THE ASTROPHYSICAL JOURNAL, 650:529–537, 2006 October 20

21 cm TOMOGRAPHY WITH FOREGROUNDS

XIAOMIN WANG
Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637; and Department of Physics, University of Pennsylvania, Philadelphia, PA 19104; xiaomin@cfcp.uchicago.edu

MAX TEGMARK
Department of Physics, Massachusetts Institute of Technology, 70 Vassar Street, Room 37-626B, Cambridge, MA 02139; and Department of Physics, University of Pennsylvania, Philadelphia, PA 19104

MÁRIO G. SANTOS
CENTRA, Instituto Superior Tecnico, Universidade Técnica de Lisboa, 1049-001 Lisboa, Portugal; and Department of Physics, University of California at Davis, Davis, CA 95616

AND

LLOYD KNOX
Department of Physics, University of California at Davis, Davis, CA 95616
Received 2005 January 25; accepted 2006 June 8

ABSTRACT

Twenty-one centimeter tomography is emerging as a powerful tool to explore the reionization epoch and cosmological parameters, but it will only be as good as our ability to accurately model and remove astrophysical foreground contamination. Previous treatments of this problem have focused on the angular structure of the signal and foregrounds and what can be achieved with limited spectral resolution (channel widths in the 1 MHz range). In this paper we introduce and evaluate a "blind" method to extract the multifrequency 21 cm signal by taking advantage of the smooth frequency structure of the Galactic and extragalactic foregrounds. We find that 21 cm tomography is typically limited by foregrounds on scales of $k \ll 1\ h\ \mathrm{Mpc}^{-1}$ and is limited by noise on scales of $k \gg 1\ h\ \mathrm{Mpc}^{-1}$, provided that the experimental channel width can be made substantially smaller than 0.1 MHz. Our results show that this approach is quite promising even for scenarios with rather extreme contamination from point sources and diffuse Galactic emission, which bodes well for upcoming experiments such as LOFAR, MWA, PAST, and SKA.

Subject headings: cosmology: theory — diffuse radiation — ISM: atoms — methods: analytical — methods: data analysis — radio lines: general

Online material: color figures

1. INTRODUCTION

Twenty-one centimeter tomography is one of the most promising cosmological probes, with the potential to complement and perhaps ultimately eclipse the cosmological parameter constraints from the cosmic microwave background (CMB; Bowman et al. 2006; McQuinn et al. 2006). It is also a unique probe of the epoch of reionization, which is now one of the least understood aspects of modern cosmology. There are various techniques to explore the epoch of reionization at $5 < z < 20$. Apart from the CMB (Holder et al. 2003; Knox 2003; Kogut et al. 2003; Santos et al. 2003), radio astronomical measurement of 21 cm radiation from neutral hydrogen has been shown theoretically to be a powerful tool to study this period (Madau et al. 1997; Tozzi et al. 2000). Lots of work has been done in recent years on various theoretical and experimental aspects of 21 cm radiation (e.g., Barkana & Loeb 2005b; Carilli et al. 2002, 2004; Ciardi & Madau 2003; Di Matteo et al. 2002, 2004; Furlanetto et al. 2004a, 2004b; Gnedin & Shaver 2004; Iliev et al. 2002, 2003; Loeb & Zaldarriaga 2004; McQuinn et al. 2006; Morales 2005; Oh & Mack 2003; Pen et al. 2004; Santos et al. 2005; Shaver et al. 1999; Wyithe & Loeb 2004a, 2004b; Zaldarriaga et al. 2004).

However, this 21 cm tomography technique will only be as good as our ability to accurately model and remove astrophysical foreground contamination, since the high-redshift signal one is looking for is quite small and can be easily swamped by foreground emission from our galaxy or others. With much effort going into upcoming experiments such as the Mileura Wide-Field Array (MWA),[1] the Low Frequency Array (LOFAR; Röttgering et al. 2003),[2] the Primeval Structure Telescope (PAST),[3] and the Square Kilometer Array (SKA),[4] which are aimed at gathering redshifted 21 cm signal from the sky and probing the epoch of reionization, it is therefore timely to study the foreground problem in detail.

Although 21 cm foregrounds have been discussed in some previous papers (e.g., Di Matteo et al. 2002, 2004; Morales & Hewitt 2004; Oh & Mack 2003; Santos et al. 2005; Zaldarriaga et al. 2004), the questions on how to remove foregrounds and noise from observations of the 21 cm signal, how well it can be done, and how reliable it is are still wide open. Previous papers have focused on the angular power spectrum of the signal, usually assuming a rather limited spectral resolution (Di Matteo et al. 2002, 2004; Oh & Mack 2003; Santos et al. 2005; Zaldarriaga et al. 2004). In this paper, we develop a method to remove the foregrounds along the line of sight, taking advantage of the fact that most astrophysical contaminants have much smoother frequency

[1] See http://web.haystack.mit.edu/MWA/MWA.html.
[2] See http://www.lofar.org.
[3] See http://astrophysics.phys.cmu.edu/~past.
[4] See http://www.skatelescope.org.

spectra than the cosmological signal one is looking for. The two approaches are complementary, and we argue that they are best used in combination: our technique can be used both to identify point sources and other highly contaminated angular regions to be discarded and to clean out residual contamination from those angular regions that are not discarded. This multifrequency approach is more powerful here than for typical CMB applications (Bennett et al. 2003; Tegmark 1998; Tegmark et al. 2003) because of the potentially much better spectral resolution and the dramatically oscillating 21 cm signal compared with smooth foregrounds along the frequency direction.

