

21CMFAST: a fast, seminumerical simulation of the high-redshift 21-cm signal

Andrei Mesinger,¹*† Steven Furlanetto² and Renyue Cen¹

¹*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*

²*Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA*

Accepted 2010 September 16. Received 2010 September 15; in original form 2010 March 19

ABSTRACT

We introduce a powerful seminumeric modelling tool, 21CMFAST, designed to efficiently simulate the cosmological 21-cm signal. Our code generates 3D realizations of evolved density, ionization, peculiar velocity and spin temperature fields, which it then combines to compute the 21-cm brightness temperature. Although the physical processes are treated with approximate methods, we compare our results to a state-of-the-art large-scale hydrodynamic simulation, and find good agreement on scales pertinent to the upcoming observations ($\gtrsim 1$ Mpc). The power spectra from 21CMFAST agree with those generated from the numerical simulation to within 10s of per cent, down to the Nyquist frequency. We show results from a 1-Gpc simulation which tracks the cosmic 21-cm signal down from $z = 250$, highlighting the various interesting epochs. Depending on the desired resolution, 21CMFAST can compute a redshift realization on a single processor in just a few minutes. Our code is fast, efficient, customizable and publicly available, making it a useful tool for 21-cm parameter studies.

Key words: methods: numerical – galaxies: high redshift – dark ages, reionization, first stars – diffuse radiation – early Universe – large-scale structure of Universe.

1 INTRODUCTION

Through challenging observational efforts, the high-redshift frontier has been incrementally pushed back in recent years. Glimpses of the $z \sim 6$ –8 Universe were provided by quasars (e.g. Fan et al. 2006), candidate Lyman break galaxies (e.g. Bouwens et al. 2008, 2009; McLure et al. 2009; Ouchi et al. 2009), Lyman alpha emitters (LAEs; e.g. Kashikawa et al. 2006; Shimasaku et al. 2006) and GRBs (e.g. Cusumano et al. 2006; Greiner et al. 2009; Salvaterra et al. 2009). Unfortunately, these precious observations currently provide only a limited set of relatively bright objects. Luckily, we will soon be inundated with observations probing this and even earlier epochs. These observations should include infrared spectra from the *James Webb Space Telescope* (JWST), the Thirty Meter Telescope (TMT), the Giant Magellan Telescope (GMT), the European Extremely Large Telescope (E-ELT), wide-field LAE surveys from the Subaru HyperSupremeCam, as well as the E-mode CMB polarization power spectrum measured by the *Planck* satellite. Some of the most important information will come in the form of the redshifted 21-cm line of neutral hydrogen. Several interferometers will attempt to observe the cosmological 21-cm signal, including the Mileura Wide Field Array (MWA; Bowman, Morales & Hewitt

2005),¹ the Low Frequency Array (LOFAR; Harker et al. 2010),² the Giant Metrewave Radio Telescope (GMRT; Pen et al. 2008), the Precision Array to Probe the Epoch of Reionization (PAPER; Parsons et al. 2009) and eventually the Square Kilometer Array (SKA).³ The first generation of these instruments, most notably LOFAR and MWA, are not only scheduled to become operational within a year, but should also yield insight into the 3D distribution of HI, provided the systematics can be overcome (see the recent reviews of Furlanetto, Oh & Briggs 2006a; Morales & Wyithe 2009).

However, interpreting this data will be quite challenging and no doubt controversial, as foreshadowed by the confusion surrounding the scant, currently available observations. There are two main challenges to overcome: (1) an extremely large parameter space, due to our poor understanding of the high-redshift Universe; (2) an enormous dynamic range (i.e. range of relevant scales).

Theoretically, the dawn of the first astrophysical objects and reionization could be modelled from first principles using numerical simulations, which include the complex interplay of many physical processes. In practice however, simulating these epochs requires enormous simulation boxes. Gigaparsec scales are necessary to statistically model ionized regions and absorption systems or create accurate mock spectra from the very rare high-redshift quasars.

*E-mail: mesinger@astro.princeton.edu

†Hubble Fellow.

¹ <http://web.haystack.mit.edu/arrays/MWA/>

² <http://www.lofar.org>

³ <http://www.skatelescope.org/>

However, the simulations also require high enough resolution to resolve the underlying sources and sinks of ionizing photons and the complex small-scale feedback mechanisms which regulate them. Thus one is forced to make compromises: deciding which physical processes can be ignored, and how the others can be parametrized and efficiently folded into large-scale models. Furthermore, large-scale simulations are computationally costly (even when they sacrifice completeness for speed by ignoring hydrodynamic processes) and thus are inefficient in large parameter studies.

On the other hand, analytic models, while more approximate, are fast and can provide physical insight into the import of various processes. However, analytical models are hard-pressed to go beyond the linear regime, and beyond making fairly simple predictions such as the mean 21-cm signal (Furlanetto 2006), the probability density function (PDF; Furlanetto, Zaldarriaga & Hernquist 2004a) and power spectrum (Pritchard & Furlanetto 2007; Barkana 2009). The 21-cm tomographic signal should be rich in information, accommodating many additional, higher-order statistical probes, such as the bispectrum (Pritchard et al., in preparation), the difference PDF (Barkana & Loeb 2008), etc.

In this paper, we follow a path of compromise, attempting to preserve the most useful elements of both analytic and numeric approaches. We introduce a self-consistent, seminumerical⁴ simulation, specifically optimized to predict the high-redshift 21-cm signal. Through a combination of the excursion-set formalism and perturbation theory, our code can generate full 3D realizations of the density, ionization, velocity, spin temperature and ultimately 21-cm brightness temperature fields. Although the physical processes are treated with approximate methods, our results agree well with a state-of-the-art hydrodynamic simulation of reionization. However, unlike numerical simulations, realizations are computationally cheap and can be generated in a matter of minutes on a single processor, with modest memory requirements. Most importantly, our code is publicly available at <http://www.astro.princeton.edu/~mesinger/Sim.html>. We name our simulation *21CMFAST*.

Seminumerical approaches have already proved invaluable in reionization studies (Zahn et al. 2005; Mesinger & Furlanetto 2007; Geil & Wyithe 2008; Alvarez et al. 2009; Choudhury, Haehnelt & Regan 2009; Thomas et al. 2009). Indeed, *21CMFAST* is a more specialized version of our previous code, *DEXM* (Mesinger & Furlanetto 2007; hereafter MF07). The difference between the two is that *21CMFAST* bypasses the halo finding algorithm, resulting in a faster code with a larger dynamic range and more modest memory requirements. In this work, we also introduce some new additions to our code, mainly to compute the spin temperature.

In Section 2, we compare predictions from *21CMFAST* with those from hydrodynamic simulations of the various physical components comprising the 21-cm signal in the post-heating regime. Density, ionization, peculiar velocity gradient and full 21-cm brightness temperature fields are explored in Sections 2.1, 2.2, 2.3, 2.4, respectively. In Section 3, we introduce our method for computing the spin temperature fields, with results from the complete calculation (including the spin temperature) presented in Section 3.3. Finally in Section 4, we summarize our findings.

Unless stated otherwise, we quote all quantities in comoving units. We adopt the background cosmological parameters $(\Omega_A, \Omega_M, \Omega_b, n, \sigma_8, H_0) = (0.72, 0.28, 0.046, 0.96, 0.82, 70 \text{ km s}^{-1} \text{ Mpc}^{-1})$,

⁴ By ‘seminumerical’ we mean using more approximate physics than numerical simulations, but capable of independently generating 3D realizations.

matching the 5-yr results of the *WMAP* satellite (Komatsu et al. 2009).

2 21-cm TEMPERATURE FLUCTUATIONS POST-HEATING ($T_S \gg T_\gamma$)

Our ultimate goal is to compute the 21-cm background, which requires a number of physics components. To identify them, note that the offset of the 21-cm brightness temperature from the CMB temperature, T_γ , along a line of sight (LOS) at observed frequency ν , can be written as (cf. Furlanetto et al. 2006a):

$$\begin{aligned} \delta T_b(\nu) &= \frac{T_S - T_\gamma}{1+z} (1 - e^{-\tau_{\nu_0}}) \\ &\approx 27 x_{\text{H}1} (1 + \delta_{\text{nl}}) \left(\frac{H}{dv_r/dr + H} \right) \left(1 - \frac{T_\gamma}{T_S} \right) \\ &\quad \times \left(\frac{1+z}{10} \frac{0.15}{\Omega_M h^2} \right)^{1/2} \left(\frac{\Omega_b h^2}{0.023} \right) \text{ mK}, \end{aligned} \quad (1)$$

where T_S is the gas spin temperature; τ_{ν_0} is the optical depth at the 21-cm frequency ν_0 ; $\delta_{\text{nl}}(\mathbf{x}, z) \equiv \rho/\bar{\rho} - 1$ is the evolved (Eulerian) density contrast; $H(z)$ is the Hubble parameter; dv_r/dr is the comoving gradient of the LOS component of the comoving velocity; and all quantities are evaluated at redshift $z = \nu_0/\nu - 1$. The final approximation makes the assumption that $dv_r/dr \ll H$, which is generally true for the pertinent redshifts and scales, though we shall return to this issue in Section 2.3.

In this comparison section, we make the standard, simplifying assumption of working in the post-heating regime: $T_S \gg T_\gamma$. For fiducial astrophysical models, this is likely a safe assumption during the bulk of reionization (Furlanetto 2006; Chen & Miralda-Escudé 2008; Santos et al. 2008; Baek et al. 2009). We will however revisit this assumption in Section 3, where we introduce our method for computing the spin temperature field.

The remaining components of equation 1 are the density, δ_{nl} , the ionization, $x_{\text{H}1}$, and the velocity gradient, dv_r/dr . Below, we study these in turn, comparing *21CMFAST* to the hydrodynamic cosmological simulation of Trac, Cen & Loeb (2008), *using the same initial conditions (ICs)*. We perform ‘by-eye’ comparisons at various redshifts/stages of reionization, as well as one- and two-point statistics: the PDFs (smoothed on several scales), and the power spectra. Since our code is designed to simulate the cosmological 21-cm signal from neutral hydrogen, we study the regime before the likely completion of reionization, $z \gtrsim 7$ (though the present data are even consistent with reionization completing at $z \lesssim 6$; Lidz et al. 2007; Mesinger 2009).

The simulations of Trac et al. (2008) are the current ‘state-of-the-art’ reionization simulations. They include simultaneous treatment of dark matter (DM) and gas, five-frequency radiative transfer (RT) on a 512^3 grid, and they resolve $M_{\text{halo}} \gtrsim 10^8 M_\odot$ ionizing sources with $\gtrsim 40$ DM particles in a 143-Mpc box.

2.1 Evolved density field

We calculate the evolved density field in the same fashion as in the ‘parent’ code, *DEXM*, outlined in MF07. In short, we generate density and velocity ICs in initial (Lagrangian) space, in roughly the same manner as numerical cosmological simulations. We then approximate gravitational collapse by moving each initial matter particle (whose mass is the total mass in the corresponding IC cell) according to first-order perturbation theory (Zel’Dovich 1970). First-order perturbation theory is very computationally convenient,




Figure 1. Slices through the density field, $\Delta(\mathbf{x}, z)$, for the gas; DM, linearly extrapolated ICs; and 21CMFAST, at $z = 7$ (clockwise from top left). Each slice is 143 Mpc on a side and 0.19 Mpc thick.

as the displacement field is a separable function of space and time, so the spatial component need only be computed once for each realization/box. There is no separate treatment of baryons and DM. Readers interested in more details concerning this approach are encouraged to check MF07.

This approach to generating large-scale density fields was also adopted by Choudhury et al. (2009) and Santos et al. (2009), who briefly showed that the resulting fields at high z traced the DM distribution from an N -body code fairly well. Here we perform more extensive comparisons. The Lagrangian space ICs used for 21CMFAST were initialized at $z = 300$ on a 1536^3 grid. The velocity fields used to perturb the ICs, as well as the resulting density fields presented below are 768^3 .⁵ We show results from both 768^3 and 256^3 boxes. We

highlight here that it takes ~ 10 min to generate the 768^3 21CMFAST density field from the 1536^3 ICs at $z = 7$ on a single-processor Mac Pro desktop computer. To put this into perspective, a hydrodynamical simulation of this single realization would take approximately 3 d to run down to $z = 7$ on a 1536-node supercomputing cluster.

In Fig. 1, we show a slice through the evolved density field, $\Delta(\mathbf{x}, z) \equiv 1 + \delta_{\text{nl}}$, at $z = 7$. Density fields computed from the gas, DM, linearly extrapolated ICs, and 21CMFAST are shown clockwise from the top left panel. It is evident that the Zel'Dovich approximation works quite well for this purpose, accurately reproducing the cosmic web. We do not capture baryonic physics, and so the 21CMFAST output looks more similar to the DM than the gas.

