, in low-density regimes (IGM) and for highly luminous sources, it is necessary to use the full speed of light to properly model the propagation of I-fronts. Most AMR simulations use a density criterion to initiate refinement. The VSLA we introduce here interpolates between low values of the speed of light and high values of the speed of light depending on the AMR refinement level. This can make the simulation far less expensive than using the full speed of light everywhere while properly modelling I-front propagation in low-density regimes.

A1 Implementation

To implement the VSLA, we first identify the grid resolution and density at which we wish to use the full speed of light. In general, we set this as the level of the base grid in the simulation (in our case Level 8). Choosing a higher level would result in a slower, albeit more conservative approach. For each higher refinement level, we successively divide the speed of light by a factor of 2 up to the highest resolved level, or until we no longer wish to further reduce the speed of light. We choose to divide the speed of light by a factor of 2 at each higher level in order to keep the global time-step constant across all levels such that

$$c_{\text{sim}}(l) = \frac{c}{2^{l-l_{\text{min}}}}, \quad (\text{A1})$$

where l is the level of interest and l_{min} is the level where we set $c_{\text{sim}} = c$.

With oct-based refinement, as is used in *RAMSES*, upon refinement, the cell length is divided by a factor of 2. Dividing c_{sim} by a factor of 2 therefore also keeps t_{RT} constant. Dividing c by a larger factor would mean lower, less resolved levels act on a shorter time-step than the more refined levels. Dividing by a factor less than 2 would

correspond to a more conservative approach but would not fully exploit the speed-up possible with the VSLA, and we therefore choose to use factors of 2 in this work.

Consider the specific case where our base grid is resolved at Level = l_{base} , which is where $c_{\text{sim}}(l_{\text{base}}) = c$, and we impose seven additional levels of refinement so that at Level = $l_{\text{base}} + 7$, $c_{\text{sim}}(l_{\text{base}} + 7) = c/2^7 = 7.8 \times 10^{-3}c$. Using the RSLA, this simulation would, in principle, be $\sim 2^7 = 128$ times faster than using the full speed of light at all refinement levels; however, this would not model the I-fronts properly in the IGM. With the VSLA, our maximum time-step for our most refined levels is the same as in the RSLA with a similar gain, in principle, in terms of speed-up. What prohibits the VSLA from achieving the same speed-up as the RSLA is that, by definition, we abandon multisteping in time (i.e. all levels are run on the same time-step, which is generally dictated by the most refined cells). Depending on what fraction of all cells are refined, abandoning the multisteping advantage of AMR can correspond to different costs. If the computational load is completely dominated by the highest resolution cells (i.e. a zoom-in simulation), then the VSLA is unlikely to become much more expensive than the RSLA. In the case where very few cells are at the highest level of refinement, abandoning multisteping can become more costly at the expense of more accurate RT, and the refinement criteria must be optimized.⁹

We can improve the speed-up even further by introducing RT sub-cycling in the simulation where the radiative time-step is semi-decoupled from the hydro time-step. Even with the RSLA, the RT time-step is often much shorter than the corresponding hydro time-step, and therefore a speed-up can be achieved by sub-cycling multiple times over the RT calculation for every single hydro time-step. This has been implemented in the main version of *RAMSES*; however, because the hydro time-steps are adaptive, one has to adopt suitable boundary conditions at the interfaces between coarse and refined cells because multiple fine steps are executed without the coarse cells being fully updated (see Commerçon et al. 2014). Nevertheless, even with these special boundary conditions, we have confirmed that the method is accurate to a few per cent. By contrast, using the VSLA forces all hydro time-steps to be of the same length, and we are no longer forced to sub-cycle the RT calculation individually at each level. Instead, we have restructured *RAMSES* so that the RT calculation is performed after every coarse time-step. This allows us to sub-cycle through all AMR levels at once. The only downside to performing the calculation this way is that for large cosmological simulations, accessing the same array multiple consecutive times is faster than cycling through all of the different levels one after another because one must continually access slower areas of memory to retrieve the arrays that hold information for other levels. We have attempted to optimize our routines to achieve the speed-up one obtains with the original version; however, there is a moderate computational price for increased accuracy.¹⁰

⁹ Note that in the newest version of the VSLA (not used in this work), it is no longer required that multisteping be abandoned for the AMR. Rather, hydro and radiative time-steps can be adaptive, and we instead update the number of sub-cycles in the radiation sub-steps depending on the speed of light in the cells on that level.