In this paper, we describe the method for removing foregrounds in frequency space, show examples of using this method in different scenarios, and discuss its promising applications for future experiments. In § 2, we introduce the reionization model that we use throughout the text, then give a brief overview of the 21 cm emission/absorption and computational formalism on how we calculate the 21 cm angular power spectrum in l-space, projected one-dimensional (1D) and two-dimensional (2D) power spectra in k-space, and the simulated 1D frequency spectrum in real space. In § 3, we describe our foreground-removing strategy, and we also show the foreground model we use in our calculations. In § 4 we give several applications of our method under different assumptions about foregrounds and noise. We summarize our results in § 5.

2. REIONIZATION MODEL AND FORMALISM

The reionization model we use throughout this paper is from Haiman & Holder (2003) and Santos et al. (2003), shown as the solid curve in Figure 3 in Santos et al. (2003). Although the most recent results from the *Wilkinson Microwave Anisotropy Probe* (*WMAP*; Spergel et al. 2006) favor a lower optical depth, the results of this paper are rather insensitive to the detailed choice of the reionization model and associated assumptions, since we are focused on foregrounds rather than on the cosmological 21 cm signal. For more information about various reionization models, see, for example, Haiman & Holder (2003), Holder et al. (2003), and Santos et al. (2003).

Below we give a brief overview of the 21 cm emission power spectrum and our calculational method. The detailed information on 21 cm radiation (emission/absorption) phenomena can be found in the literature (e.g., Madau et al. 1997; Santos et al. 2005; Shaver et al. 1999; Tozzi et al. 2000; Zaldarriaga et al. 2004).

2.1. *Three-dimensional Power Spectrum*

The differential antenna temperature observed at Earth between the neutral hydrogen patch and the CMB can be approximated as (Shaver et al. 1999; Tozzi et al. 2000)

$$\delta T_b \approx (0.016\ \mathrm{K})\frac{1}{h}(1+\delta)(1-x)\frac{\Omega_b h^2}{0.02}\left(\frac{1+z}{10}\frac{0.3}{\Omega_m}\right)^{1/2}, \quad (1)$$

where $\delta = \rho/\bar{\rho} - 1$ is the fluctuation of the density field. We write the ionization fraction x as a sum of two terms (Santos et al. 2003),

$$x = x_e(1+\delta_{x_e}), \quad (2)$$

where x_e is the average ionization fraction and δ_{x_e} is the fluctuation of the ionization fraction across the sky. Thus, the ionization fraction x is not only a function of redshift, but also dependent on its position in the sky.

Assuming $\delta \ll 1$ and $\delta_{x_e} \ll 1$ and neglecting all second- and higher order terms, we obtain the three-dimensional (3D) power spectrum for 21 cm emission,

$$P_{\mathrm{3D}}(k,z) = (0.016\ \mathrm{K})^2\frac{1}{h^2}\left(\frac{\Omega_b h^2}{0.02}\right)^2\frac{1+z}{10}\frac{0.3}{\Omega_m} \times\left(1-2x_e+x_e^2+b^2e^{-k^2R^2}x_e^2\right)P_{\mathrm{matter}}(k), \quad (3)$$

where $P_{\mathrm{matter}}(k)$ is the matter power spectrum. The power spectrum for the ionized fraction $P_{\delta_{x_e}}(k)$ is defined as in Santos et al. (2003),

$$P_{\delta_{x_e}}(k) = b^2 P_{\mathrm{matter}}(k)e^{-k^2R^2}, \quad (4)$$

where b is the mean bias weighted by the different halo properties. The mean radius R of the ionized patches in H II regions is modeled as

$$R = \left(\frac{1}{1-x_e}\right)^{1/3}R_p, \quad (5)$$

where R_p is the comoving size of the fundamental patch, with $R_p \sim 100$ kpc (Santos et al. 2003).

We assume a cosmological concordance model (Seljak et al. 2005; Spergel et al. 2003; Tegmark et al. 2004; $\Omega_k = 0$, $\Omega_\Lambda = 0.71$, $\Omega_b = 0.047$, $h = 0.72$, $n_s = 0.99$, $\sigma_8 = 0.9$) throughout our calculations.

2.2. *Projected 1D Power Spectra*

The 21 cm signal changes with redshift for two separate reasons, one slow and one fast:

1. The average properties of the universe (x, T_k, T_s, etc.) evolve on a timescale of $\Delta z \sim 1$.
2. The local properties of the universe change on much smaller scales of Δz, corresponding to the sizes and separations between ionized regions.

Across a very small redshift range $z_0 - \Delta z < z < z_0 + \Delta z$, where $\Delta z < 1$, we make the approximation of ignoring the former and including only the latter, approximating parameters like x, T_k, and T_s by their values at z_0. This enables us to linearize the relations between frequency ν, redshift z, and comoving radial distance D.