⁵ Our code allows the ICs to be sampled on to a high-resolution, HI-RES-DIM³, grid, while the subsequent evolved density, ionization, etc. fields can be created at lower resolution, LOW-RES-DIM³. This allows efficient use of available memory, with the code storing at most HI-RES-DIM³ + $4 \times$ LOW-RES-DIM³ floating point numbers in RAM. However, the Zel'Dovich perturbation approach, just as numerical simulations, requires

that the evolved fields are adequately resolved by the high-resolution grid. Failure to do so can substantially underestimate the fluctuations in the density field. We roughly find that the high-resolution grid should have cell sizes $\lesssim 1$ Mpc to accurately model density fields at redshifts of interest ($z \lesssim 40$). For larger cell sizes, $\gtrsim 10$ Mpc, the linear theory evolution option of 21CMFAST should be used.

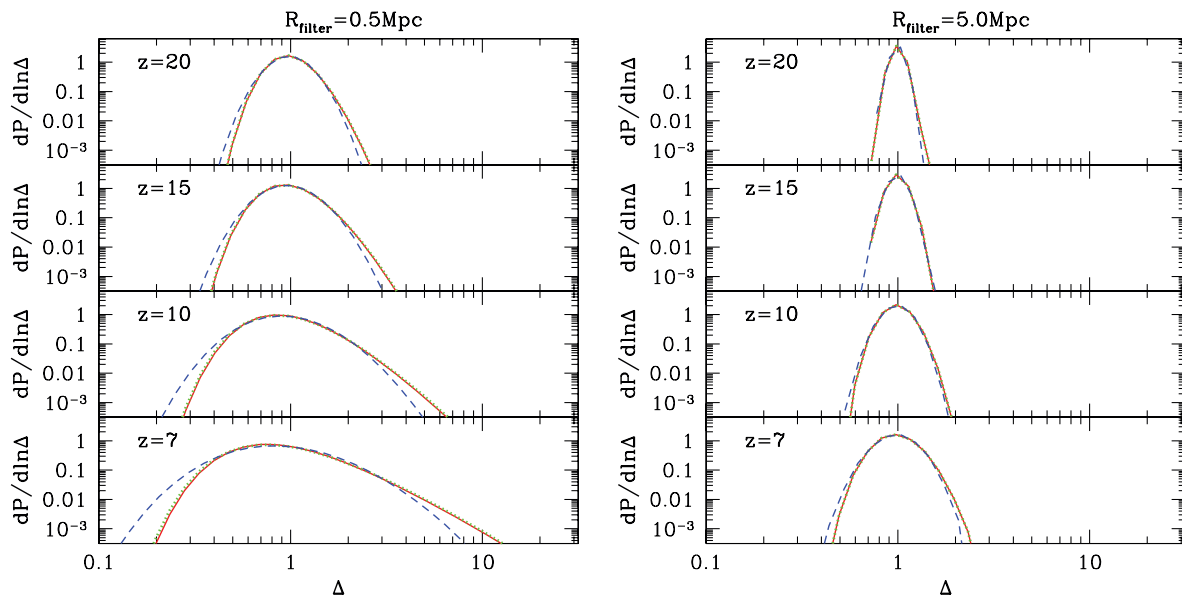


Figure 2. PDFs of the density fields smoothed on scale R_{filter} , computed from the gas (solid red curves), DM (dotted green curves) and 21CMFAST (dashed blue curves) fields, at $z = 20, 15, 10, 7$ (top to bottom). The left-hand panel corresponds to $R_{\text{filter}} = 0.5 \text{ Mpc}$; the right-hand panel corresponds to $R_{\text{filter}} = 5 \text{ Mpc}$. All smoothing was performed with a real-space, top-hat filter.

However, hydrodynamics is not included even in most of the present-day cosmological ($\gtrsim 100 \text{ Mpc}$) simulations (e.g. Iliev et al. 2006; McQuinn et al. 2007; see the recent review in Trac & Gnedin 2009).

Also, note that the linear density field is only accurate on large scales ($\gtrsim 10 \text{ Mpc}$ at $z \sim 7$). Thus care should be taken in applying tools which rely on the linear density field (such as the standard excursion set formalism) at smaller scales. Nevertheless, we include in 21CMFAST a feature to evolve the density field using fully linear evolution, instead of the perturbation approach discussed above. This allows one to generate extremely large boxes at low resolution. When using this feature, one should make sure that the chosen cell size is indeed in the linear regime at the redshift of interest. Some results making use of this feature are presented below.

In Fig. 2, we show the PDFs of the density fields smoothed on scale R_{filter} , computed from the gas (solid red curves), DM (dotted green curves) and 21CMFAST (dashed blue curves) fields, at $z = 20, 15, 10, 7$ (top to bottom). From the left-hand panel ($R_{\text{filter}} = 0.5 \text{ Mpc}$), we see that as structure formation progresses, we tend to increasingly overpredict the abundance of small-scale underdensities, and underpredict the abundance of large-scale overdensities. However, even at $z = 7$, our PDFs are accurate at the per cent level to over a dex around the mean density. Understandably, the agreement between the PDFs becomes better with increasing scale (see the right-hand panel corresponding to $R_{\text{filter}} = 5 \text{ Mpc}$). Interestingly, the DM distributions match the gas quite well, although this is somewhat of a coincidence, as we shall see from the power spectra below.

In Fig. 3, we present the density power spectra, defined as $\Delta_{\delta\delta}^2(k, z) = k^3 / (2\pi^2 V) \langle |\delta(k, z)|^2 \rangle_k$. The solid red, dotted green and dashed blue curves correspond to the gas, DM and 21CMFAST fields, respectively. On small scales ($k \gtrsim 5 \text{ Mpc}^{-1}$), the three fields have different power. The collapse of gas is initially delayed with respect to the pressureless DM, resulting in less small-scale power. The perturbation theory approach of 21CMFAST is closer in spirit to the DM evolution, but does not capture virialized structure. In fact the close agreement at $z = 7$ between the gas and 21CMFAST den-




Figure 3. Density power spectra, $\Delta_{\delta\delta}^2(k, z)$, of the gas (solid red curves), DM (dotted green curves) and 21CMFAST (dashed blue curves) fields, at $z = 20, 15, 10, 7$ (top to bottom).

sity power spectra is a coincidence, with the small-scale flattening of the 21CMFAST power attributable to ‘shell-crossing’ by the matter particles in the Zel’Dovich approximation. During reionization, the evolution of the gas is very complicated, since the power spectrum on small scales is sensitive to the thermal history of the reionization model (e.g. Hui & Gnedin 1997).

On large scales ($k \lesssim 0.5 \text{ Mpc}^{-1}$), all three power spectra agree remarkably at all epochs. To put this into perspective, neither the MWA nor LOFAR has sufficient signal-to-noise to detect the 21-cm signal beyond $k \gtrsim 2 \text{ Mpc}^{-1}$ (e.g. McQuinn et al. 2006).

2.2 Ionization field

We use a new, refined, seminumeric algorithm, FFRT, presented in Zahn et al. (2010) to generate ionization fields, $x_{\text{H I}}(\mathbf{x}, z)$. The ionization fields have been exhaustively compared against cosmological RT codes in Zahn et al. (2010), yielding a good agreement across a broad range of statistical diagnostics on moderate to large scales. Thus, we will not present further tests here. Instead, we merely outline the procedure, and motivate some aspects with regards to the goals of 21CMFAST: speed and efficiency.

We use the excursion-set approach for identifying H II regions, pioneered by Furlanetto, Zaldarriaga & Hernquist (2004b). The foundation of this approach is to require that the number of ionizing photons inside a region be larger than the number of hydrogen atoms it contains. Then ionized regions are identified via an excursion-set approach, starting at large scales and progressing to small scales, analogous to the derivation of the Press–Schechter (PS) mass functions (Bond et al. 1991; Lacey & Cole 1993).

Specifically, we flag fully ionized cells in our box as those which meet the criteria $f_{\text{coll}}(\mathbf{x}, z, R) \geq \zeta^{-1}$, where ζ is some efficiency parameter and $f_{\text{coll}}(\mathbf{x}, z, R)$ is the collapse fraction smoothed on decreasing scales, starting from a maximum R_{max} and going down to the cell size, R_{cell} . Additionally, we allow for partially ionized cells by setting the cell’s ionized fraction to $\zeta f_{\text{coll}}(\mathbf{x}, z, R_{\text{cell}})$ at the last filter step for those cells which are not fully ionized.⁶ The ionizing photon horizon, R_{max} , is a free parameter which can be chosen to match the extrapolated ionizing photon mean free path, in the ionized IGM, at $z \sim 7\text{--}10$ (e.g. Storrie-Lombardi et al. 1994; Miralda-Escudé 2003; Choudhury, Ferrara & Gallerani 2008). The photon sinks dominating the mean free path of ionizing photons are likely too small to be resolved in reionization simulations. An effective horizon due to photon sinks can delay the completion of reionization (e.g. Choudhury et al. 2009; Furlanetto & Mesinger 2009), and cause a drop in large-scale 21-cm power, as we shall see below.

There are two main differences between FFRT used in 21CMFAST and the previous incarnation of our H II bubble finder used in DEXM (MF07): (1) the use of the halo finder to generate ionization fields in DEXM (MF07) versus using the evolved density field and conditional PS to generate ionization fields in 21CMFAST; and (2) the bubble-flagging algorithm, which in MF07 is taken to paint the entire spherical region enclosed by the filter as ionized (‘flagging-the-entire-sphere’), whereas for 21CMFAST we just flag the central cell as ionized (‘flagging-the-central-cell’; for more information, see MF07; Zahn et al. 2007; and the appendix in Zahn et al. (2010)). These default settings of 21CMFAST were chosen to maximize speed and dynamic range, while minimizing the memory requirements. Nevertheless, they are left as user-adjustable options in the codes.

The first difference noted above means that 21CMFAST *does not explicitly resolve source haloes*. In MF07, we made use of a seminumerically generated halo field, which accurately reproduces N -body halo fields down to non-linear scales (MF07; Mesinger et al., in preparation). As in numerical reionization simulations, these haloes

were assumed to host ionizing sources. However, the intermediate step of generating such halo source fields adds additional computation time, and generally requires many GB of RAM for typical cosmological uses. Also as in numerical simulations, this memory requirement means that simulation boxes are limited to $\lesssim 200$ Mpc if they wish to resolve atomically cooled haloes at $z = 7\text{--}10$, and even smaller sizes if they wish to resolve these at higher redshifts or resolve molecularly cooled haloes. Although DEXM’s halo finder is much faster than N -body codes, and can generate halo fields at a given redshift in a few hours on a single processor, extending the dynamic range even further without hundreds of GB of RAM would be very useful. Alternatives to extending the dynamic range have been proposed by McQuinn et al. (2007); Santos et al. (2009). These involve stochastically populating cells with haloes below the resolution threshold. Although computationally efficient, it is unclear if these alternatives preserve higher-order statistics of the non-Gaussian $f_{\text{coll}}(\mathbf{x}, z, R)$ field, as each cell is treated independently of the others. Furthermore, such prescriptions require generating merger trees to self-consistently track the subgrid halo populations throughout their evolution (McQuinn et al. 2007). Otherwise the evolution of the shape of the subgrid mass function due to halo mergers is not preserved, as is the case in Santos et al. (2009).

Therefore, to increase the speed and dynamic range, we use the FFRT algorithm, which uses the conditional PS formalism (Lacey & Cole 1993; Somerville & Kolatt 1999) to generate the collapsed mass (i.e. ionizing source) field. Since conditional PS operates directly on the density field, without needing to resolve haloes, one can have an enormous dynamic range with a relatively small loss in accuracy. Most importantly, when computing $f_{\text{coll}}(\mathbf{x}, z, R)$ we use the *non-linear* density field, generated according to Section 2.1, instead of the standard linear density field. The resulting ionization fields are a much better match to RT simulations than those generated from the linear density field (Zahn et al. 2010; foreshadowed also by the ICs panel in Fig. 1 above). We normalize the resulting collapsed mass field to match the Sheth–Tormen (ST) mean collapse fraction, which in turn matches numerical simulations (see equation 14 and associated discussion).

The other major difference is that by default, FFRT in 21CMFAST flags just the central filter cell, instead of the entire sphere enclosed by the filter, as in MF07. The main motivation for this switch is that the former algorithm is $\mathcal{O}(N)$, while the latter is slower: $\mathcal{O}(N)$ at $\bar{x}_{\text{H I}} \sim 1$ but approaching $\mathcal{O}(N^2)$ as $\bar{x}_{\text{H I}} \rightarrow 0$. There are some other minor differences between the FFRT and the ionization scheme in MF07, such as the use of a sharp k -space filter instead of a spherical top-hat, but these have a smaller impact on the resulting ionization maps.