¹⁰ Note again that in the newest version of VSLA (not used in this work), we have made our algorithm fully consistent with the boundary conditions adopted in Commerçon et al. (2014), which allows for a greater increase in speed. Thus, the slowdown of the VSLA compared to the RSLA reported in this work (especially for the cosmological test) is an overestimate of what we have seen with the newest version of the algorithm.




Figure A1. Simple setup of three cells to show how photon properties are advected across the refinement boundary.

A2 Conservation of flux and photon number density

To maintain the conservation properties of the code, we demonstrate a simple example of how photon flux and number density are advected across cell boundaries.

To calculate the intercell fluxes, we use the global Lax–Friedrichs (GLF) flux function, where

$$\mathcal{F}_{\text{GLF}} = \frac{\mathcal{F}_l + \mathcal{F}_{l+1}}{2} - \frac{c_{\text{sim}}(l)}{2}(\mathcal{U}_{l+1} - \mathcal{U}_l). \quad (\text{A2})$$

Here $\mathcal{U}_l = [N_l, \mathbf{F}_l]$ and $\mathcal{F}_l = [\mathbf{F}_l, c_l^2 \mathbb{P}_l]$, where N_l is the number density of photons in the cell, \mathbf{F}_l is the flux vector of the cell and \mathbb{P}_l is the pressure tensor of the cell. To calculate the pressure tensor, we have

$$\mathbb{P}_l = \mathbb{D}_l N_l. \quad (\text{A3})$$

We compute this for each of the different photon groups. \mathbb{D}_l is the Eddington tensor of the cell, and for this, we use the M1 closure such that

$$\mathbb{D}_l = \frac{1 - \chi_l}{2} \mathbf{I} + \frac{3\chi_l - 1}{2} \mathbf{n}_l \otimes \mathbf{n}_l, \quad (\text{A4})$$

where

$$\mathbf{n}_l = \frac{\mathbf{F}_l}{|\mathbf{F}_l|}, \quad \chi_l = \frac{3 + 4f_l^2}{5 + 2\sqrt{4 - 3f_l^2}}, \quad f_l = \frac{|\mathbf{F}_l|}{c_l N_l}. \quad (\text{A5})$$

If we consider the 1D case, which we will do for the rest of this derivation, $\mathbb{P}_l \rightarrow N_l$, and therefore $\mathcal{F}_l = [\mathbf{F}_l, c_l^2 N_l]$.

In Fig. A1, we show a setup of four cells with indices $l = [0 : 3]$. Each cell has a speed of light c_l , a number density of photons N_l and a photon flux \mathbf{F}_l . We begin by calculating the intercell fluxes between the central two cells, and we can use equation (A2) to calculate these values. In order to account for the differences in speed of light between the different cells, we make a slight modification to \mathcal{U} and \mathcal{F} for the cells to the left-hand side and right-hand side of the cell of interest so that

$$\mathcal{U}_{l\pm 1}(N_{l\pm 1}, \mathbf{F}_{l\pm 1}) \rightarrow \mathcal{U}_{l\pm 1}(N'_{l\pm 1}, \mathbf{F}_{l\pm 1}) \quad (\text{A6})$$

and

$$\mathcal{F}_{l\pm 1}(\mathbf{F}_{l\pm 1}, c_{l\pm 1}^2 N_{l\pm 1}) \rightarrow \mathcal{F}_{l\pm 1}(\mathbf{F}_{l\pm 1}, c_l^2 N'_{l\pm 1}), \quad (\text{A7})$$

where

$$N'_{l\pm 1} = \frac{c_{l\pm 1}}{c_l} N_{l\pm 1}. \quad (\text{A8})$$

Starting with the photon number densities, the intercell flux at b as seen by $l = 1$ is

$$\begin{aligned} \mathcal{F}_{\text{GLF}, b, l=1}^{N, l=1} &= \frac{\mathbf{F}_1 + \mathbf{F}_2}{2} - \frac{c_1}{2}(N'_2 - N_1) \\ &= \frac{\mathbf{F}_1 + \mathbf{F}_2}{2} - \frac{c_2}{2}N_2 + \frac{c_1}{2}N_1. \end{aligned} \quad (\text{A9})$$

Likewise, we can calculate the same value at b from the perspective of $l = 2$ as

$$\begin{aligned} \mathcal{F}_{\text{GLF}, b}^{N, l=2} &= \frac{\mathbf{F}_2 + \mathbf{F}_1}{2} - \frac{c_2}{2}(N_2 - N'_1) \\ &= \frac{\mathbf{F}_2 + \mathbf{F}_1}{2} - \frac{c_2}{2}N_2 + \frac{c_1}{2}N_1. \end{aligned} \quad (\text{A10})$$

Clearly, the fluxes calculated in equations (A9) and (A10) are identical, and, thus, it is trivial to show that the total number of photons is conserved, regardless of the speed of light in each of the cells, as long as we substitute the appropriate value of N into equation (A2). This will be detailed further below.