If we make the above-mentioned approximation and ignoring redshift space distortions, the 21 cm signal near a given z_0 has an isotropic 3D power spectrum P_{3D} that we can project into 1D (radial) power spectra $P_{\mathrm{1D}}(k,z_0)$ (Hui et al. 1999; Peacock 1999):

$$P_{\mathrm{1D}}(k, z\approx z_0) = \frac{1}{2\pi}\int_k^\infty P_{\mathrm{3D}}(k', z\approx z_0)k'\,dk'. \quad (6)$$

Figure 1 shows the line-of-sight 1D 21 cm emission power spectra for the fiducial reionization model at different reionization epochs. For comparison, we also plot the power spectrum for the neutral medium at $x = 8.09$ (*thick solid line*).

2.3. *Simulated Signal in Real Space from 1D Power Spectrum*

We generate and analyze our simulations with fast Fourier transforms (FFTs). The simulated signal in real space in the region $0 < x < L$ is

$$f(x) = \frac{1}{\sqrt{N}}\sum_{q=0}^{N-1}\left[A_q\cos\left(\frac{2\pi x}{L}q\right) + B_q\sin\left(\frac{2\pi x}{L}q\right)\right], \quad (7)$$

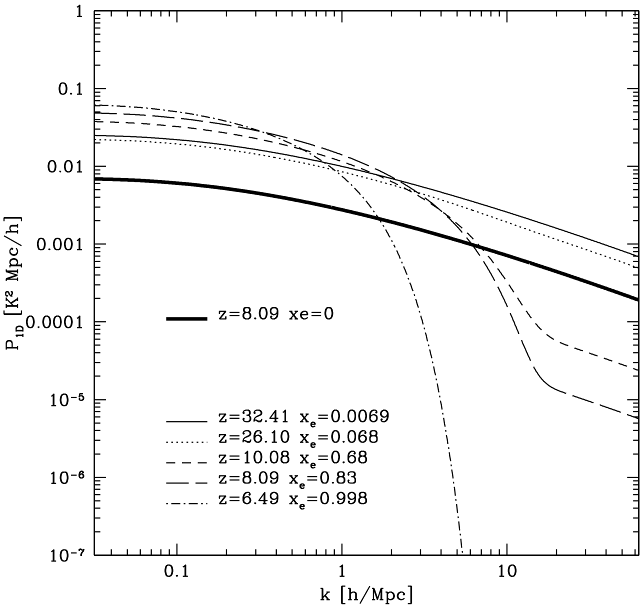

Fig. 1.—Line-of-sight 1D 21 cm power spectra at different redshifts and ionization fractions. The thick solid line indicates the power spectrum for the neutral intergalactic medium (IGM) at redshift 8.09.

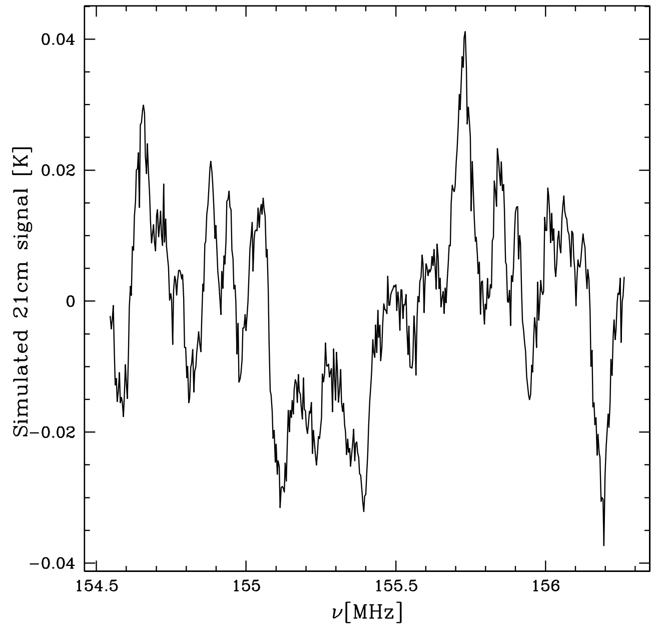

Fig. 2.—Simulated 21 cm signal in frequency space corresponding to $z \approx 8.09$ and $x_e \approx 0.83$.

where A_q and B_q are Gaussian random variables with zero mean and standard deviations $\Delta A_q = \Delta B_q = [P_{1D}(k, z_0)/2]^{1/2} = [P_{1D}(2\pi q/L, z_0)/2]^{1/2}$, and N is chosen to be a large enough integer that all information from $P_{1D}(k)$ is included in the summation ($k_{\max} = 2\pi N/L$). The box size L should be small enough that the range of L satisfies $\Delta z \ll 1$.

Figure 2 shows our simulated 21 cm signal versus frequency around 155 MHz, corresponding to an epoch around redshift $z \approx 8.09$. If we plotted the *observed* signal in the relevant frequency range, the expected contribution from foregrounds would lie far above the cosmological signal shown in Figure 2. The key to doing 21 cm cosmology is, therefore, removing foregrounds using multifrequency information, as emphasized by, for example, Zaldarriaga et al. (2004), and we now turn to this subject.

3. METHOD FOR FOREGROUND REMOVAL IN FREQUENCY SPACE

Because of its small frequency cross-correlations, the 21 cm signal oscillates dramatically along the frequency direction. The foregrounds, on the other hand, are generally quite smooth over the short frequency range we consider. This slowly varying nature of the foregrounds compared to the signal is a great advantage when removing them (Gnedin & Shaver 2004; McQuinn et al. 2006; Morales et al. 2006), and it is the main reason that our foreground-removal method works so well. Our method described here is insensitive to the reionization model and the redshift range we choose, since we are focused on the foregrounds rather than the cosmological 21 cm signal.