2.3 Peculiar velocity gradient field

Redshift space distortions, accounted for with the dv_r/dr term in equation (1),⁷ are often ignored when simulating the 21-cm signal. In the linear regime, redshift space distortions of the 21-cm field are similar to the well-studied Kaiser effect (e.g. Kaiser 1987), and the power spectrum of fluctuations is enhanced on all scales by a geometric factor of 1.87 (Bharadwaj & Ali 2004; Barkana &

⁶ Our algorithm can also optionally account for Poisson fluctuations in the halo number, when the mean collapse fraction becomes small, $f_{\text{coll}}(\mathbf{x}, z, R_{\text{cell}}) \times M_{\text{cell}} < 50M_{\text{min}}$, where M_{cell} is the total mass within the cell and our faintest ionizing sources correspond to a halo mass of M_{min} . This last step is found to be somewhat important when the cell size increases to $\gtrsim 1$ Mpc (see the appendix in Zahn et al. 2010). This is left as an option since turning off such stochastic behaviour allows the user to better track the deterministic redshift evolution of a single realization.

⁷ Note that this expression is exact, as long as the dv_r/dr field is constant over the width of the 21-cm line and $dv_r/dr \ll H(z)$. Alternatively, one can apply redshift space distortions when converting the comoving signal from the simulation box to an observed frequency. However, for the sake of consistency, we perform all of our calculations in comoving space.




Figure 4. PDFs of the comoving LOS derivative of v_r [in units of $H(z)$], smoothed on scale $R_{\text{filter}} = 0.5 \text{ Mpc}$ (left) and 5.0 Mpc (right). Solid red curves are generated from the hydrodynamic simulation, while the dashed blue curves are generated by 21CMFAST . Redshifts corresponding to $z = 20, 15, 10, 7$ are shown from top to bottom. All smoothing was performed with a real-space, top-hat filter. The dotted magenta curves were generated on comparable scales by 21CMFAST with different initial conditions; however, they assume linear evolution of the density field, instead of the perturbation theory approach (see Section 2.1).

Loeb 2005a). However, the small-scale overdensities, where redshift space distortions are most important, are also the regions whose 21-cm emission is first erased by ‘inside-out’ reionization. Preliminary studies therefore concluded that redshift space distortions would only be noticeable before reionization and in its early stages (McQuinn et al. 2006, MF07). As we are interested in accurately simulating the 21-cm signal from all cosmological epochs, including pre-reionization, here we will compare the velocity gradient term from 21CMFAST and hydrodynamic simulations.

Using the Zel’Dovich approximation on our 3D realizations, we can again efficiently move beyond the linear regime into the quasi-linear regime, and take into account correlations in the velocity gradient field. In this first-order perturbation theory, the velocity field can be written as

$$\mathbf{v}(\mathbf{k}, z) = \frac{i\mathbf{k}}{k^2} \dot{D}(z) \delta(\mathbf{k}), \quad (2)$$

and so the derivative of the LOS velocity, v_r , where \mathbf{r} for simplicity is oriented along a basis vector, can be written in k -space as

$$\frac{dv_r}{dr}(\mathbf{k}, z) = ik_r v_r(\mathbf{k}, z) \quad (3)$$

$$\approx -\frac{k_r^2}{k^2} \dot{D}(z) \delta_{\text{nl}}(\mathbf{k}), \quad (4)$$

where differentiation is performed in k -space. The last approximation is used for 21CMFAST , while the first, exact expression is used for the numerical simulation.⁸

In Fig. 4 we show the PDFs of the comoving LOS derivative of v_r [in units of $H(z)$], smoothed on scale $R_{\text{filter}} = 0.5 \text{ Mpc}$ (left)

and 5.0 Mpc (right). Solid red curves are generated from the hydrodynamic simulation, while the dashed blue curves are generated by 21CMFAST . Redshifts corresponding to $z = 20, 15, 10, 7$ are shown top to bottom. We see that our perturbation theory approach again does remarkably well in reproducing results from the hydrodynamic simulation. The velocity gradients agree even better than the density fields, since the velocity field is coherent over larger scales. The shape of the distributions is noticeably non-linear on small scales and late times. The curves resemble PDFs of the sign-flipped non-linear density field, δ_{nl} , which is understandable from equation (4). The dotted magenta curves in the bottom right-hand panels were generated on comparable scales by 21CMFAST with different initial conditions; however, they assume linear evolution of the density field, instead of the perturbation theory approach. As expected, linear evolution results in a symmetric Gaussian PDF.

Do we reproduce the geometric, scale-free enhancement of the power spectrum on linear scales? In the top panel of Fig. 5, we plot *dimensionless* 21-cm power spectra, $\Delta_{21}^2(k, z) = k^3 / (2\pi^2 V) |\delta_{21}(\mathbf{k}, z)|^2_k$ where $\delta_{21}(\mathbf{x}, z) \equiv \delta T_b(\mathbf{x}, z) / \delta T_b(z) - 1$. The spectra are generated by 21CMFAST in the limit of $T_s \gg T_\gamma$ and assuming $\bar{x}_{\text{HI}} = 1$. The solid red curves correspond to a 5-Gpc box with $\Delta x = 10 \text{ Mpc}$ cells, while the dot-dashed blue and dashed green curves correspond to a 1-Gpc box with different resolutions. The upper set of curves were computed including peculiar velocities, while the lower set were computed without peculiar velocities. The bottom three panels show the ratios of the power spectra that include redshift space distortions to those that do not.

Indeed the red curves in Fig. 5, which were evolved linearly, accurately capture the enhancement factor of 1.87, shown with a dotted horizontal line. The other two curves, which include first-order non-linear effects, show an enhancement of power *in excess of the purely geometric factor*. From equation (4), one sees that a high-value tail in the density distribution resulting from non-linear evolution would drive a corresponding negative tail in the dv_r/dr

⁸ There is an inconsistency in the above equations for 21CMFAST , as equation (4) is applied to the non-linear density field, whereas equation (2) assumes a linear δ . Nevertheless, as we shall show below, equation (4) reproduces the non-linear velocity gradient field from the numerical simulations remarkably well.




Figure 5. Top panel: Dimensionless power spectra from 21CMFAST, in a fully neutral universe, in the limit of $T_S \gg T_\gamma$. The upper set of curves were computed including peculiar velocities, while the lower set were computed without peculiar velocities. Bottom three panels: Ratios of the dimensionless power spectra with peculiar velocities to those not including peculiar velocities. The linear regime geometric enhancement of 1.87 is demarcated by the upper horizontal dotted line. In all panels, the dot-dashed blue and dashed green curves were generated from the same ICs in an $L = 1$ Gpc box, sampled with an initial resolution of $\Delta x = 0.56$ Mpc, with the evolved density, velocity and ionization fields generated at lower resolutions of $\Delta x = 3.3$ Mpc (dashed green curves) and 10 Mpc (dot-dashed blue curves). The solid red curves were generated from an $L = 5$ Gpc box with a single resolution of $\Delta x = 10$ Mpc. However, the density field used for the solid red curves was generated assuming linear evolution, while the others were generated with first-order perturbation theory (see Section 2.1).

distributions, which in turn enhances the 21-cm signal through the $1/(1+dv_r/dr/H)$ term in equation 1. Although the dv_r/dr distributions are zero-mean, the distributions of $1/(1+dv_r/dr/H)$ are not. The bias to higher values is further enhanced when weighted by the local density as in the δT_b expression, $\Delta/(1+dv_r/dr/H)$. Intuitively, infall in overdense regions causes photons emitted there to travel farther in order to reach a fixed relative redshift; therefore the optical depth and δT_b are increased in $\delta > 0$ regions (Barkana & Loeb 2005a).

To further explore this effect and compare our results with simulations, in Fig. 6, we plot the ratio of the *dimensional* 21-cm power spectra, $\delta \bar{T}_b(z)^2 \Delta_{21}^2(k, z)$, computed including peculiar velocities to those without peculiar velocities. The thin solid red curve in the upper panel corresponds to a fully neutral universe at $z = 9$. The hydrodynamic simulations go down to much smaller scales than plotted in Fig. 5, which are more non-linear and hence show a larger enhancement of power, though we confirm that most of this is due to the evolution in the mean signal, $\delta \bar{T}_b(z)^2$.

This enhanced 21-cm power from non-linear peculiar velocities obviously merits more investigation beyond the scope of this paper. Therefore we defer further analysis to future work. We caution however that it is unclear how well we can estimate this enhancement, due to the misuse of the $1/(1+dv_r/dr/H)$ term in equation (1). This expression diverges at $dv_r/dr = -H(z)$. To compensate for this be-




Figure 6. Ratios of the *dimensional* power spectra, $\delta \bar{T}_b(z)^2 \Delta_{21}^2(k, z)$, computed including peculiar velocities to those without peculiar velocities. Panels correspond to $(z, \bar{x}_{\text{HI}}) = (9.00, 0.86)$, $(7.73, 0.65)$, $(7.04, 0.38)$ and $(6.71, 0.20)$, (top to bottom). Solid red curves are generated from the hydrodynamic simulation, while the dashed blue curves are generated by 21CMFAST. The thin solid red curve in the upper panel corresponds to a fully neutral universe at $z = 9$. The linear regime geometric enhancement of 1.87 is demarcated by the upper horizontal dotted line.

haviour, we impose a maximum value of $|dv_r/dr| = 0.5H(z)$, and confirm that our results are only weakly sensitive to this choice in the $\sim 0.1H(z) - 0.7H(z)$ range. Similar misuses of the mapping from real space to redshift space have already been noted in the context of galaxy surveys, although the derivation of the $1/(1+dv_r/dr/H)$ term for the classic Kaiser effect is different (see Scoccimarro 2004, section IV B). Although the correlation of the field at large distances is small, the velocity gradient at a given point need not be small. Furthermore, equation (1) assumes that the brightness temperature is proportional to the 21-cm optical depth $\tau(v_{\text{obs}})$, which is only true when $\tau(v_{\text{obs}}) \ll 1$. Therefore, our treatment quickly becomes problematic as $dv_r/dr \rightarrow -H(z)$. If the user is interested in more accurate predictions of the 21-cm signal as observed with 21-cm interferometers, we recommend to turn-off the velocity gradient correction in 21CMFAST, and generate the final 21-cm intensity maps from the component fields. Already there are several steps required to transform a comoving simulation box to a simulated frequency-dependent signal for a particular experiment (e.g. Harker et al. 2010; Matejek et al., in preparation). The discussion below in the remainder of this section should not be sensitive to the above inconsistency in applying redshift-space corrections.

The remaining curves in Fig. 6 do not artificially set $\bar{x}_{\text{HI}} = 1$, but use the values of \bar{x}_{HI} from the numerical simulations. Solid red curves are generated from the hydrodynamic simulation, while the dashed blue curves are generated by 21CMFAST. We confirm the results of MF07: that the enhancement of power due to redshift space distortions vanishes in the early stages of reionization, and subsequently only affects small scales.

When the density field is linear, the dv_r/dr term is narrowly centred around zero. Therefore the enhancement factor from equation (1): $1/(1+dv_r/dr/H)$ can be Taylor expanded to first order,

becoming $1 - dv_r/dr/H$. This enhancement factor in the linear limit pre-reionization is symmetric around unity and, although the mean value of δT_b does not change, the associated ~ 1.867 enhancement in power (e.g. Bharadwaj & Ali 2004; Barkana & Loeb 2005a) is driven equally by over- and under-densities. When the density field is non-linear, however, the $(dv_r/dr)/H$ term has a corresponding high-value tail to negative values and is no longer symmetric (see Fig. 4). The $(dv_r/dr)/H < 0$ tail can rapidly enhance δT_b due to the divergence at $dv_r/dr = -H$. This is the cause of the large enhancement of power in the non-linear case when ignoring the ionization field. However, as reionization progresses, H II regions quickly ‘cover-up’ the most non-linear regions, due to the inside-out nature of reionization. Hence, the enhancement in power due to redshift space distortions fades as reionization progresses in Fig. 6.

Using linear theory, one can relate the velocity gradient field to the density field in k -space via equation 4. When converting to real space, this essentially means that the velocity gradient is proportional to the negative weighted density field with the k_r^2/k^2 window function enhancing the contributions of the LOS modes. Therefore, underdense regions will generally have a weaker 21-cm signal [negative $\delta \rightarrow$ positive $(dv_r/dr)/H \rightarrow$ decrease in δT_b]. Since reionization is inside out on large scales, the 21-cm signal comes increasingly from underdense regions as reionization progresses. The velocity gradients in those regions not covered by the ionized bubbles should become increasingly positive according to the above, and the enhancement to δT_b fades away and becomes negative. Hence the curves in the two bottom panels in Fig. 6 dip below unity at low k . This means that *during the final stages of reionization, peculiar velocity effects actually decrease power on moderate to large scales*. We confirm that both the mean signal and the large-scale power show this decrease due to peculiar velocities in the advanced stages of reionization, and it appears in both the simulations and the 21CMFAST. We defer a detailed analytic treatment of redshift space distortions to a future work.