We can perform a similar exercise with the fluxes, and the intercell flux at b as computed by $l = 1$ is

$$\begin{aligned} \mathcal{F}_{\text{GLF}, b}^{F, l=1} &= \frac{c_1^2 N'_2 + c_1^2 N_1}{2} - \frac{c_1}{2}(\mathbf{F}_2 - \mathbf{F}_1) \\ &= \frac{c_1 c_2 N_2 + c_1^2 N_1}{2} - \frac{c_1}{2}(\mathbf{F}_2 - \mathbf{F}_1). \end{aligned} \quad (\text{A11})$$

In contrast, if we compute the intercell flux at $l = 2$, we have

$$\begin{aligned} \mathcal{F}_{\text{GLF}, b}^{F, l=2} &= \frac{c_2^2 N_2 + c_2^2 N'_1}{2} - \frac{c_2}{2}(\mathbf{F}_2 - \mathbf{F}_1) \\ &= \frac{c_2^2 N_2 + c_1 c_2 N_1}{2} - \frac{c_2}{2}(\mathbf{F}_2 - \mathbf{F}_1). \end{aligned} \quad (\text{A12})$$

Thus, there is an asymmetry between the intercell flux of the flux computed by $l = 1$ and that computed by $l = 2$ such that

$$\mathcal{F}_{\text{GLF}, b}^{F, l=2} = \frac{c_2}{c_1} \mathcal{F}_{\text{GLF}, b}^{F, l=1}. \quad (\text{A13})$$

In order to assure that no spurious flux is created or destroyed at the intercell boundary, one must update the cells at the boundary with a different flux. For the cell $l = 1$, we use $\mathcal{F}_{\text{GLF}, b}^{F, l=1}$ and for $l = 2$, we use $\frac{c_2}{c_1} \mathcal{F}_{\text{GLF}, b}^{F, l=1}$.

This becomes clearer when we plug in steady-state values for this setup, where $c_3 = c_2 = c$ and $c_1 = c_0 = \frac{c}{2}$, as we would use in the cosmological simulation. In this case, $\mathbf{F}_0 = \mathbf{F}_1 = \mathbf{F}_2 = \mathbf{F}_3 = \mathbf{F}$

and $N_1 = N_0 = N$ and $N_3 = N_2 = \frac{N}{2}$.¹¹ The updated state of a cell is computed as

$$\mathcal{U}_l^{n+1} = \mathcal{U}_l^n + \frac{\Delta t}{\Delta x} (\mathcal{F}_{\text{GLF},l-1/2}^n - \mathcal{F}_{\text{GLF},l+1/2}^n). \quad (\text{A14})$$

If we consider the photon density update for cell $l = 2$, we have computed $\mathcal{F}_{\text{GLF},l-1/2}^n$ in equation (A10) and $\mathcal{F}_{\text{GLF},l+1/2}^n$ is easily computed as

$$\mathcal{F}_{\text{GLF},e}^{N,l=2} = \frac{F_2 + F_3}{2} - \frac{c_2}{2} (N'_3 - N_2). \quad (\text{A15})$$

Plugging in the steady-state values in the equations for $\mathcal{F}_{\text{GLF},l-1/2}^n$ and $\mathcal{F}_{\text{GLF},l+1/2}^n$, we find the updated number density in cell $l = 2$ to be

$$\begin{aligned} \mathcal{U}_{l,N}^{n+1} &= \mathcal{U}_{l,N}^n + \frac{\Delta t}{\Delta x} \left(\frac{F_1 + F_2}{2} - \frac{c_2}{2} N_2 + \frac{c_1}{2} N_1 \right. \\ &\quad \left. - \frac{F_3 + F_2}{2} + \frac{c_2}{2} N'_3 - \frac{c_2}{2} N_2 \right) \\ &= \mathcal{U}_{l,N}^n + \frac{\Delta t}{\Delta x} \left(F - \frac{c}{2} \frac{N}{2} + \frac{c}{4} N - F + \frac{c}{2} \frac{N}{2} - \frac{c}{2} \frac{N}{2} \right) \\ &= \mathcal{U}_{l,N}^n. \end{aligned} \quad (\text{A16})$$