3.1. *Foreground Removal Method*

Our basic approach is to subtract foregrounds separately in each angular direction in the sky by fitting their total intensity dependence on frequency by a log-log polynomial. Note that since we are fitting the total foreground spectrum separately pixel by pixel (fitting not only for the amplitude, but also for the spectral index and the running of the spectral index), we are unaffected by the possible complication of huge variations of the foreground spectral index across the sky. (If the foregrounds lacked both frequency coherence and spatial coherence, i.e., fluctuated randomly with both frequency and position, then we would be unable to identify and remove them and could merely average them down, as we do with detector noise.)

There are two separate steps in our analysis: (1) simulation and (2) cleaning. We treat them as completely independent. In other words, our cleaning algorithm is "blind," containing no information about the foreground and noise model used in the simulation step. It is entirely specified by the single integer m, which gives the order of the log-log fitting polynomial.

In the simulation step, we simulate for each pixel the total observed frequency spectrum y_i at n different log-frequencies $x_i = \log(\nu_i)$, $i = 1, \ldots, n$. This simulated total signal $y_i = \log(I^i_{21\,\mathrm{cm}} + I^i_{\mathrm{syn}} + I^i_{\mathrm{ff}} + I^i_{\mathrm{ps}} + I^i_{\mathrm{det}})$ includes 21 cm signal $I^i_{21\,\mathrm{cm}}$, synchrotron emission foreground I^i_{syn}, free-free emission foreground I^i_{ff}, point-source foreground I^i_{ps}, and detector noise I^i_{det}. We test a variety of different assumptions for foregrounds and noise in this step.

Then we group the y_i into an n-dimensional vector $\boldsymbol{y}$ and group the x_i and their powers into an $n \times m$ matrix $\mathbf{X}$ so that the data can be modeled as

$$\boldsymbol{y} = \mathbf{X}\boldsymbol{a} + \boldsymbol{n}, \tag{8}$$

where the m-dimensional parameter vector $\boldsymbol{a}$ parameterizes the foreground contributions. It is what we need to find out in the cleaning step. In equation (8), $\boldsymbol{n}$ is the part left in the total signal that cannot be fitted by the parameters in $\boldsymbol{a}$, including the contributions from signal, detector noise, and residual foregrounds.

In all our calculations throughout this paper, we take this fitting polynomial to be quadratic; that is, we fit the total foregrounds as a single running power law. Equivalently, we fit the log intensity of the foreground as $\log I = a_3 + a_2 \log\nu + a_1(\log\nu)^2$. That is to say, the number of the fitting parameter in our computation is always $m = 3$. We found that this improved noticeably over $m = 2$, whereas increasing to $m = 4$ gave essentially no further improvement.

In the cleaning step, we find our best-fit parameter vector $\boldsymbol{a}$ by minimizing $\chi^2 = (\boldsymbol{y} - \mathbf{X}\boldsymbol{a})^t\mathbf{N}^{-1}(\boldsymbol{y} - \mathbf{X}\boldsymbol{a})$, obtaining (Tegmark 1997)

$$\boldsymbol{a} = \left(\mathbf{X}^t\mathbf{N}^{-1}\mathbf{X}\right)^{-1}\mathbf{X}^t\mathbf{N}^{-1}\boldsymbol{y}, \tag{9}$$

where $\mathbf{N}$ is the covariance matrix of the contributions from the detector noise and the 21 cm signal. We then subtract the total fitted foreground contribution $\mathbf{X}\boldsymbol{a}$ from the simulated measurement vector, thus obtaining what we will refer to as the cleaned signal $\boldsymbol{y}$.

Although this cleaning technique is only optimal if $\mathbf{N}$ is known and the contributions from noise and 21 cm signals are Gaussian (Tegmark 1997), we use equation (9) anyway and quantify the residual noise using our simulations. Since equation (9) minimizes the rms residual even in the presence of non-Gaussianity (Tegmark 1997), it is a robust general-purpose fit that does not require detailed foreground or signal modeling. We simply set $\mathbf{N} = \mathbf{I}$, the identity matrix, which will be essentially optimal if white detector noise dominates. If desired, this can be further improved by modeling the foreground power spectrum found in real data and iterating.

Since the cleaning step uses a single polynomial in log-log space, it cannot fit exactly a simulation including detector noise or more than one foreground component (since adding the exponential of two polynomials does not give the exponential of a polynomial). We will see that this simple cleaning algorithm is nonetheless very successful, able to fit any of our foreground models well over the limited frequency range that is relevant.

3.2. *Foreground and Detector Noise Models*

The foregrounds we consider in this paper include Galactic synchrotron emission, Galactic free-free (thermal) emission, and extragalactic point sources. For more information on models of foregrounds in the 100 MHz range, please see, for example, Di Matteo et al. (2002, 2004), Haslam et al. (1982), Haverkorn et al. (2003), Morales & Hewitt (2004), Oh & Mack (2003), Platania et al. (2003), Santos et al. (2005), and Shaver et al. (1999). As emphasized above, the foreground models we describe here are used only in our simulation step, not for the cleaning process.