In general, the relative enhancement due to peculiar velocity effects shown in Fig. 6 is well predicted by 21CMFAST on moderate to large scales, $k \lesssim 1 \text{ Mpc}^{-1}$, but is underpredicted at small scales and in a very neutral universe. This is attributable in part to the differences in the ionization field. Our ionization field algorithm, FFRT, does not properly capture ionization fronts and small-scale H II structure (Zahn et al. 2010). Thus statistics such as this one which are sensitive to small scales are not accurately reproduced. 21CMFAST’s underprediction of the power spectrum enhancement at small scales is likely also attributable in part to the fact that the Nyquist frequency corresponds to larger scales in the 21CMFAST boxes, since these are directly computed on a 256^3 grid, whereas the RT simulation is smoothed down from a 512^3 grid. This means that our shot noise on small scales (approaching the $L/256$ cell size) is larger than in the numerical simulation, and so fractional enhancements in power contribute less.

2.4 Full post-heating comparison of 21-cm emission

We now combine the terms from equation (1) to provide a full comparison of 21CMFAST and numerical simulations, *with the assumption of $T_S \gg T_\gamma$* . For the purposes of these comparisons, the two codes only share the same initial density field; the evolved density, velocity and ionization fields for 21CMFAST are all generated self-consistently as explained above.

In Fig. 7, we plot slices through the δT_b signal, generated from hydrodynamic simulation, the algorithm outlined in MF07,⁹ and 21CMFAST, left to right columns. Rows correspond to $(z, \bar{x}_{\text{H I}}) = (9.00, 0.86), (7.73, 0.65), (7.04, 0.38)$ and $(6.71, 0.20)$, top to bottom.

As already shown in Section 2.1, the density fields are modelled quite accurately with perturbation theory. One can also see that both seminumerical schemes reproduce the large-scale H II region morphology (shown in black) of the RT simulations. Differences emerge at moderate to small scales, with the FFRT ionization algorithm of 21CMFAST generally resulting in H II regions which too are connected. This difference is mostly attributable to the bubble-flagging algorithm; in general the ‘flagging-the-entire-sphere’ algorithm of MF07 better reproduces H II morphological structure than the ‘flagging-the-central-cell’ algorithm of Zahn et al. (2007) (e.g. MF07).

In Fig. 8, we show the PDFs of δT_b for the hydrodynamic simulation (red solid curves), 21CMFAST (blue dashed curves) and MF07 (magenta dotted curves). Panels correspond to $(z, \bar{x}_{\text{H I}}) = (9.00, 0.86), (7.73, 0.65), (7.04, 0.38)$ and $(6.71, 0.20)$, top to bottom. The left-hand panel was generated using the unfiltered δT_b field with cell length $\Delta x = 143/256 \text{ Mpc}$ (effectively $R \sim 0.35 \text{ Mpc}$), while the right-hand panel was generated from the δT_b field, filtered on $R_{\text{filter}} = 5 \text{ Mpc}$ scales.

From the left-hand panel, we see that we underpredict the number of ‘almost’ fully ionized cells, $\delta T_b \lesssim 10 \text{ mK}$. This can be traced to our algorithm for determining the partially ionized fraction in the remaining neutral cells. Our algorithm assumes that cells are partially ionized by subgrid sources chewing away at their host cell’s H I. Instead, partially ionized cells on these small scales generally correspond to unresolved ionization fronts from non-local sources (see appendix in Zahn et al. 2010). This discrepancy decreases as the cell size increases, since then the fraction of cells which are ionized by sources internal to the cell increases, and the assumption implicit in our FFRT algorithm becomes increasingly accurate. Aside from this, the distributions in the left-hand panel agree very well. This should not be surprising, since for comparison the ionization efficiency of the seminumerical schemes was chosen so that the mean ionized fraction at these epochs agrees with the numerical simulation (i.e. the spikes at $\delta T_b = 0 \text{ mK}$ match).¹⁰ The remainder of the signal at $\delta T_b \gtrsim 10 \text{ mK}$ merely reflects the density distribution of the neutral cells, and we have already demonstrated in Section 2.1 that our density fields match the hydrodynamic simulation quite well.¹¹

The right-hand panel of Fig. 8 shows the δT_b distributions, smoothed on $R_{\text{filter}} = 5 \text{ Mpc}$ scales. As evidenced by the smaller relative spike at $\delta T_b = 0 \text{ mK}$, the ionization fields on these scales are not as binary (i.e. either fully ionized or fully neutral) as those in the left-hand panel. Thus the PDFs encode more information on the

⁹ This algorithm is not exactly the same as in MF07, since partially ionized cells are still allowed. However, this does not affect our results, since on scales as small as our cell size, the ionization field produced by cosmological RT simulations can be treated as binary (i.e. either fully ionized or fully neutral; see the appendix of Zahn et al. 2010). The ionization field becomes less binary in the late stages of reionization, or when hard spectra dominate reionization.

¹⁰ Strictly speaking, the partially ionized cells do impact this calibration, but since most cells are either fully ionized or fully neutral, this is a small effect.

¹¹ The δT_b distributions include the additional check of sampling the density field of the remaining neutral cells, instead of the entire density field studied in Section 2.1. Therefore the close agreement of the $\delta T_b \gtrsim 10 \text{ mK}$ PDFs contains an additional confirmation of the accuracy of our ionization algorithms.




Figure 7. δT_b maps. The slices are generated from the hydrodynamic simulation, DEXM (MF07), and 21CMFAST, left to right columns. All slices are 143 Mpc on a side and 0.56-Mpc thick, and correspond to $(z, \bar{x}_{H1}) = (9.00, 0.86)$, $(7.73, 0.65)$, $(7.04, 0.38)$ and $(6.71, 0.20)$, top to bottom.

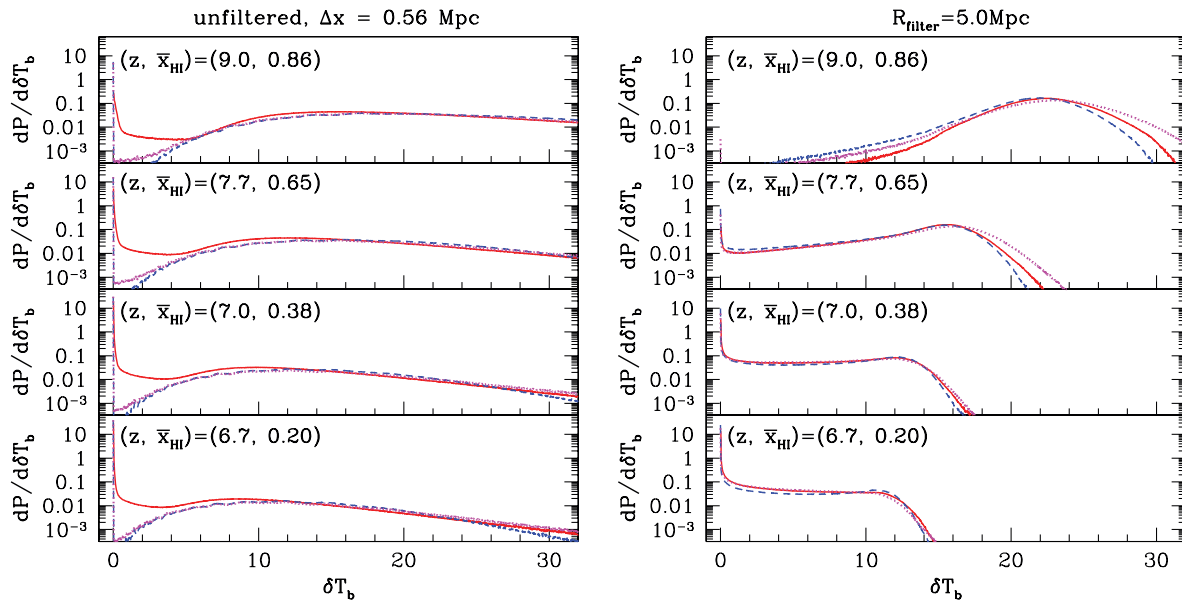


Figure 8. PDFs of δT_b created using equation (1) for the hydrodynamic simulation (red solid curves), 21cmFAST (blue dashed curves) and MF07 (magenta dotted curves). Panels correspond to $(z, \bar{x}_{\text{HI}}) = (9.0, 0.86)$, $(7.7, 0.65)$, $(7.0, 0.38)$ and $(6.7, 0.20)$, top to bottom. The left-hand panel was generated using the unfiltered δT_b field with cell length $\Delta x = 143/256$ Mpc (effectively $R \sim 0.35$ Mpc), while the right-hand panel was generated from the δT_b field, filtered on $R_{\text{filter}} = 5$ Mpc scales. All fields invoke the simplifying assumption of $T_S \gg T_\gamma$.

ionization algorithms. The top panel at $\bar{x}_{\text{HI}} = 0.86$, shows that the predicted distributions of δT_b agree well around the mean signal, but the seminumerical distributions diverge from the RT in the wings. As mentioned before, the ‘flagging-the-entire-sphere’ ionization algorithm from MF07 results in less connected H II regions, and so there are more isolated 5.0-Mpc neutral patches. This results in an increased number of high- δT_b regions. The converse is true for the FFRT ionization scheme which is the fiducial setting of 21cmFAST. The agreement between the schemes improves as reionization progresses.

In Fig. 9, we compare the power spectra of these δT_b boxes. Again, the hydrodynamic simulation is shown with red solid curves, MF07 is shown with dotted magenta curves, and 21cmFAST is shown with dashed blue curves.

At all scales, the power spectra agree with each other at the 10s of per cent level.¹² At moderate to large scales, agreement is best, with MF07 performing slightly better than the FFRT algorithm which is default in 21cmFAST.¹³ On smaller scales, MF07 predicts too much power, while 21cmFAST underpredicts the power. It was shown in Zahn et al. (2010) that the FFRT ionization algorithm used in 21cmFAST overpredicts the correlation of the ionization and density fields on small scales, due to the fact that it operates directly on the evolved density field. This strong cross-correlation results in

¹² The seminumerical simulations show an increase in power approaching the Nyquist frequency, which is most likely just shot noise of our 256^3 boxes. The numerical simulations do not show the same upturn, since they were generated on higher-resolution grids (512^3 for the RT and 768^3 for the density), and subsequently smoothed down to 256^3 ; numerical simulations generated directly on the same scale 256^3 grids show similar shot noise upturns in power on these scales (see fig. 7 in Zahn et al. 2010).

¹³ Note that the FFRT results shown here are not precisely analogous to those in Zahn et al. (2010), since there the evolved density field was taken from an N -body simulation, whereas in 21cmFAST, we self-consistently generate the density field according to Section 2.1.




Figure 9. Comparison of 21-cm power spectra obtained from the hydrodynamic cosmological simulation (solid red curves), and the seminumerical algorithms in MF07 (dotted magenta curves) and 21cmFAST (dashed blue curves). Sets of curves correspond to $(z, \bar{x}_{\text{HI}}) = (9.0, 0.86)$, $(7.7, 0.65)$, $(7.0, 0.38)$ and $(6.7, 0.20)$, top to bottom at high k .

an underprediction of 21-cm power on these scales. The converse is true of the MF07 scheme, which although using discrete source haloes, paints entire filtered regions as ionized, thus underpredicting the cross-correlation of the ionization and density fields.

Most importantly, the differences in the power spectra between the models is much smaller than the differences due to the progress of reionization. Therefore, one might naively predict that the seminumerical schemes are accurate enough to estimate \bar{x}_{HI} from

the power spectra to $\pm \lesssim 0.1$, or even better if the behaviour of the models is understood. However, there are many astrophysical uncertainties associated with prescriptions for sources and sinks of ionizing photons during the epoch of reionization, and it will likely be these which regulate the achievable constraints on \bar{x}_{HI} . *Therefore it is imperative for models to be fast and be able to span large regions of parameter space.* A single 21CMFAST realization of the δT_{b} fields shown in this section (generated from 1536³ ICs) takes ~ 30 min to compute on a single-processor computer.

3 THE SPIN TEMPERATURE

We now relax the requirement in Section 2 of $T_{\text{S}} \gg T_{\gamma}$, and derive the full 21-cm brightness temperature offset from equation (1), including the spin temperature field. As mentioned previously, models predict that the heating epoch concluded well before the bulk of reionization, at $z \gg 10$ (Furlanetto 2006; Chen & Miralda-Escudé 2008; Santos et al. 2008; Baek et al. 2009). However, the second-generation 21-cm interferometers, such as SKA, might be able to peek into this high-redshift regime of the dark ages. Furthermore, the astrophysical quantities at high- z are uncertain, and we do not really know how robust is the assumption of $T_{\text{S}} \gg T_{\gamma}$ even during the early stages of reionization. Therefore, for many applications, especially parameter studies, it is important to compute the spin temperature field. Unfortunately, there is currently no numerical simulation that includes the computationally expensive radiative transfer of both X-rays and Ly α photons from atomically or molecularly cooled sources (though see the recent work of Baek et al. 2010, who perform RT simulations on a small subset of sources, with $M \gtrsim 10^{10} M_{\odot}$). Therefore we cannot directly compare our spin temperature fields to numerical simulations.