Similarly, for the flux, we have computed $\mathcal{F}_{\text{GLF},l-1/2}^n$ in equation (A12) and $\mathcal{F}_{\text{GLF},l+1/2}^n$ is computed as

$$\mathcal{F}_{\text{GLF},b}^{F,l=2} = \frac{c_2^2 N_2 + c_2^2 N'_1}{3} - \frac{c_2}{2} (F_3 - F_2). \quad (\text{A17})$$

Plugging into equation (A14), we find the updated flux in cell $l = 2$ as

$$\begin{aligned} \mathcal{U}_{l,F}^{n+1} &= \mathcal{U}_l^n + \frac{\Delta t}{\Delta x} \left(\frac{c_2^2 N_2 + c_2 c_1 N_1}{2} \right. \\ &\quad \left. - \frac{c_2}{2} (F_2 - F_1) - \frac{c_2^2 N_2 + c_2^2 N'_3}{2} + \frac{c}{2} (F_3 - F_2) \right) \\ &= \mathcal{U}_{l,F}^n + \frac{\Delta t}{\Delta x} \left(\frac{c^2 \frac{N}{2} + \frac{c^2}{2} N}{2} - \frac{c}{2} F + \frac{c}{2} F \right. \\ &\quad \left. - \frac{c^2 \frac{N}{2} + c^2 \frac{N}{2}}{2} - \frac{c}{2} F + \frac{c}{2} F \right) \\ &= \mathcal{U}_{l,F}^n. \end{aligned} \quad (\text{A18})$$

Clearly, the steady state is maintained, which demonstrates the validity of the VSLA algorithm. A similar exercise can easily be computed for the cell $l = 1$. Recall that when updating the flux, this is done asymmetrically depending on which side of the boundary one is on, which is key to maintaining the steady state.

We present here a series of idealized tests that verify that our new VSLA implementation exhibits the properties we have described.

A2.1 Test 1: point source in an optically thin box

In this test, we set up a low-resolution square box with sides of length 0.8 kpc and put a point source that emits 10^{50} photons s^{-1} at position (0, 0.4, 0.4) kpc such that it is emitting from the left-hand wall of the box. All walls have outflow boundary conditions so that any radiation that reaches the edges of the box is advected out

of the computational domain. Cells with $x < 0.15$ kpc are refined to Level 5 (32^3) and the rest are refined to Level 4 (16^3), and this refinement map is kept static throughout the test. The source luminosity is split into three bins with mean energies of 18.85, 35.05 and 65.67 eV, respectively, and the number of photons emitted in each bin is consistent with that of a 10^5 K blackbody spectrum. The box is entirely filled with hydrogen gas with a uniform density of $n_{\text{H}} = 0.1 \text{ cm}^{-3}$ and a temperature of 100 K. In this test, all hydrogen atom cross-sections are set to zero so that the photons travel unimpeded across the box as if the radiation is propagated in a vacuum.

We begin by running this test four times. In our first experiment, we set $c_{\text{sim}} = 0.004c$ everywhere in the box, regardless of the level of refinement, which represents the control run. With this setup, the effective crossing time of the box is $\sim 650\,000$ yr. The simulation is run for 4 Myr, much longer than the crossing time, so that the simulation reaches a steady state. In the next two experiments, we increase c_{sim} in the coarse cells (i.e. at $x \geq 0.15$ kpc) by a factor of 2 (our default VSLA implementation), by a factor of 10 and by a factor of 100 in order to demonstrate the conservation of flux and photon number density in all cases.

Fig. A2 shows the fractional deviation of the total photon flux ($\tilde{c}N_{\gamma}$) in the second three experiments compared with the reference control run ($|\text{Control} - \text{Experiment}|/\text{Control}$). The total photon flux is well conserved, regardless of the speed of light we use in the less refined region. In all cases, the luminosity of the source is the same, and therefore the flux, which depends only on luminosity and distance, should be the same, regardless of the value of the speed of light we choose in the simulation (assuming the simulation has reached a steady state). We have confirmed that we conserve photon flux in each individual direction to a precision of one part in $\sim 10^{-13}$ when a flux limiter is not imposed. When a second-order flux limiter¹² is used, the conservation changes to less than one part in 10^{-2} and is slightly different for the x -direction compared to y and z . This is almost certainly due to the geometry of the problem. The photon source is located in a position such that all photons inherit a positive flux in the x -direction, while this is not the case for both the y - and z -directions, and the refinement boundary is placed at a fixed x value. In the case of the VSLA, it may be useful to adopt a lower order slope limiter for the radiation while keeping it on for the hydro as we expect the radiation to be a much smoother field.