3.2.1. *Galactic Synchrotron Radiation*

For Galactic synchrotron emission, which probably causes most contamination of all foregrounds (perhaps of order 70% at 150 MHz; Platania et al. 2003; Shaver et al. 1999), we assume its intensity to be a running power law in frequency,

$$I_{\rm syn} = A_{\rm syn}\left(\frac{\nu}{\nu_*}\right)^{-\alpha_{\rm syn}-\Delta\alpha_{\rm syn}\log(\nu/\nu_*)}, \tag{10}$$

with a spectral index $\alpha_{\rm syn} = 2.8$ (Tegmark et al. 2000) and a spectral "running" index $\Delta\alpha_{\rm syn} = 0.1$ (Haverkorn et al. 2003; Platania et al. 2003; Shaver et al. 1999; Tegmark et al. 2000). Here $\nu_* \equiv 150$ MHz. We assume the amplitude of the synchrotron foreground to be $A_{\rm syn} = 335.4$ K, an extrapolation from Haverkorn et al. (2003) and Tegmark et al. (2000). We also explore other normalizations that are $A_{\rm syn}$ orders of magnitude higher than the value we define here in our calculations. Similarly, we try other values of the spectral index and spectral running index in the calculation to test the robustness of our method. We discuss the details in § 4.

3.2.2. *Galactic Free-Free Emission*

We model the Galactic free-free emission (which might contribute a contamination of order 1% at 150 MHz; Shaver et al. 1999) as a running power law as well (Haverkorn et al. 2003; Platania et al. 2003; Tegmark et al. 2000):

$$I_{\rm ff} = A_{\rm ff}\left(\frac{\nu}{\nu_*}\right)^{-\alpha_{\rm ff}-\Delta\alpha_{\rm ff}\log(\nu/\nu_*)}, \tag{11}$$

where $\alpha_{\rm ff} = 2.15$, $\Delta\alpha_{\rm ff} = 0.01$, and $A_{\rm ff} = 33.5$ K, extrapolated from Haverkorn et al. (2003) and Tegmark et al. (2000).

3.2.3. *Extragalactic Point Sources*

Point sources have been estimated to cause about 30% of the contamination at 150 MHz (Shaver et al. 1999) and are typically less smooth in frequency than the Galactic foregrounds. When looking in a given direction, we are observing the same point sources as we change frequency, so there are not small-scale fluctuations in the same sense as when we change observing directions.[5] A serious complication compared to the Galactic synchrotron and free-free cases is that when we observe many point sources in a pixel, they can each have quite different spectral indices, possibly making their combined intensity a quite complicated function of frequency.

One approach would be to model this complicated function as a running power law over the narrow frequency range involved, just as we did for the synchrotron and free-free foregrounds:

$$I_{\rm ps} = A_{\rm ps}\left(\frac{S_{\rm cut}}{\rm mJy}\right)^{\beta}\left(\frac{\nu}{\nu_*}\right)^{-\alpha_{\rm ps}-\Delta\alpha_{\rm ps}\log(\nu/\nu_*)}, \tag{12}$$

where $\alpha_{\rm ps} = 2.81$, $\Delta\alpha_{\rm ps} = 0.25$, and $\beta = 0.125$ (Tegmark et al. 2000).

However, we wish to be as conservative as possible in this paper, and we therefore adopt a more complicated point source model in our simulations. We therefore simulate a large number of random point sources $i = 1, \ldots$ in the pixel that we are considering and sum their intensity contributions in units of kelvins:

$$I_{\rm ps} = \left(\frac{dB}{dT}\right)^{-1}\Omega_{\rm sky}^{-1}\sum_i S_i^*\left(\frac{150\ \rm MHz}{\nu}\right)^{\alpha_i} - \langle I_{\rm ps}\rangle, \tag{13}$$

where

$$\langle I_{\rm ps}\rangle = \left(\frac{dB}{dT}\right)^{-1}\int_0^{S_{\rm cut}} S\frac{dN}{dS}\,dS\int\left(\frac{150}{\nu}\right)^{\alpha} f(\alpha)\,d\alpha \tag{14}$$

is the average value of the point-source foreground intensity. The conversion factor $dB/dT = 6.9\times10^5$ mJy K^{-1}. The assumed sky area per pixel is approximately $\Omega_{\rm sky} = 12$ arcmin^2. In equation (13), S_i^* is the flux of the ith point source at 150 MHz. It is generated randomly from the source count distribution $dN/dS = 4(S/1\ \rm Jy)^{-1.75}$ (Di Matteo et al. 2004), which is truncated at a maximum flux $S_{\rm max} = S_{\rm cut} = 0.1$ mJy, above which we assume that point sources can be detected and their pixels

[5] There is, however, the subtle effect of off-beam point sources dimming toward higher frequencies because the beam gets narrower (Oh & Mack 2003; Zaldarriaga et al. 2004). For the narrow frequency intervals $\Delta\ln\nu$ that we are considering, this effect will be around the percent level for individual sources, in the same ballpark as the intensity change due to the frequency dependence of the emission mechanism. This means that it will not imprint sharp spectral features in the total foreground emission and should be well fitted by our blind cleaning algorithm. We have not included this complication in the present analysis; it would be worth incorporating it in a more detailed foreground analysis, particularly in one including explicit modeling of the sky pixelization.

discarded. In other words, when we talk about the contamination from point sources below, we refer only to the contribution from unresolved point sources. To avoid having to generate infinitely many point sources, we also truncated the distribution at a minimum flux $S_{\rm min} = 10^{-3}$ mJy, since we find that the total flux contribution has converged by then. We generate α_i, the spectral index of the ith point source, randomly from the Gaussian distribution

$$f(\alpha) = \frac{1}{\sqrt{(2\pi)}\sigma_\alpha} \exp\left[-\frac{(\alpha-\alpha_0)^2}{2\sigma_\alpha^2}\right], \tag{15}$$

with the spectral index α in the range of $[\alpha_0 - \Delta\alpha, \alpha_0 + \Delta\alpha]$, where $\Delta\alpha = 5\sigma_\alpha$. To be conservative, we allow the spectral index to vary in a fairly large region, $\sigma_\alpha = 10$, through our calculations.