Our derivations in this section are similar to other semi-analytic models (Furlanetto 2006; Pritchard & Furlanetto 2007; Santos et al. 2008). However, unlike Santos et al. (2008, 2009), we do not explicitly resolve the halo field as an intermediary step. Instead we operate directly on the evolved density fields, using excursion set formalism to estimate the mean number of sources inside spherical shells corresponding to some higher redshift. As discussed above, bypassing the halo field allows the code to be faster, with modest memory requirements. Below we go through our formalism in detail.

The spin temperature can be written as (e.g. Furlanetto et al. 2006a)

$$T_{\text{S}}^{-1} = \frac{T_{\gamma}^{-1} + x_{\alpha} T_{\alpha}^{-1} + x_{\text{c}} T_{\text{K}}^{-1}}{1 + x_{\alpha} + x_{\text{c}}}, \quad (5)$$

where T_{K} is the kinetic temperature of the gas and T_{α} is the colour temperature, which is closely coupled to the kinetic gas temperature, $T_{\alpha} \approx T_{\text{K}}$ (Field 1959). There are two coupling coefficients in the above equation. The collisional coupling coefficient can be written as

$$x_{\text{c}} = \frac{0.0628 \text{ K}}{A_{10} T_{\gamma}} \left[n_{\text{HI}} \kappa_{1-0}^{\text{HH}}(T_{\text{K}}) + n_{\text{e}} \kappa_{1-0}^{\text{eH}}(T_{\text{K}}) + n_{\text{p}} \kappa_{1-0}^{\text{pH}}(T_{\text{K}}) \right], \quad (6)$$

where $A_{10} = 2.85 \times 10^{-15} \text{ s}^{-1}$ is the spontaneous emission coefficient; n_{HI} , n_{e} and n_{p} are the number density of neutral hydrogen, free electrons and protons, respectively; and $\kappa_{1-0}^{\text{HH}}(T_{\text{K}})$, $\kappa_{1-0}^{\text{eH}}(T_{\text{K}})$, and $\kappa_{1-0}^{\text{pH}}(T_{\text{K}})$ are taken from Zygelman (2005), Furlanetto & Furlanetto (2007) and Furlanetto & Furlanetto (2007), respectively. The Wouthuysen–Field (Wouthuysen 1952; Field 1958; WF) cou-

pling coefficient can be written as

$$x_{\alpha} = 1.7 \times 10^{11} (1+z)^{-1} S_{\alpha} J_{\alpha}, \quad (7)$$

where S_{α} is a correction factor of order unity involving detailed atomic physics, and J_{α} is the Lyman α background flux in units of $\text{pcm}^{-2} \text{ s}^{-1} \text{ Hz}^{-1} \text{ sr}^{-1}$. We compute T_{α} and S_{α} according to Hirata (2006) and denote proper units with the prefix ‘p’.

According to the above equations, there are two main fields governing the spin temperature: (1) the kinetic temperature of the gas, $T_{\text{K}}(\mathbf{x}, z)$, and (2) the Ly α background, $J_{\alpha}(\mathbf{x}, z)$. We address these in Sections 3.1 and 3.2, respectively.

3.1 The kinetic temperature

3.1.1 Evolution equations

To calculate the kinetic temperature, one must keep track of the inhomogeneous heating history of the gas. We begin by writing down the evolution equation for $T_{\text{K}}(\mathbf{x}, z)$ and the local ionized fraction in the ‘neutral’ (i.e. outside of the ionized regions discussed in Section 2.2) IGM, $x_{\text{e}}(\mathbf{x}, z)$:

$$\frac{dx_{\text{e}}(\mathbf{x}, z')}{dz'} = \frac{dt}{dz'} \left[\Lambda_{\text{ion}} - \alpha_{\text{A}} C x_{\text{e}}^2 n_{\text{b}} f_{\text{H}} \right], \quad (8)$$

$$\begin{aligned} \frac{dT_{\text{K}}(\mathbf{x}, z')}{dz'} &= \frac{2}{3k_{\text{B}}(1+x_{\text{e}})} \frac{dt}{dz'} \sum_p \epsilon_p \\ &+ \frac{2T_{\text{K}}}{3n_{\text{b}}} \frac{dn_{\text{b}}}{dz'} - \frac{T_{\text{K}}}{1+x_{\text{e}}} \frac{dx_{\text{e}}}{dz'}, \end{aligned} \quad (9)$$

where $n_{\text{b}} = \bar{n}_{\text{b},0}(1+z')^3[1 + \delta_{\text{nl}}(\mathbf{x}, z')]$ is the total (H + He) baryonic number density at (\mathbf{x}, z') , $\epsilon_p(\mathbf{x}, z')$ is the heating rate *per baryon*¹⁴ for process p in erg s^{-1} , Λ_{ion} is the ionization rate per baryon, α_{A} is the case-A recombination coefficient, $C \equiv \langle n^2 \rangle / \langle n \rangle^2$ is the clumping factor on the scale of the simulation cell, k_{B} is the Boltzmann constant, f_{H} is the hydrogen number fraction¹⁵ and we distinguish between the output redshift of interest, z , and some arbitrary higher redshift, z' .¹⁶ We also make the accurate assumption that single ionized helium and hydrogen are ionized to the same degree, $x_{\text{e}}(\mathbf{x}, z')$, inside the mainly neutral IGM (e.g. Wyithe & Loeb 2003).

In order to speed up our calculation, we extrapolate the cell’s density to higher redshifts assuming linear evolution from z (the desired output redshift at which we compute the non-linear density field with perturbation theory): $\delta_{\text{nl}}(\mathbf{x}, z') \approx \delta_{\text{nl}}(\mathbf{x}, z) D(z')/D(z)$, where $D(z)$ is the linear growth factor. In principle, we should follow the non-linear redshift evolution of each cell’s density, $\delta_{\text{nl}}(\mathbf{x}, z')$. However, linearly extrapolating backwards from z is a good approximation, considering that the majority of cells sized for cosmological simulations are in the linear or quasi-linear regime at very high redshifts where heating is important (see, e.g., Fig. 2). Additionally, the evolution of structure, and thus also of heating, is extremely rapid

¹⁴ Note that our notation is different than that in Furlanetto (2006) and Pritchard & Furlanetto (2007), who present heating and ionization rates per proper volume.

¹⁵ Equation (8) ignores helium recombinations, which is a good approximation given that most He recombining photons will cause ionizations of H I or He I.

¹⁶ For clarity of presentation, we will only explicitly show dependent variables of functions on the left-hand side of equations. Where it is obvious, we also do not explicitly show dependences on (\mathbf{x}, z') .

during this epoch, and is dominated by redshifts not much higher than z . Therefore, we can rewrite equations (8) and (9) as

$$\frac{dx_e(\mathbf{x}, z')}{dz'} = \frac{dt}{dz'} \Lambda_{\text{ion}} - \frac{dt}{dz'} \alpha_A C x_e^2 f_{\text{H}} \bar{n}_{b,0} (1+z')^3 \left[1 + \delta_{\text{nl}}(\mathbf{x}, z) \frac{D(z')}{D(z)} \right], \quad (10)$$

$$\frac{dT_{\text{K}}(\mathbf{x}, z')}{dz'} = \frac{2}{3k_B(1+x_e)} \frac{dt}{dz'} \sum_p \epsilon_p + \frac{2T_{\text{K}}}{1+z'} + \frac{2T_{\text{K}}}{3} \frac{dD(z')/dz'}{D(z)/\delta_{\text{nl}}(\mathbf{x}, z) + D(z')} - \frac{T_{\text{K}}}{1+x_e} \frac{dx_e}{dz'}. \quad (11)$$

The first term in equation (11) corresponds to the heat input, which for our purposes is dominated by the heating processes discussed below. The second term corresponds to the Hubble expansion, the third corresponds to adiabatic heating and cooling from structure formation, and the fourth corresponds to the change in the total number of gas particles due to ionizations.

3.1.2 Heating and ionization rates

At very high redshifts, Compton scattering between the CMB photons and the residual free electrons sets $T_{\text{K}} = T_{\gamma}$. After decoupling, the gas temperature evolves through adiabatic cooling, $T_{\text{K}}(z') = 2.73 \times (1+z_{\text{dec}})[(1+z')/(1+z_{\text{dec}})]^2$, where the decoupling redshift is given by $z_{\text{dec}} \approx 137(\Omega_b h^2/0.022)^{0.4} - 1$ (Peebles 1993). However, as this is only approximate, we shall use the actual Compton heating rate (e.g. Seager, Sasselov & Scott 2000) in the sum in equation (11):

$$\frac{2}{3k_B(1+x_e)} \epsilon_{\text{comp}} = \frac{x_e}{1+f_{\text{He}}+x_e} \frac{8\sigma_{\text{T}} u_{\gamma}}{3m_e c} (T_{\gamma} - T_{\text{K}}), \quad (12)$$

where f_{He} is the helium number fraction, u_{γ} is the energy density of the CMB and σ_{T} is the Thomson cross-section. The initial conditions for x_e and T_{K} in 21CMFAST can either be provided by the user, or taken from the publicly available code RECFAST¹⁷ (Seager, Sasselov & Scott 1999).

We now proceed to outline our procedure for estimating X-ray heating, which is the dominant heating process in this epoch.¹⁸ We compute the heating rate per particle by summing contributions from sources located in concentric spherical shells around (\mathbf{x}, z') . First, we take the standard ansatz in assuming that sources emit photons with a rate proportional to the growth of the total mass fraction inside DM haloes, f_{coll} . Note that while this may not be strictly true, we are averaging over large volumes and many sources, and so it is probably a decent assumption in practice. Thus the total X-ray emission rate (in s^{-1}) per redshift interval from luminous sources located between z'' and $z'' + dz''$ (where $z'' \geq z'$) can be

written as

$$\frac{d\dot{N}_{\text{X}}}{dz''} = \zeta_{\text{X}} f_{*} \Omega_b \rho_{\text{crit},0} \left(1 + \delta_{\text{nl}}^{R''} \right) \frac{dV}{dz''} \frac{df_{\text{coll}}}{dt}, \quad (13)$$

where the efficiency, ζ_{X} , is the number of photons per solar mass in stars and the remaining terms on the RHS correspond to the total star formation rate inside the spherical shell demarcated by z'' and $z'' + dz''$. Specifically, f_{*} is the fraction of baryons converted to stars, and $dV(z'')$ is the comoving volume element at z'' (i.e. volume of the shell). The collapsed fraction is computed according to a hybrid prescription, similar to Barkana & Loeb (2004, 2008):

$$f_{\text{coll}}(\mathbf{x}, z'', R'', S_{\text{min}}) = \frac{\bar{f}_{\text{ST}}}{\bar{f}_{\text{PS,nl}}} \text{erfc} \left[\frac{\delta_c - \delta_{\text{nl}}^{R''}}{\sqrt{2[S_{\text{min}} - S^{R''}]}} \right], \quad (14)$$

where R'' is the comoving, null-geodesic distance between z' and z'' ; S_{min} and $S^{R''}$ are the mass variances on the scales corresponding to the smallest-mass sources and R'' , respectively; $\delta_{\text{nl}}^{R''}$ is the evolved density¹⁹ smoothed on scale R'' , which we again linearly extrapolate from z : $\bar{\delta}_{\text{nl}}^{R''}(\mathbf{x}, z'') = \bar{\delta}_{\text{nl}}^{R''}(\mathbf{x}, z) D(z'')/D(z)$; δ_c is the critical linear density corresponding to virialization; and $\bar{f}_{\text{ST}}(z'', S_{\text{min}})$ is the mean ST collapsed fraction (with the Jenkins et al. 2001 normalization) while $\bar{f}_{\text{PS,nl}}(z'', S_{\text{min}}, R'')$ is the mean PS collapse fraction, averaged over the evolved density field, $\delta_{\text{nl}}^{R''}$. Therefore the normalization factor, $\bar{f}_{\text{ST}}/\bar{f}_{\text{PS,nl}}$, assures that the mean collapse fraction matches the ST collapse fraction (in agreement with N -body simulations; e.g. Jenkins et al. 2001; McQuinn et al. 2007; Trac & Cen 2007), while the fluctuations are sourced by the conditional PS model applied on the evolved density field.²⁰ There is an implicit assumption in equation (13) that $dD(z'')/dz'' \ll df_{\text{coll}}/dz''$, which is quite accurate for all pertinent epochs, as the large-scale density evolves much slower than the exponential growth of the collapsed fraction.