We also show in Fig. A3 how the photon number density changes across the boundary such that the photoionization rate remains constant. Here we see the fractional deviations expected for changing the speed of light by a factor of 2, 10 and 100, respectively.

Because the total photon flux is conserved across the boundary, the number density of photons in the cells is different by the same factor as the change in \tilde{c} . However, because the absorption cross-section is scaled with the speed of light, for a constant density medium, the same number of absorptions will occur even with a change in photon number density across the refinement regions. This is further shown in Appendix A2.2.

We have also checked to make sure that our implementation is robust to changes of the location of the source as well as the direction of the refinement boundary. We have run the same test with the exact reverse setup by putting the source at position (0.8, 0.4, 0.4) kpc and reversing the orientation of the refinement boundary, and obtained

¹¹ The change in N is so that the photoionization rate remains constant, which is what expected for a steady state.

¹² This determines how the cell properties are interpolated from coarse to fine cells and likewise in the reverse.




Figure A2. Fractional deviations of the VSLA run with a 2 (left-hand panel), 10 (middle panel) and 100 (right-hand panel) times faster speed of light in the coarse cells compared to the reference RSLA run. The snapshots are taken at 4 Myr, long after the crossing time of the box and at a point where the simulation should be in a steady state. The total photon flux is well conserved between all the runs. The vertical dashed line shows the location where the refinement level changes. The coordinates of the box have been recentred on (0.4, 0.4, 0.4) kpc for these plots.




Figure A3. Fractional deviations of the VSLA run with a 2 (left-hand panel), 10 (middle panel) and 100 (right-hand panel) times faster speed of light in the coarse cells compared to the reference RSLA run. The snapshots are taken at 4 Myr, long after the crossing time of the box and at a point where the simulation should be in a steady state. The number density of photons changes as expected in order to maintain a constant photoionization rate. The coordinates of the box have been recentred on (0.4, 0.4, 0.4) kpc for these plots.

Table A1. Wall clock time for test 1.

| c_{coarse} | Wall clock time (s) |
|---------------------|---------------------|
| 0.004c | 27.45 |
| 0.004c ^a | 33.77 |
| 0.008c | 32.80 |
| 0.04c | 155.30 |
| 0.4c | 1686.92 |

^aWe have run a second control run where we turn off adaptive time-stepping (i.e. n_{subcycle} is set to 1 for all levels).

very similar results. Furthermore, we have run the test both in the y-direction and in the z-direction and once again found similar results.

These tests have all been run on four cores, and we provide the timing results in Table A1. All tests have exactly the same domain decomposition and number of outputs. We have run an additional control where adaptive time-stepping is turned off (i.e. all levels are run on the fine time-step and thus there is no hydro sub-cycling), and

we see that the timing is the same for the run with the twice enhanced speed of light. Because this test is at such low resolution and we use more cores than necessary (this was done to test the robustness of our algorithm to MPI calls), the timing between the control with and without adaptive time-stepping is fairly comparable. In the run with a 10 times enhanced speed of light, the time-step is five times shorter than the control run, and we see that it takes 4.6 times as long to run. Likewise, the run with the 100 times enhanced speed of light has a time-step that is 50 times shorter than the control run, and this takes 55 times as long to run. Note that some of the time is spent on I/O, and the amount of time needed to print an output is independent of the value of the speed of light used. For the shorter runs, this becomes more important in timing than for longer runs.

A2.2 Test 2: point source in a low-density medium

The setup for this test is exactly the same as in the previous test, except that the H I cross-sections are no longer set to zero, rather the energy-weighted and photoionization cross-sections are set as

Table A2. Wall clock time for test 2. The control experiment has adaptive time-stepping turned off.

| c_{coarse} | Wall clock time (s) |
|---------------------|---------------------|
| 0.004c | 32.43 |
| 0.008c | 31.42 |
| 0.04c | 162.08 |
| 0.4c | 1423.24 |

appropriate for a 10^5 K blackbody spectrum. We have set the RAMSES-RT parameter ‘static=false.’ so that the hydro quantities can be updated. Furthermore, we decrease the density in the box to $n_{\text{H}} = 10^{-3} \text{ cm}^{-3}$ so that the Strömgen radius is much larger than the box size. In addition to the control where $c_{\text{sim}} = 0.004c$, we perform a VSLA run where we enhance c_{sim} by a factor of 2 at the coarse level. After 4 Myr, the box is completely ionized and optically thin to radiation. Flux in all directions is conserved to a few parts in 10^{-6} , and the maximal deviation in ionized fraction is

2×10^{-6} , which occurs right at the edge of the far side of the box and away from the source. Timings for this test can be found in Table A2.