3.2.4. *Detector Noise*

We treat detector noise as white noise. In the Rayleigh-Jeans limit, the rms detector noise in a pixel can be approximated as

$$\sigma_T = \frac{\lambda^2}{2k_{\rm B}} B = \frac{\lambda^2}{2k_{\rm B}} \frac{S}{A}, \tag{16}$$

where $k_{\rm B}$ is the Boltzmann constant and λ is the redshifted wavelength of 21 cm emission. The specific brightness B is related to the point-source sensitivity S by dividing it with the pixel area A.

At redshift 8.47, $\nu = 150$ MHz, $\lambda = 2m$, with the LOFAR virtual core configuration,[6] for a 5.′2 pixel with 4 MHz bandpass and 1 hr integration, the sensitivity S is approximately 0.17 mJy, and from equation (16) we get

$$\sigma_T^{\rm LOFAR} = (108 \text{ mK}) \left(\frac{4 \text{ MHz}}{\Delta\nu}\right)^{0.5} \left(\frac{1 \text{ hr}}{t}\right)^{0.5}, \tag{17}$$

where $\Delta\nu$ is the channel width and t is the total integration time. Similarly, for the MWA experiment,[7] a 4.′6 pixel with 32 MHz bandpass and 1 hr integration has a point-source sensitivity of $S = 0.27$ mJy, so we get the MWA detector noise of

$$\sigma_T^{\rm MWA} = (218 \text{ mK}) \left(\frac{32 \text{ MHz}}{\Delta\nu}\right)^{0.5} \left(\frac{1 \text{ hr}}{t}\right)^{0.5}. \tag{18}$$

We should mention that although at 4 MHz bandwidth, the sensitivity for MWA is worse than that for LOFAR, MWA has a larger bandpass and field of view. This larger field of view leads to vastly more pixels, which is an advantage for foreground removal, as we will see in later sections. The detector thermal noise is only one of the many concerns in the experiment, such as calibration, systematics, etc. Thus, it should not be considered as the only criterion to judge an experiment.

The 1D power spectrum of the detector noise can then be written as

$$P_{\rm det} = 2\pi\sigma_T^2. \tag{19}$$

In our simulation, we consider two scenarios. One scenario assumes a fiducial future experiment with Gaussian random detector noise down to the $\sigma_T = 1$ mK level. The other scenario assumes a currently achievable detector noise level of $\sim$200 mK. This is based on equations (17) and (18) for the LOFAR and MWA experiments, assuming 1000 hr of integration time and 4–8 kHz frequency resolutions, respectively.

[6] See http://www.lofar.org.

[7] See http://web.haystack.mit.edu/MWA/MWA.html.

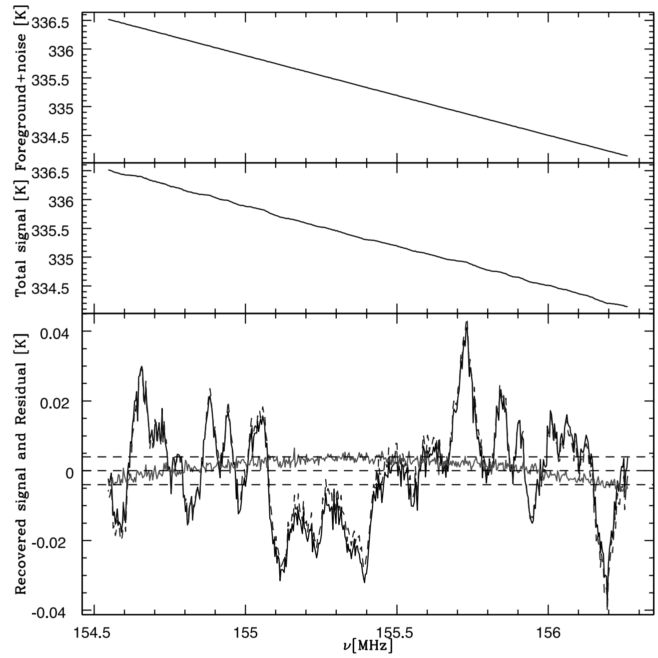

Fig. 3.—Spectrum in a single pixel before and after foreground cleaning. The top panel shows the total contaminant signal, consisting of synchrotron radiation ($A_{\rm syn} = 335.4$ K, $\alpha_{\rm syn} = 2.8$, $\Delta\alpha_{\rm syn} = 0.1$), free-free emission foreground ($A_{\rm ff} = 33.5$ K, $\alpha_{\rm ff} = 2.15$, $\Delta\alpha_{\rm ff} = 0.01$), extragalactic point sources ($\sigma_\alpha = 10$), and detector noise ($\sigma = 1$ mK). The middle panel has the cosmological 21 cm signal added. The bottom panel shows the recovered 21 cm signal (*dashed curve*) compared with the true simulated signal (*solid curve*) and the residual (recovered minus simulated 21 cm signal; *gray curve*). The three horizontal black dashed lines correspond to −0.004, 0, and 0.004 K, respectively. (Note the different vertical axis limits.) The small-scale wiggles in the residual represent detector noise, whereas the smoothed parabola-shaped component of the residual indicates the error in the foreground fitting. [*See the electronic edition of the Journal for a color version of this figure.*]

4. RESULTS

As we showed previously in Figure 2, the signal wiggles rapidly with frequency. This is the key advantage of removing foregrounds in frequency space, since foregrounds are typically relatively smooth functions of frequency.