We assume that the X-ray luminosity of sources can be characterized with a power law of the form, $L_e \propto (\nu/\nu_0)^{-\alpha}$, with ν_0 being the lowest X-ray frequency escaping into the IGM. We can then write the arrival rate [i.e. number of photons $\text{s}^{-1} \text{Hz}^{-1}$ seen at $(\mathbf{x}, z')]$ of X-ray photons with frequency ν , from sources within z'' and $z'' + dz''$ as

$$\frac{d\phi_{\text{X}}(\mathbf{x}, \nu, z', z'')}{dz''} = \frac{d\dot{N}_{\text{X}}}{dz''} \alpha \nu_0^{-1} \left(\frac{\nu}{\nu_0} \right)^{-\alpha-1} \left(\frac{1+z''}{1+z'} \right)^{-\alpha-1} e^{-\tau_{\text{X}}}, \quad (15)$$

where the last term accounts for IGM attenuation. For computation efficiency, we compute the IGM X-ray mean free path through the

¹⁹ Although the collapse criterion for conditional PS was derived using the linear density field, Zahn et al. (2010) find that using the evolved density in equation (14) when computing ionization fields yields a better match to radiative transfer simulations. However, this distinction is less important in the context of computing X-ray flux fields than ionizing flux fields, since the density is more linear on the most pertinent scales R'' , due to the larger mean free path of X-ray photons.

²⁰ Also note that equation (14) implicitly assumes that the sources are evenly distributed within R'' , which is the same assumption inherent in the ionization algorithm. Although the rate of change of the collapse fraction in each shell can be computed from PS, we hesitate to apply this prescription to X-rays, since the ionization algorithm which uses equation (14) has already been rigorously tested and found to agree well with RT simulations of reionization (Zahn et al. 2010).

¹⁷ <http://www.astro.ubc.ca/people/scott/recfast.html>

¹⁸ There are two other processes that may be important. The first is shock heating in the IGM: with such a small temperature, a population of weak shocks could substantially change the thermal history. This appears to be the case in at least one simulation (Gnedin & Shaver 2004), although other simulations (Kuhlen, Madau & Montgomery 2006) and analytic models (Furlanetto & Loeb 2004) predict much smaller effects. The second is any exotic process, such as dark matter decay or annihilation, that produces X-rays or hot electrons (Furlanetto, Oh & Pierpaoli 2006b).

mean²¹ IGM:

$$\tau_X(\nu, z', z'') = \int_{z''}^{z'} d\hat{z} \frac{c dt}{d\hat{z}} \bar{x}_{\text{H I}}(\hat{z}) \bar{n}(\hat{z}) \bar{\sigma}(z', \hat{\nu}), \quad (16)$$

where the photoionization cross-section is weighted over species, $\bar{\sigma}(z', \hat{\nu}) \equiv f_{\text{H}}(1 - \bar{x}_e)\sigma_{\text{H}} + f_{\text{He I}}(1 - \bar{x}_e)\sigma_{\text{He I}} + f_{\text{He II}}\bar{x}_e\sigma_{\text{He II}}$ and is evaluated at $\hat{\nu} = \nu(1 + \hat{z})/(1 + z')$. In practice, the contribution of He II can be ignored. We also remind the reader that $\bar{x}_{\text{H I}}$ is the volume-filling factor of neutral regions during reionization.

Finally, one obtains the X-ray heating rate per baryon, ϵ_X from equation (11) by integrating over frequency and z'' :

$$\epsilon_X(\mathbf{x}, z') = \int_{\nu_0}^{\infty} d\nu \sum_i (h\nu - E_i^{\text{th}}) f_{\text{heat}} f_i x_i \sigma_i \int_{z'}^{\infty} dz'' \frac{d\phi_X/dz''}{4\pi r_p^2}, \quad (17)$$

where r_p is the proper, null-geodesic separation of z' and z'' , and the frequency integral includes a sum over species, $i = \text{H I}, \text{He I}$ or He II , in which f_i is the number fraction; x_i is the cell's species ionization fraction [which for H I and He I is $(1 - x_e)$, and for He II is x_e]; σ_i the ionization cross-section; and E_i^{th} is the ionization threshold energy of species i . The factor $f_{\text{heat}}[h\nu - E_i^{\text{th}}, x_e(\mathbf{x}, z')]$ is defined as the fraction of the *electron energy*, $h\nu - E_i^{\text{th}}$, deposited as heat. We use the new results of Furlanetto & Stoever (2009) (who provide a more complete treatment than the previous work of Shull & van Steenberg 1985), to compute f_{heat} , as well as f_{ion} and $f_{\text{Ly}\alpha}$ below. These fractions take into account the cell's local ionization state, x_e , as opposed to the global IGM value, \bar{x}_e , used to calculate the optical depth in equation (16).

Unfortunately, the double integral equation (17) is slow to evaluate due to the attenuation term, which depends on both redshift and frequency (equation 16). To speed up computation, we make the additional approximation that all photons with optical depth $\tau_X \leq 1$ are absorbed, while no photons with optical depth $\tau_X > 1$ are absorbed. Such a step-function attenuation has been shown to yield fairly accurate ionizing photon flux probability distributions (see fig. 2 in Mesinger & Furlanetto 2009, although comparisons were limited to a single set of parameters). This approximation allows us to separate the frequency and redshift integrals in equation (17), removing the exponential attenuation term from $d\phi_X/dz''$ and setting the lower bound of the frequency integral to either ν_0 or the frequency corresponding to an optical depth of unity, $\nu_{\tau=1}(\bar{x}_e, z', z'')$, whichever is larger. Expanding and grouping the terms, we obtain

$$\epsilon_X(\mathbf{x}, z') = \zeta_X \alpha c \nu_0^{-1} f_* \Omega_b \rho_{\text{crit},0} (1 + z')^{\alpha+1} \int_{\text{Max}[\nu_0, \nu_{\tau=1}]}^{\infty} d\nu \left(\frac{\nu}{\nu_0}\right)^{-\alpha-1} \sum_i (h\nu - E_i^{\text{th}}) f_{\text{heat}} f_i x_i \sigma_i \int_{z'}^{\infty} dz'' (1 + z'')^{-\alpha+2} \left(1 + \bar{\delta}_{\text{nl}}^{R''}\right) \frac{df_{\text{coll}}}{dz''}. \quad (18)$$

²¹ Note that our framework for computing the spin temperature starts to break down in the advanced stages of reionization. When H II regions become sizable, there will be large sightline-to-sightline fluctuations in the X-ray optical depth, which can only be taken into account with approximate radiative transfer. In our fiducial model (see below), heating has already saturated at $\bar{x}_{\text{H I}} \gtrsim 0.99$, so this is not a concern. However, some extreme models might be able to push the heating regime well into the bulk of reionization; we caution the user against overinterpreting the results from 21CMFAST in that regime.

Now the integrand in both integrals only depends on a single variable, and the entire frequency integral can be treated as a function of z'' .

Analogously, we can also express the ionization rate per particle in equation (10) as

$$\Lambda_{\text{ion}}(\mathbf{x}, z') = \int_{\text{Max}[\nu_0, \nu_{\tau=1}]}^{\infty} d\nu \sum_i f_i x_i \sigma_i F_i \int_{z'}^{\infty} dz'' \frac{d\phi_X/dz''}{4\pi r_p^2}, \quad (19)$$

$$F_i = (h\nu - E_i^{\text{th}}) \left(\frac{f_{\text{ion,H I}}}{E_{\text{H I}}^{\text{th}}} + \frac{f_{\text{ion,He I}}}{E_{\text{He I}}^{\text{th}}} + \frac{f_{\text{ion,He II}}}{E_{\text{He II}}^{\text{th}}} \right) + 1,$$

where $f_{\text{ion},j}[h\nu - E_i^{\text{th}}, x_e(\mathbf{x}, z'), j]$ is now the fraction of the electron's energy going into *secondary* ionizations of species j , with the unity term inside the sum accounting for the primary ionization of species i .

3.2 The Ly α background

The Lyman α background has two main contributors: X-ray excitation of H I, $J_{\alpha,X}$; and direct stellar emission of photons between Ly α and the Lyman limit, $J_{\alpha,*}$. The former can easily be related to the X-ray heating rate, assuming that the X-ray energy injection rate is balanced by photons redshifting out of Ly α resonance (Pritchard & Furlanetto 2007):

$$J_{\alpha,X}(\mathbf{x}, z) = \frac{cn_b}{4\pi H(z)\nu_\alpha} \int_{z'}^{\infty} dz'' \frac{d\phi_X/dz''}{4\pi r_p^2} \int_{\text{Max}[\nu_0, \nu_{\tau=1}]}^{\infty} d\nu \sum_i (h\nu - E_i^{\text{th}}) \frac{f_{\text{Ly}\alpha}}{h\nu_\alpha} f_i x_i \sigma_i, \quad (20)$$

where $f_{\text{Ly}\alpha}[h\nu - E_i^{\text{th}}, x_e(\mathbf{x}, z)]$ is the fraction of the electron's energy going into Ly α photons.

Because of the high resonant optical depth of neutral hydrogen, photons redshifting into any Lyman- n resonance at (\mathbf{x}, z) will be absorbed in the IGM. They then quickly and locally cascade with a fraction $f_{\text{recycle}}(n)$ passing through Ly α and inducing strong coupling (Hirata 2006; Pritchard & Furlanetto 2006). Therefore, the direct stellar emission component of the Ly α background (in $\text{pc m}^{-2} \text{s}^{-1} \text{Hz}^{-1} \text{sr}^{-1}$) can be estimated with a sum over the Lyman resonance backgrounds (e.g. Barkana & Loeb 2005b):

$$J_{\alpha,*}(\mathbf{x}, z) = \sum_{n=2}^{n_{\text{max}}} J_\alpha(n, \mathbf{x}, z) = \sum_{n=2}^{n_{\text{max}}} f_{\text{recycle}}(n) \int_z^{z_{\text{max}}(n)} dz' \frac{1}{4\pi} \frac{d\phi_*^e(\nu'_n, \mathbf{x})/dz'}{4\pi r_p^2}, \quad (21)$$

where the emissivity per unit redshift (number of photons $\text{s}^{-1} \text{Hz}^{-1}$) is calculated analogously to the X-ray luminosity above:

$$\frac{d\phi_*^e(\nu'_n, \mathbf{x})}{dz'} = \varepsilon(\nu'_n) f_* \bar{n}_{b,0} \left(1 + \bar{\delta}_{\text{nl}}^{R''}\right) \frac{dV}{dz'} \frac{df_{\text{coll}}}{dt}. \quad (22)$$

Here $\varepsilon(\nu)$ is the number of photons produced per Hz per stellar baryon, and is evaluated at the emitted (rest frame) frequency:

$$\nu'_n = \nu_n \frac{1 + z'}{1 + z}. \quad (23)$$

The upper limit of the redshift integral in equation (21) corresponds to the redshift of the next Lyman resonance:

$$1 + z_{\text{max}}(n) = (1 + z) \frac{1 - (n + 1)^{-2}}{1 - n^{-2}}. \quad (24)$$

Following Barkana & Loeb (2005b), we truncate the sum at $n_{\max} = 23$, and use their Population II and Population III spectral models for $\varepsilon(\nu)$. For computational efficiency, one can rearrange the terms in equation (21), placing the sum over Lyman transitions inside the redshift integral. Substituting in equation (22) and simplifying, we obtain

$$J_{\alpha,*}(\mathbf{x}, z) = \frac{f_* \bar{n}_{b,0} c}{4\pi} \int_z^\infty dz' (1+z')^3 \left(1 + \delta_{\text{nl}}^{R''}\right) \frac{df_{\text{coll}}}{dz'} \sum_{n=2}^{n(z')} f_{\text{recycle}}(n) \varepsilon(\nu'_n), \quad (25)$$

where the contribution from the sum over the Lyman transitions is a function of z' , and is zero at $z' > z_{\max}(n=2)$.

The total Lyman α background is then just the sum of the above components:

$$J_{\alpha,\text{tot}}(\mathbf{x}, z) = J_{\alpha,X}(\mathbf{x}, z) + J_{\alpha,*}(\mathbf{x}, z). \quad (26)$$

In our fiducial model, we do not explicitly take into account other soft-UV sources of Ly α such as quasars, assuming that these are subdominant to the stellar emission. However, our framework makes it simple to add additional source terms to the integrand of equation (21), if the user wishes to explore such scenarios (e.g. Volonteri & Gnedin 2009).

3.3 Results: complete δT_b evolution

All of the results in this section are from an $L = 1$ Gpc simulation, whose ICs are sampled on a 1800^3 grid, with the final low-resolution boxes being 300^3 (3.33 Mpc cells). Our fiducial model below assumes $f_* = 0.1$, $\zeta_X = 10^{57} M_\odot^{-1}$ (~ 1 X-ray photon per stellar baryon)²², $h\nu_0 = 200$ eV, $\alpha = 1.5$, $T_{\text{vir,min}} = 10^4$ K for all sources (X-ray, Lyman α and ionizing), $C = 2$, $R_{\max} = 30$ Mpc, $\zeta_{\text{ion}} = 31.5^{23}$ and the stellar emissivity, ε , of Pop II stars from Barkana & Loeb (2005b) normalized to 4400 ionizing photons per stellar baryon. The free parameters pertaining to the spin temperature evolution were chosen to match those in Furlanetto (2006) and Pritchard & Furlanetto (2007), to facilitate comparison. It is trivial to customize the code to add, for example, redshift or halo mass dependences to these free parameters. The impressive length of the above list of uncertain astrophysical parameters (which itself is only a simplified description of the involved processes) serves well to underscore the need for a fast, portable code, capable of quickly scrolling through parameter space.