In Fig. A4, we show the propagation of the I-front at 0.1, 0.2, 0.4, 0.9 and 4 Myr. The position of the I-front is very similar in the left-hand panels of Fig. A4 since the I-front is only just beginning to cross into the coarse cells where the speed of light differs. In the left-hand centre panel of the figure, we see a small mushroom effect where it appears that the radiation has expanded a bit further in the y-direction after the refinement boundary. This is due to the low resolution of the cells and the rendering that comes with averaging the quantities along the z-axis. The mushroom effect is enhanced in the lower left-hand panels because the speed of light is faster so the radiation should be diffusing quicker than in the control run. We can see in the four right-hand panels of Fig. A4 that the I-front of the control run is lagging behind the VSLA runs as expected. This is further tested in Appendix A4.1. After 4 Myr, the I-front is far beyond the regions of the box and the simulation is in a steady state, as described above.




Figure A4. Evolution of the I-front at 0.1, 0.2, 0.4, 0.9 and 4 Myr for the control run (top row) and VSLA runs (bottom three rows). The refinement boundary occurs at $x = -0.25$, where to the left-hand side of the boundary, the speed of light is set to be that of the control run, while to the right-hand side of the boundary, we vary the speed of light as indicated on the plot. Hence, the I-fronts travel faster when the speed of light is increased.

Table A3. Refinement scheme for Iliev 2009 test 6.

| Level | N_{cells} | Radius |
|-------|--------------------|-----------------------------|
| 6 | 64^3 | $r > 0.25$ kpc |
| 7 | 128^3 | $0.25 \geq r > 0.125$ kpc |
| 8 | 256^3 | $0.125 \geq r > 0.0625$ kpc |
| 9 | 512^3 | $0.0625 \geq r$ kpc |

A3 Additional tests

In addition to the previous two tests, we have checked that our algorithm is robust to adaptive refinement, both refining and de-referring cells on-the-fly. Furthermore, imposing multiple levels of refinement also poses no issues.

A4 VSLA comparison tests

Having demonstrated that the VSLA successfully conserves both the number of photons and the flux across level boundaries, we now test our new VSLA implementation in a series of controlled experiments that demonstrate the advantage of the VSLA over both the RSLA and using the full speed of light. We present three tests demonstrating that the physical properties of the I-front are conserved in regimes where the velocity of the I-front (the rate at which the radius H_{II} region grows) is much slower than the speed of light and that the VSLA better reproduces the results of the full speed of light in low-density regimes at a much reduced computational cost.

A4.1 Iliev 2009 test 6

Iliev et al. (2009) test 6 is designed to represent a radiative source emitting at the centre of a dense cloud of gas. A luminous source ($\dot{N}_\gamma = 5 \times 10^{50}$ photons s^{-1} summed over three radiation bins) mimicking a 10^5 K blackbody is placed at the corner of a box of length 0.8 kpc and centred on a dense cloud of gas with the following density profile:

$$n_H(r) = \begin{cases} n_0 & \text{if } r \leq r_0 \\ n_0(r_0/r)^2 & \text{if } r \geq r_0, \end{cases} \quad (\text{A19})$$




Figure A5. Left-hand panels: position (top) and velocity (bottom) of the I-front as a function of time for each of the three different simulations for Iliev test 6. Right-hand panels: position (top) and velocity (bottom) of the I-front as a function of time for each of the three different simulations for Iliev test 6 reversed. The fluctuations in velocity are due to sampling.

where $n_0 = 3.2 \text{ cm}^{-3}$ and $r_0 = 91.5 \text{ pc}$. The initial temperature of the entire box is set to 100 K and the simulation contains only hydrogen. The boundary conditions of the three walls touching the radiation source are reflective, while the other three are outflow. The original test uses a fixed 128^3 Cartesian mesh; however, since the VSLA attempts to keep the light-crossing time of all cells the same, regardless of resolution, we have modified the original test slightly to better explore the properties of VSLA. Rather than a fixed uniform mesh, we initialize a fixed grid as listed in Table A3, which uses a base grid of 64^3 cells with three additional levels of refinement at specified radii.