We simulate the 21 cm signal as a Gaussian random field, although in reality, the signal is of course highly non-Gaussian. We make this Gaussianity approximation for simplicity, since the key quantity that we are interested in (the power spectra of the residual noise and foregrounds) depends mainly on the power spectra of the signal, foregrounds, and noise, not on whether the statistics are Gaussian or not.

4.1. *Baseline Example 1: Long-Term Potential (Noise $\ll$ Signal)*

The results for the baseline example with noise much smaller than the signal are shown in Figure 3. The top panel shows the total contaminant in a pixel, including Galactic synchrotron radiation, Galactic free-free emission, extragalactic point sources, and detector noise with $\sigma = 1$ mK, which is the fiducial value for a future-generation experiment. The foregrounds are modeled as in the previous section, with parameters (given in the figure caption) corresponding to a rather pessimistic assumption about the foreground properties.

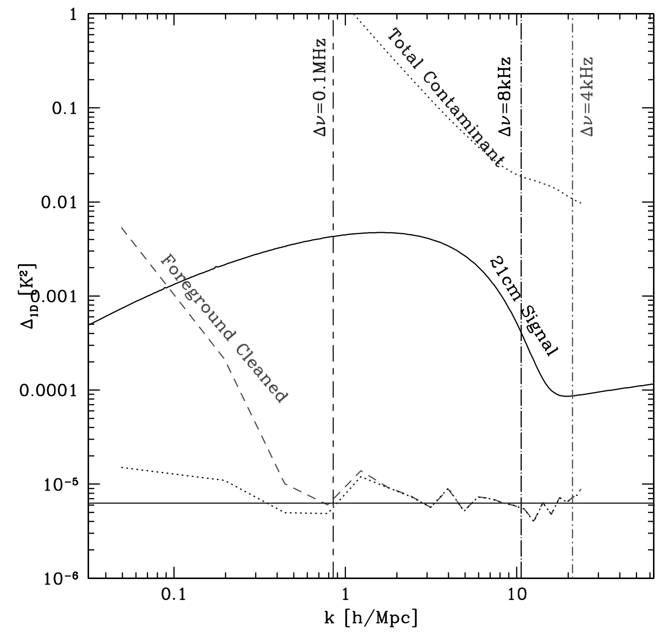

FIG. 4.—One-dimensional power spectra of the 21 cm signal (*solid curve*), the total contaminant from Fig. 3 (*upper dotted curve*), the residual contaminant after foreground cleaning (*dashed gray curve*), and the detector noise alone (*lower dotted curve*). The horizontal solid line shows the white noise power spectrum used for the detector noise simulation. The three vertical lines correspond to the different channel widths for different experiments. From left to right, they are 0.1 MHz (fiducial; *short-dash–long-dashed line*), 8 kHz (for MWA; *dot–long-dashed line*), and 4 kHz (for LOFAR; *dot–short-dashed gray line*). [*See the electronic edition of the Journal for a color version of this figure.*]

The total foreground is so huge that the top panel looks quite similar to the middle panel (which includes the 21 cm signal). From the middle panel, we cannot really tell if there is any signature of 21 cm emission. However, the bottom panel shows that the foregrounds can be effectively removed, with the residual (the recovered 21 cm signal minus the input "true" signal) being more than 5 orders of magnitude below the foregrounds in amplitude.

By transforming the detector noise and the residual from the bottom panel of Figure 3 back to k-space, we are able to compare them with the 21 cm 1D power spectrum, shown in Figure 4. Before foreground cleaning, the total contaminant (*upper dotted curve*) is seen to dominate over the 21 cm power spectrum (*solid curve*). The foreground power spectrum is seen to rise toward the left, reflecting its rather smooth frequency dependence. After foreground cleaning, the residual contaminant (*gray dashed curve*) is significantly below the original contaminant, except on scales of $k \ll 0.1\ h\ \mathrm{Mpc}^{-1}$. The flat section of the residual power spectrum for $k \gtrsim 1\ h\ \mathrm{Mpc}^{-1}$ is seen to correspond to detector noise.

The three vertical lines in Figure 4 correspond to different minimum channel widths for different upcoming experiments. From left to right, they are 0.1 MHz (fiducial), 8 kHz (for MWA[8]), and 4 kHz (for LOFAR[9]). Information to the right of this minimum channel width line is lost, roughly corresponding to an exponential blow-up of the effective detector noise. In other words, the effective detector noise goes much higher above the signal to the right of those lines, so little information can be extracted there. So in order to take advantage of our method of foreground removal, the channel width needs to be small enough to reach the noise-dominated region.

[8] See http://web.haystack.mit.edu/MWA/MWA.html.
[9] See http://www.lofar.org.