We also note that the T_S calculations outlined in Section 3 are the slowest part of the 21CMFAST code (as they involve tracking evolution down to the desired redshift), and therefore should only be used in the regime where they are important ($z \gtrsim 17$ in our fiducial model). For example, generating a δT_b box, assuming $T_S \gg T_\gamma$, on a 300^3 grid takes only a few minutes on single processor (depending on the choice of higher resolution for sampling the ICs). However, including the spin temperature field takes an additional day of computing time. Nevertheless, once the spin temperature evolution is computed for a given realization at z , all of the intermediate outputs at $z' > z$ can be used to compute δT_b at those redshifts at no additional computation cost.

²² This number was chosen to match the total X-ray luminosity per unit star formation rate at low redshifts (see Furlanetto 2006 and references therein for details).

²³ This emissivity was chosen so that the mid-point of reionization is $z \sim 10$ and the end is $z \sim 7$.




Figure 10. Evolution of the mean temperatures from 21CMFAST in our fiducial model. Solid, dashed and dotted curves show T_S , T_K and T_γ , respectively.

Before showing detailed results, it would be useful to summarize the various evolutionary stages (cf. section 3.1 in Pritchard & Furlanetto 2007). The reader is encouraged to refer to the evolution of the mean temperatures shown in Fig. 10 and/or view the full movie available at <http://www.astro.princeton.edu/~mesinger/Movies/delT.mov> while reading the following.

(i) **Collisional coupling; $\bar{T}_K = \bar{T}_S \leq T_\gamma$:** At high redshifts, the IGM is dense, so the spin temperature is collisionally coupled to the gas kinetic temperature. The gas temperature is originally coupled to the CMB, but after decoupling cools adiabatically as $\propto (1+z)^{-2}$, faster than the CMB. The 21-cm brightness temperature offset from the CMB in this regime starts at zero, when all three temperatures are equal, and then becomes increasingly negative as T_S and T_K diverge more and more from T_γ . The fluctuations in δT_b are driven by the density field, as collisional coupling is efficient everywhere. In our fiducial model, this epoch corresponds to $100 \lesssim z$.

(ii) **Collisional decoupling; $\bar{T}_K < \bar{T}_S < T_\gamma$:** The IGM becomes less dense as the Universe expands. The spin temperature starts to decouple from the kinetic temperature, and begins to approach the CMB temperature again, thus δT_b starts rising towards zero. Decoupling from T_K occurs as a function of the local gas density, with underdense regions decoupling first. The power spectrum initially steepens, as small-scale density fluctuations drive the additional fluctuations of the collisional coupling coefficient. As the spin temperature in even the overdense regions finally decouples from the kinetic temperature, the power spectrum flattens again, and the mean signal drops. In our fiducial model, this epoch corresponds to $35 \lesssim z \lesssim 100$.

(iii) **Collisional decoupling \rightarrow WF coupling transition; $\bar{T}_K < \bar{T}_S \approx T_\gamma$:** As the spin temperature throughout the IGM decouples from the kinetic temperature, the mean signal is faint and might disappear, if the first sources wait long enough to ignite. In our fiducial model, this transition regime does not really exist. In fact our first sources turn on before the spin temperature fully decouples from the kinetic temperature.

(iv) **WF coupling; $\bar{T}_K < \bar{T}_S < T_\gamma$:** The first astrophysical sources turn on, and begin coupling the spin temperature of the nearby IGM to the kinetic temperature through the WF effect (Ly α coupling).

As the requirements for Ly α coupling are more modest than those to heat the gas through X-ray heating, the kinetic temperature keeps decreasing in this epoch. The mean brightness temperature offset from the CMB starts becoming more negative²⁴ again and can even reach values of $\delta T_b < -100$ mK. In our fiducial model, this epoch corresponds to $25 \lesssim z \lesssim 35$.

(v) **WF coupling \rightarrow X-ray heating transition;** $\bar{T}_k \sim \bar{T}_s < T_\gamma$: Ly α coupling begins to saturate as most of the IGM has a spin temperature which is strongly coupled to the kinetic temperature. The mean spin temperature reaches a minimum value, and then begins increasing. A few underdense voids are left only weakly coupled as X-rays from the first sources begin heating the surrounding gas in earnest, raising its kinetic temperature. The 21-cm power spectrum steepens dramatically as small-scale overdensities now host hot gas, while on large scales the gas is uniformly cold as Ly α coupling saturates. As inhomogeneous X-ray heating continues, the large-scale power comes back up. In our fiducial model, this transition occurs around $z \sim 25$.

(vi) **X-ray heating;** $\bar{T}_k = \bar{T}_s < T_\gamma$: X-rays start permeating the IGM. The fluctuations in δT_b are now at their maximum, as regions close to X-ray sources are heated above the CMB temperature, $\delta T_b > 0$, while regions far away from sources are still very cold, $\delta T_b < 0$. A ‘shoulder’ in the power spectrum, similar to that seen in the epoch of reionization (e.g. McQuinn et al. 2007), moves from small scales to large scales. X-rays eventually heat the entire IGM, and 21-cm can only be seen in emission. The power spectrum falls as this process nears completion. In our fiducial mode, this epoch corresponds to $18 \lesssim z \lesssim 25$.

(vii) **X-ray heating \rightarrow reionization transition;** $\bar{T}_k = \bar{T}_s > T_\gamma$: X-rays have heated all of the IGM to temperatures above the CMB. The 21-cm signal becomes insensitive to the spin temperature. Emission in 21-cm is now at its strongest before reionization begins in the earnest. The 21-cm power spectrum is driven by the fluctuations in the density field. In our fiducial model, this epoch corresponds to $16 \lesssim z \lesssim 18$.

(viii) **Reionization:** Ionizing photons from early generations of sources begin permeating the Universe, wiping-out the 21-cm signal inside ionized regions. The power spectrum initially drops on large scales at $\bar{x}_{\text{H I}} \gtrsim 0.9$ as the first regions to be ionized are the small-scale overdensities (McQuinn et al. 2007). The mean signal decreases as H II regions grow, and the power spectrum is governed by H II morphology. This epoch can have other interesting features depending on the detailed evolution of the sources and sinks of ionizing photons, as well as feedback processes, but as the focus of this section is the pre-reionization regime, we shall be brief in this point. In our fiducial model, this epoch corresponds to $7 \lesssim z \lesssim 16$.

These milestones are fairly general, and should appear in most regions of astrophysical parameter space. However, the details of the signal, as well as the precise timing and duration of the above epochs depends sensitively on uncertain astrophysical parameters. For example, note that the above epochs in our fiducial model are shifted to higher redshifts than the analogous ones in Furlanetto (2006) and Pritchard & Furlanetto (2007). This is because the source abundances in those works were computed with PS, which underpredicts the abundances of $T_{\text{vir}} > 10^4$ K haloes by over an order of magnitude at high redshifts (e.g. Trac & Cen




Figure 11. Top panel: Evolution of the mean temperatures from 21CMFAST in our fiducial, $\zeta_X = 10^{57} M_\odot^{-1}$ model (blue curves), and a model with a hundred times weaker X-ray heating, $\zeta_X = 10^{55} M_\odot^{-1}$ (red curves). Solid, dashed and dotted curves show T_s , T_k and T_γ , respectively. Bottom panel: The corresponding evolution in \bar{x}_e and $\bar{x}_{\text{H I}}$.

2007). A similar effect is seen when compared to the recent numerical simulations of Baek et al. (2010), who were only able to resolve haloes with mass $> 10^{10} M_\odot$, which is 2–3 orders of magnitude away from the atomic cooling threshold at these redshifts. With such rare sources, they heat the gas in their $100 h^{-1}$ Mpc boxes to above the CMB temperature much later.

Thorough investigation of the available parameter space is beyond the scope of this work. However, just to briefly show an alternative evolution, in Fig. 11, we include a model where the X-ray efficiency is 2 orders of magnitude lower than in our fiducial model. As one would expect, the Ly α pumping epoch is unaffected. However, X-ray heating is delayed by $\Delta z \sim 7$. In such an extreme model, the 21-cm signal would be seen in strong absorption against the CMB for a long time, and the X-ray heating epoch would overlap with the early stages of the reionization epoch.

Finally, in Fig. 12, we show slices through our fiducial δT_b box (left), and the corresponding 3D power spectra (right). The slices were chosen to highlight various epochs in cosmic 21-cm signal discussed above: the onset of Ly α pumping, the onset of X-ray heating, the completion of X-ray heating and the mid-point of reionization are shown from top to bottom. We encourage the interested reader to see more evolutionary stages through the movie at <http://www.astro.princeton.edu/~mesinger/Movies/delT.mov>. When normalized to the same epoch, our power-spectrum evolution agrees fairly well with the analytical model of Pritchard & Furlanetto (2007), as well as its application to a numerical (Santos et al. 2008) and a seminumerical (MF07) simulation (Santos et al. 2009, cf. their fig. 11).

One interesting feature worth mentioning is that our reionization power spectra show a drop in power on large scales, which persists throughout reionization. In the advanced stages of reionization ($\bar{x}_{\text{H I}} \lesssim 0.9$), the strongest imprint on the 21-cm power spectrum is from H II morphology, with a ‘shoulder’ feature quickly propagating from small to large scales, and flattening the power spectrum (see, for example, Fig. 9 and McQuinn et al. 2007). However there

²⁴ Note that we discuss global trends here. Locally around each Ly α source, there are partially ionized regions hosting hotter gas (e.g. Cen 2006).

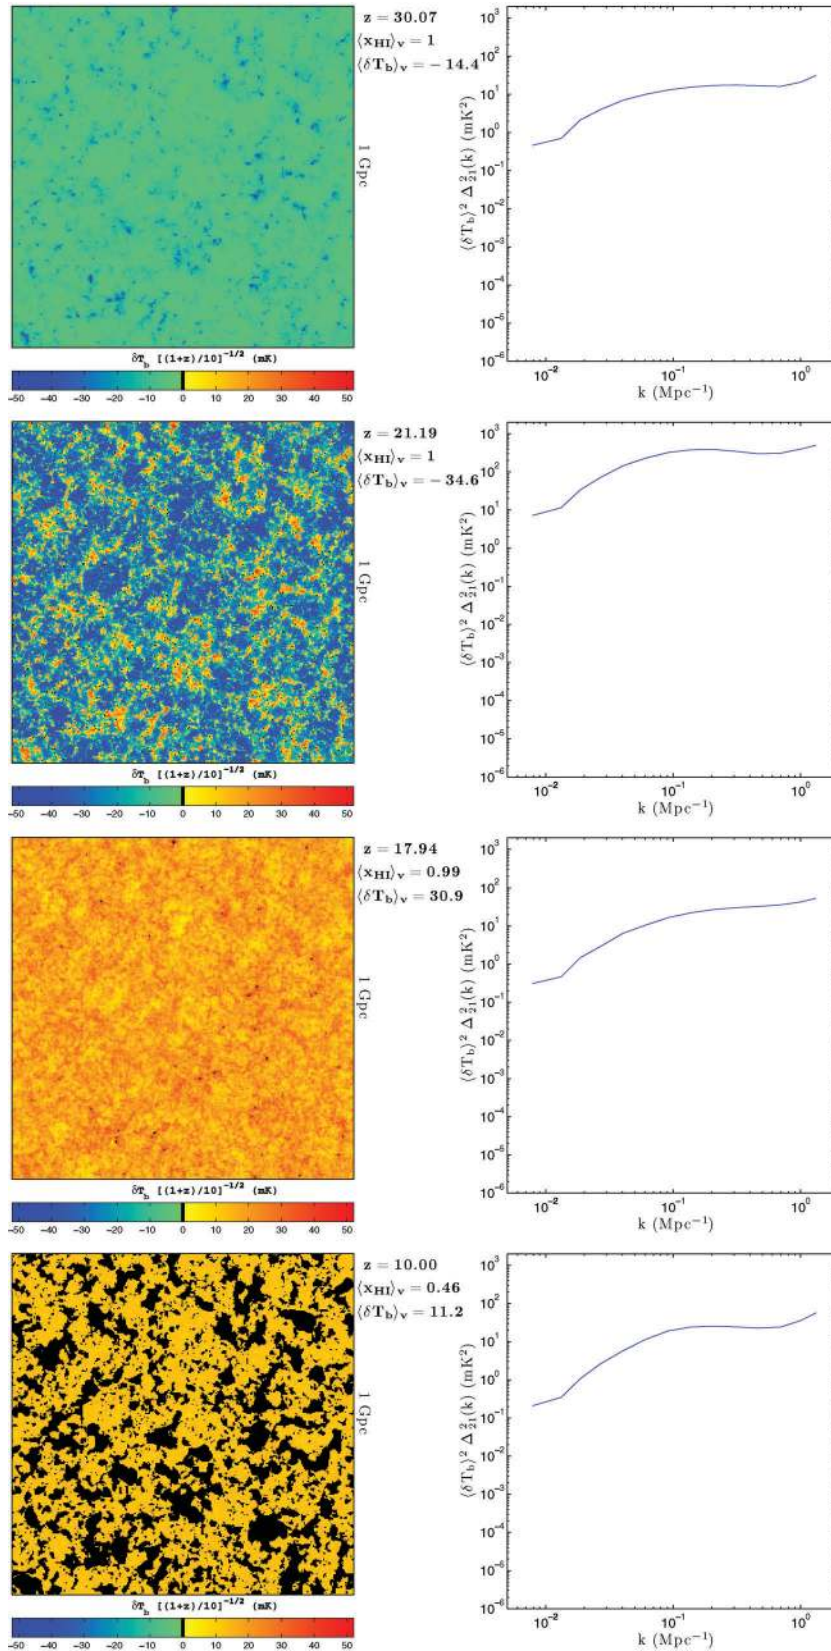


Figure 12. Slices through our δT_b simulation box (left), and the corresponding 3D power spectra (right), for our fiducial model at $z = 30.1, 21.2, 17.9, 10.0$ (top to bottom). The slices were chosen to highlight various epochs in the cosmic 21-cm signal (cf. the corresponding mean evolution in Fig. 10): the onset of Ly α pumping, the onset of X-ray heating, the completion of X-ray heating and the mid-point of reionization are shown from top to bottom. All slices are 1 Gpc on a side and 3.3 Mpc deep. For a movie of this simulation, see <http://www.astro.princeton.edu/~mesinger/Movies/delT.mov>.