We provide three versions of this simulation: (1) $c_{\text{sim}} = 0.01c$; (2) $c_{\text{sim}} = 0.08c$; and (3) the VSLA implementation where $c_{\text{sim}} = 0.01c$ on Level 9 and is increased by a factor of 2 at each lower level up to $c_{\text{sim}} = 0.08c$ on Level 6. In this test, the velocity of the I-front is much slower than the speed of light, and we therefore do not expect any differences between the three different runs. This test is simply meant to demonstrate that in the regime where the I-front velocity is slow relative to the speed of light, our VSLA method reproduces what is expected. In the left-hand panel of Fig. A5, we show the position and velocity of the front as a function of time for the first ~ 15 Myr of the simulation. As expected, we see very good agreement between all three simulations, despite the change in the speed of light. In the bottom left-hand panel of Fig. A5, we see that the velocity of the I-front is much shorter than the values of the speed of light used in any of the three experiments.

A4.2 Iliev 2009 test 6 reversed

In our next test, we change the location of the radiation source such that it begins emitting in much lower density gas while irradiating a much higher density region, much like one galaxy shining on to another. The simulation is very similar to the setup in the previous test in that we use the same density profile for the gas, initial temperature, box size and source luminosity, but we move the radiation source from $(x, y, z) = (0, 0, 0)$ to $(x, y, z) = (0.8, 0.8, 0.8)$. The refinement scheme is exactly the same as in the previous test and is given in Table A3. Because the density profile is convex with respect to the source location, we measure the position of the I-front

with respect to the opposite corner and subtract the distance from the diagonal. We then use the change in distance to get a velocity.

Similar to the previous test, we run three versions of the simulation: (1) $c_{\text{sim}} = 0.01c$; (2) $c_{\text{sim}} = 0.08c$; and (3) the VSLA implementation where $c_{\text{sim}} = 0.01c$ on Level 9 and is increased by a factor of 2 at each lower level up to $c_{\text{sim}} = 0.08c$ on Level 6. The situation here is very different from the previous example because the radiation source begins emitting in a much lower density region. In this case, we expect the I-front to move considerably faster, much as it would in the low-density IGM, and thus the particular value we choose for c_{sim} will affect the propagation of the I-front.

In the right-hand panel of Fig. A5, we see that the VSLA run agrees well with the $c_{\text{sim}} = 0.08c$ simulation in both the position of the I-front and its velocity while the $c_{\text{sim}} = 0.01c$ run lags behind. This is expected since the VSLA has $c_{\text{sim}} = 0.08c$ at Level 6, where much of the initial propagation takes place. When the I-front reaches the higher density region of the box, it slows down considerably to a value much less than the reduced speed of light that allows the I-front in run with $c_{\text{sim}} = 0.01c$ to catch up with the other two runs.

What this test reveals is that using the VSLA technique to set the speed of light to always be slightly faster than the expected velocity of the I-front will save significantly on computational cost while also reproducing the correct behaviour of the I-front expansion. Comparing the timings of this run, we find that the runs with $c_{\text{sim}} = 0.01c$, $c_{\text{sim}} = 0.08c$ and VSLA took 4319.16, 34 425.47 and 14 479.48 s, respectively. As expected, the run with $c_{\text{sim}} = 0.01c$ is nearly exactly eight times faster than the run with $c_{\text{sim}} = 0.08c$. Abandoning the multistepping costs an extra factor of 3.35 over the run with $c_{\text{sim}} = 0.01c$; however, the VSLA run is still 2.38 times faster than the run with $c_{\text{sim}} = 0.08c$, despite the fact that the evolution of the I-front is indistinguishable between the two runs.

A4.3 Low-resolution cosmological simulation

Our final test aims to see how the VSLA algorithm performs when used in a low-resolution cosmological simulation. In order to do this, we set up a cosmological box of size 230^3 Mpc with a base grid on Level 8 with 256^3 dark matter particles and allow for refinements in the grid up to Level 13. This gives a maximum physical resolution of ~ 3.5 kpc at $z = 7$. We run three simulations: a ‘Full c’ run where $c_{\text{sim}} = c$ at all levels, a VSLA run where we set $c_{\text{sim}} = c$ on Level 8 and divide it by a factor of 2 up to Level 13, where $c_{\text{sim}} = 0.03125c$, and a final RSLA run where $c_{\text{sim}} = 0.03125c$ on all levels. We set the star formation and escape fraction of radiation from star particles to be the same for all three simulations and run the simulation to $z = 7$. The crossing time of light across the box at this redshift is ~ 3 Gyr for the RSLA run, while it is ~ 94 Myr for the ‘Full c’ run. The former is much longer than the age of the Universe at this redshift, while the latter is much less than the age of the Universe. In a box with only a few strong and clustered sources, as we have in this simulation, one would expect that the RSLA simulations lag far behind the ‘Full c’ simulation.

In Fig. A6, we plot the volume-weighted H II filling factor as a function of redshift. We can see that the VSLA simulation tracks the ‘Full c’ run beautifully, and there is near-perfect agreement. Both these simulations have the entire box completely ionized at $z \sim 7.5$. By contrast, the RSLA simulation is only a few per cent ionized at these redshifts.

We should emphasize that the setup of this simulation is not completely realistic and designed to exploit the potential of VSLA compared to RSLA in terms of agreement with the ‘Full c’ run. The star particles that form and ionize the box are strongly clustered




Figure A6. Volume-filling factor of H II as a function of redshift in a 230^3 Mpc box using ‘Full c’, VSLA and RSLA. As this simulation is at extremely low resolution, only massive haloes form with extremely massive and luminous star particles. This extreme choice is meant to mimic a quasar-dominated reionization scenario, and in this case, the VSLA works extremely well.

around a few different locations, and only a small fraction of cells are refined to the maximum level. What potentially matters for this is how the sources are distributed throughout the box. If there are many weaker sources equally distributed around the volume, then the RSLA run will not lag very far behind the ‘Full c’ run in terms of redshift of reionization. This might correspond to a case where mini-haloes are the dominant sources of photons during reionization. The case we have just simulated here where there are a few very strong sources around the box exacerbates the problem of using the RSLA because the time until the overlap of the Strömgren spheres is much longer. We know that the RSLA lags behind until the Strömgren radius is reached. In this type of simulation, the Strömgren radii are very large for these sources, which makes the lag in the RSLA run worse. This might correspond to a physical scenario where quasars are the dominant sources of photons during reionization, and it is well established that for these very bright sources, the I-fronts never reach their Strömgren radius (Shapiro & Giroux 1987).

Comparing the timing between the two runs, the RSLA run took 2283 s to reach $z = 7$, while it took 32 852 and 42 271 s for the VSLA and ‘Full c’ runs, respectively, on 64 cores. In this case, the VSLA run takes ~ 13.7 times longer than the RSLA, while the ‘Full c’ run takes ~ 17.7 times longer. We can ask why there is such a difference in the timings? In theory, the ‘Full c’ run should be 32 times faster than the RSLA run. However, there is a time period before the first star particle forms when the on-the-fly RT is turned off, where the simulations take the same amount of time. Since we have run the simulations for only a very short period of time, this eats into how much speed-up one can achieve just by using the RSLA. Running the simulation for longer should lead to a better speed-up for the RSLA. Furthermore, we have not made an attempt to optimize the number of cores for the computation. Looking towards the VSLA, this factor is much higher than one would naively expect, given our previous tests. We can expect that abandoning multistepping in time costs roughly a factor of ~ 3.5 in runtime, as observed in the previous section. We still have to account for another factor of 4 in runtime. This is due to the fact that star particles form before the maximum level of resolution is reached. If this occurs, the standard version of VSLA will be a factor of 2 slower than the RSLA for every lower resolution level above the maximum level that star particles form. This is because we have set the speed of light to divide by

a factor of 2 on each subsequent level. In this specific very low resolution simulation, star particles start forming on Levels 10 and 11 before the simulation has refined down to Level 13. This can slow down the VSLA considerably. One can mitigate this effect in a few different ways. One might enforce that star particles can form only when the maximum level of the simulation is reached, thereby preventing the slowdown. This solution is potentially useful when running simulations that refine many levels and, in particular, zoom simulations. Alternatively, one could change the algorithm slightly so that the highest resolved cells always have the same speed of light as the RSLA run, and as more levels are introduced into the simulation, the speed of light increases at the lower resolution levels. As an example, in this specific case, one could start the VSLA run with $c_{\text{sim}} = 0.03125c$ on Level 8. When the simulation refines to Level 9, we would set $c_{\text{sim}} = 0.03125c$ on Level 9 and $c_{\text{sim}} = 0.0625c$

on Level 8. Likewise, when the simulation refines to Level 10, we would set $c_{\text{sim}} = 0.03125c$ on Level 10, $c_{\text{sim}} = 0.0625c$ on Level 9 and $c_{\text{sim}} = 0.0125c$ on Level 8. If one allows about seven levels of refinement, by the time the simulation reaches the maximum level, we will be back in the original form of the VSLA algorithm. If one sets the simulation to always resolve the same physical scale (i.e. releasing a new level when the scalefactor increases by a factor of 2), then this latter algorithm should perform very similarly to the original VSLA and also have the added bonus of a significant speed-up compared to the original algorithm. All in all, the way to take advantage of the VSLA is to configure it to the problem of interest, and there is no one-size-fits-all method.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.