should still be a drop in large-scale power beyond either the H II bubble scale, or the photon mean free path in the ionized IGM,²⁵ whichever is smallest. This is an interesting feature from which one can deduce the ionizing photon mean free path in the late stages of reionization (Mesinger & Dijkstra 2008; Furlanetto & Mesinger 2009). And since it occurs at $k \lesssim 0.1 \text{ Mpc}^{-1}$ (for our fiducial choice of parameters), it is beyond the dynamic range of present-day numerical simulations hoping to resolve atomically cooled source haloes.

4 CONCLUSIONS

We introduce a powerful new seminumeric modelling tool, 21CMFAST, designed to efficiently simulate the cosmological 21-cm signal. Our approach uses perturbation theory, excursion set formalism and analytic prescriptions to generate evolved density, ionization, peculiar velocity and spin temperature fields, which it then combines to compute the 21-cm brightness temperature. This code is based on the seminumerical simulation, DEXM (MF07). However, here we bypass the halo finder and operate directly on the evolved density field, thereby increasing the speed and decreasing the memory requirements. In the post-heating regime, 21CMFAST can generate a realization in a few minutes on a single processor, compared to many days on >1000-node supercomputing cluster required to generate the same resolution boxes using state-of-the-art numerical simulations. 21CMFAST realizations in the pre-heating regime require ~ 1 d of computation time. Conversely, RT simulations of the pre-heating regime which resolve most sources currently do not exist, as they are too computationally expensive. Our code is publicly available at <http://www.astro.princeton.edu/~mesinger/Sim.html>

We compare maps, PDFs and power spectra from 21CMFAST, with corresponding ones from the hydrodynamic numerical simulations of Trac et al. (2008), generated from the same initial conditions. We find good agreement with the numerical simulation on scales pertinent to the upcoming observations ($\gtrsim 1$ Mpc). The power spectra from 21CMFAST agree with those generated from the numerical simulation to within 10s of per cent down to the Nyquist frequency.

We find evidence that non-linear peculiar velocity effects enhance the 21-cm power spectrum, beyond the expected geometric, linear value. This enhancement quickly diminishes during the onset of reionization, remaining only on small scales at $\bar{x}_{\text{H I}} \lesssim 0.7$. Interestingly, we also find that the large-scale power is *decreased* as a result of peculiar velocities in the advanced stages of reionization. This is due to the ‘inside-out’ nature of reionization on large scales: the remaining H I regions are preferentially underdense, in which peculiar velocities decrease the 21-cm optical depth and brightness temperature.

Our code can also simulate the pre-reionization regime, including the astrophysical processes of X-ray heating and the WF effect. We show results from a 1-Gpc simulation which tracks the cosmic 21-cm signal down from $z = 250$, highlighting the various interesting epochs.

There are several large 21-cm interferometers scheduled to become operational soon. Interpreting their upcoming data will be difficult since we know very little about the astrophysical processes at high redshifts. Furthermore, there is an enormous range of scales involved. Conventional numerical simulations are too slow for efficient parameter exploration. 21CMFAST is not.

²⁵ Note that RT simulations of reionization generally do not explicitly include an effective ionizing photon mean free path from unresolved LLSs.

ACKNOWLEDGMENTS

We thank Hy Trac for permission to use the initial conditions and the output from his radiative transfer hydrodynamic simulation. We thank Dave Spiegel for his considerable Matlab experience, without which making 21-cm movies would have been considerably more difficult. Support for this work was provided by NASA through Hubble Fellowship grant *HST*-HF-51245.01-A to AM, awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS 5-26555. SRF was partially supported by the David and Lucile Packard Foundation and by NASA through the LUNAR program. The LUNAR consortium (<http://lunar.colorado.edu>), headquartered at the University of Colorado, is funded by the NASA Lunar Science Institute (via Cooperative Agreement NNA09DB30A) to investigate concepts for astrophysical observatories on the Moon. Support was also partially provided by NASA grants NNG06GI09G and NNX08AH31G to RC.

REFERENCES

- Alvarez M. A., Busha M., Abel T., Wechsler R. H., 2009, *ApJ*, 703, L167
 Baek S., Di Matteo P., Semelin B., Combes F., Revaz Y., 2009, *A&A*, 495, 389
 Baek S., Semelin B., Di Matteo P., Revaz Y., Combes F., 2010, *A&A*, preprint (arXiv:1003.0834)
 Barkana R., 2009, *MNRAS*, 397, 1454
 Barkana R., Loeb A., 2004, *ApJ*, 609, 474
 Barkana R., Loeb A., 2005a, *ApJ*, 624, L65
 Barkana R., Loeb A., 2005b, *ApJ*, 626, 1
 Barkana R., Loeb A., 2008, *MNRAS*, 384, 1069
 Bharadwaj S., Ali S. S., 2004, *MNRAS*, 352, 142
 Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, *ApJ*, 379, 440
 Bouwens R. J., Illingworth G. D., Franx M., Ford H., 2008, *ApJ*, 686, 230
 Bouwens R. J. et al., 2010, *ApJ*, 709, L133
 Bowman J. D., Morales M. F., Hewitt J. N., 2005, *BAAS*, 37, 1217
 Cen R., 2006, *ApJ*, 648, 47
 Chen X., Miralda Escudé J., 2008, *ApJ*, 684, 18
 Choudhury T. R., Ferrara A., Gallerani S., 2008, *MNRAS*, 385, L58
 Choudhury T. R., Haehnelt M. G., Regan J., 2009, *MNRAS*, 394, 960
 Cusumano G. et al., 2006, *Nat*, 440, 164
 Fan X. et al., 2006, *AJ*, 132, 117
 Field G. B., 1958, *Proc. Inst. Radio Eng.*, 46, 240
 Field G. B., 1959, *ApJ*, 129, 536
 Furlanetto S. R., 2006, *MNRAS*, 371, 867
 Furlanetto S. R., Furlanetto M. R., 2007, *MNRAS*, 374, 547
 Furlanetto S. R., Loeb A., 2004, *ApJ*, 611, 642
 Furlanetto S. R., Mesinger A., 2009, *MNRAS*, 394, 1667
 Furlanetto S., Stoever S. J., 2010, *MNRAS*, 404, 1869
 Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004a, *ApJ*, 613, 16
 Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004b, *ApJ*, 613, 1
 Furlanetto S. R., Oh S. P., Briggs F. H., 2006a, *Phys. Rep.*, 433, 181
 Furlanetto S. R., Oh S. P., Pierpaoli E., 2006b, *Phys. Rev. D*, 74, 103502
 Geil P. M., Wyithe J. S. B., 2008, *MNRAS*, 386, 1683
 Gnedin N. Y., Shaver P. A., 2004, *ApJ*, 608, 611
 Greiner J. et al., 2009, *ApJ*, 693, 1610
 Harker G. et al., 2010, *MNRAS*, 405, 2492
 Hirata C. M., 2006, *MNRAS*, 367, 259
 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
 Jenkins A., Frenk C. S., White S. D. M., Colberg J. M., Cole S., Evrard A. E., Couchman H. M. P., Yoshida N., 2001, *MNRAS*, 321, 372
 Kaiser N., 1987, *MNRAS*, 227, 1
 Kashikawa N. et al., 2006, *ApJ*, 648, 7

- Komatsu E. et al., 2009, *ApJS*, 180, 330
 Kuhlen M., Madau P., Montgomery R., 2006, *ApJ*, 637, L1
 Lacey C., Cole S., 1993, *MNRAS*, 262, 627
 Lidz A., McQuinn M., Zaldarriaga M., Hernquist L., Dutta S., 2007, *ApJ*, 670, 39
 McLure R. J., Dunlop J. S., Civasuolo M., Koekemoer A. M., Sabbi E., Stark D. P., Targett T. A., Ellis R. S., 2010, *MNRAS*, 403, 960
 McQuinn M., Zahn O., Zaldarriaga M., Hernquist L., Furlanetto S. R., 2006, *ApJ*, 653, 815
 McQuinn M., Lidz A., Zahn O., Dutta S., Hernquist L., Zaldarriaga M., 2007, *MNRAS*, 377, 1043
 Mesinger A., 2010, *MNRAS*, 407, 1328
 Mesinger A., Dijkstra M., 2008, *MNRAS*, 390, 1071
 Mesinger A., Furlanetto S., 2007, *ApJ*, 669, 663 (MF07)
 Mesinger A., Furlanetto S., 2009, *MNRAS*, 400, 1461
 Miralda Escudé J., 2003, *ApJ*, 597, 66
 Morales M. F., Wyithe J. S. B., 2010, *ARA&A*, 48, 127
 Ouchi M. et al., 2009, *ApJ*, 706, 1136
 Parsons A. R. et al., 2009, *AJ*, 139, 1468
 Peebles P. J. E., 1993, *Principles of Physical Cosmology*. Princeton Univ. Press, Princeton, NJ
 Pen U.-L., Chang T.-C., Hirata C. M., Peterson J. B., Roy J., Gupta Y., Odegova J., Sigurdson K., 2009, *MNRAS*, 399, 181
 Pritchard J. R., Furlanetto S. R., 2006, *MNRAS*, 367, 1057
 Pritchard J. R., Furlanetto S. R., 2007, *MNRAS*, 376, 1680
 Salvaterra R. et al., 2009, *Nat*, 461, 1258
 Santos M. G., Amblard A., Pritchard J., Trac H., Cen R., Cooray A., 2008, *ApJ*, 689, 1
 Santos M. G., Ferramacho L., Silva M. B., Amblard A., Cooray A., 2010, *MNRAS*, 406, 2421
 Scoccimarro R., 2004, *Phys. Rev. D*, 70, 083007
 Seager S., Sasselov D. D., Scott D., 1999, *ApJ*, 523, L1
 Seager S., Sasselov D. D., Scott D., 2000, *ApJS*, 128, 407
 Shimasaku K. et al., 2006, *PASJ*, 58, 313
 Shull J. M., van Steenberg M. E., 1985, *ApJ*, 298, 268
 Somerville R. S., Kolatt T. S., 1999, *MNRAS*, 305, 1
 Storrie-Lombardi L. J., McMahon R. G., Irwin M. J., Hazard C., 1994, *ApJ*, 427, L13
 Thomas R. M. et al., 2009, *MNRAS*, 393, 32
 Trac H., Cen R., 2007, *ApJ*, 671, 1
 Trac H., Gnedin N. Y., 2009, *ASL*, preprint (arXiv:0906.4348)
 Trac H., Cen R., Loeb A., 2008, *ApJ*, 689, L81
 Volonteri M., Gnedin N. Y., 2009, *ApJ*, 703, 2113
 Wouthuysen S. A., 1952, *AJ*, 57, 31
 Wyithe J. S. B., Loeb A., 2003, *ApJ*, 586, 693
 Zahn O., Zaldarriaga M., Hernquist L., McQuinn M., 2005, *ApJ*, 630, 657
 Zahn O., Lidz A., McQuinn M., Dutta S., Hernquist L., Zaldarriaga M., Furlanetto S. R., 2007, *ApJ*, 654, 12
 Zahn O., Mesinger A., McQuinn M., Trac H., Cen R., Hernquist L. E., 2010, preprint (ArXiv1003.3455)
 Zel'Dovich Y. B., 1970, *A&A*, 5, 84
 Zygelman B., 2005, *ApJ*, 622, 1356

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