In other words, 21 cm tomography is limited mainly by foregrounds for $k \ll 1\ h\ \mathrm{Mpc}^{-1}$ and limited mainly by noise for $k \gg 1\ h\ \mathrm{Mpc}^{-1}$. To take full advantage of their sensitivity by pushing residual foregrounds down to the detector noise levels, experiments should therefore be designed to have a channel width substantially smaller than 0.1 MHz. Such small channel widths are realistic and achievable for upcoming experiments, since the analysis can now practically be done by dedicated high-speed electronics, even if the software solution was not fast enough.

To test the robustness of our foreground cleaning technique, we repeated the above analysis for a wide range of foreground models with the same noise level. First we tested a suite of models with only detector noise and synchrotron radiation, changing the values of the parameters defined in equation (10). Most of the results were similar to those shown in Figure 3. Increasing the synchrotron amplitude parameter A_{syn} all the way up to a completely unrealistic value of 10^7 K had essentially no effect, because in this case the simulated spectrum had the exact same shape as the model for which we fitted in the cleaning step. Likewise, changing the spectral index α_{syn} over the extreme range $-80 < \alpha_{\mathrm{syn}} < 20$ had little effect. Complicating the synchrotron spectrum with a running of the running term $\Delta\alpha_{\mathrm{syn}}^{\mathrm{rr}} \log^2(\nu/\nu_*)$, so that the intensity of the synchrotron foreground can be written as

$$I_{\mathrm{syn}} = A_{\mathrm{syn}} \left(\frac{\nu}{\nu_*}\right)^{-\alpha_{\mathrm{syn}} - \Delta\alpha_{\mathrm{syn}} \log(\nu/\nu_*) - \Delta\alpha_{\mathrm{syn}}^{\mathrm{rr}} \log^2(\nu/\nu_*)}, \quad (20)$$

still caused a negligible increase of the residual as long as $|\Delta\alpha_{\mathrm{syn}}^{\mathrm{rr}}| \leq 50$. This is a reflection of the fact that over a fairly narrow frequency range, even a more complicated function can be accurately approximated by a parabola in log-log space; that is, by the first three terms of its Taylor expansion.

Making variations around the baseline case of multiple foregrounds, we also tried numerous examples with either higher foreground amplitudes, different spectral indices, or larger running spectral indices, again obtaining results similar to the ones shown in Figure 3. For example, for the point-source foreground, we tried different values for the parameter σ_α in equation (13), from $\sigma_\alpha = 0.2$ up to $\sigma_\alpha = 100$. We found that as long as $\sigma_\alpha \leq 60$ (a conservative cut), the residuals are rather insensitive to the distribution of point-source spectral indices.

All our tests show that, as with astrophysically plausible foreground amplitudes, the effectiveness of our simple "blind" cleaning method is almost independent of the number, shape, and amplitude of relatively smooth foregrounds. The basic reason for this robustness is easy to understand. As long as the total foreground signal can be well approximated by our fitting function (a log-log parabola) over the small frequency interval in question, then the main contribution to the residual will not be the foregrounds, but the amplitude of this log-log parabola contributed by random detector noise. Our numerical examples show that the residual indeed does have roughly the shape of our fitting function, not the shape of the main leading-order contribution from residual foregrounds (the next term in their Taylor expansion; i.e., a cubic term).

4.2. *Baseline Example 2: Near Term Situation (Noise $\gg$ Signal)*

The detector white noise we assumed in the previous example is small in comparison to the signal, in which case we can subtract the foreground easily for each pixel. Nevertheless, that level of noise might not be achieved until future next-generation

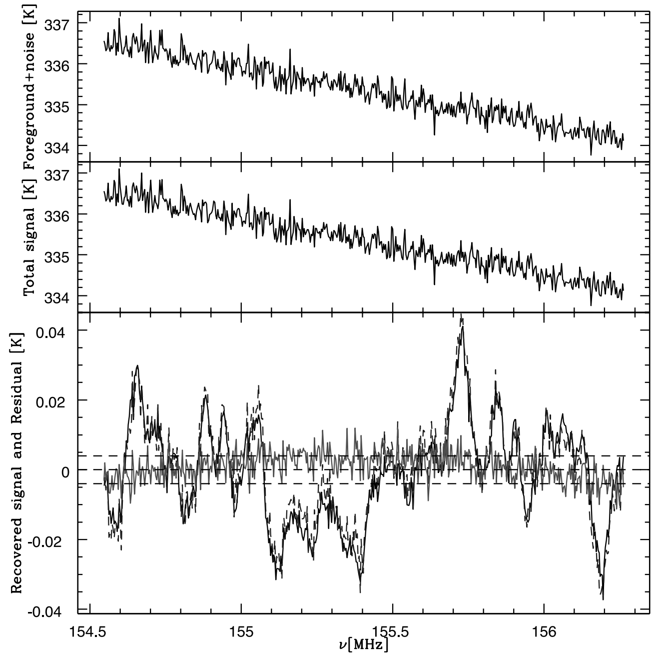

FIG. 5.—Same as Fig. 3, but with 200 times larger detector noise ($\sigma = 200$ mK). In the bottom panel, the recovered 21 cm signal (*dashed curve*) and the residual (*gray curve*) are the results of averaging 4000 pixels. All other curves in the plot are from a single pixel. [*See the electronic edition of the Journal for a color version of this figure*.]

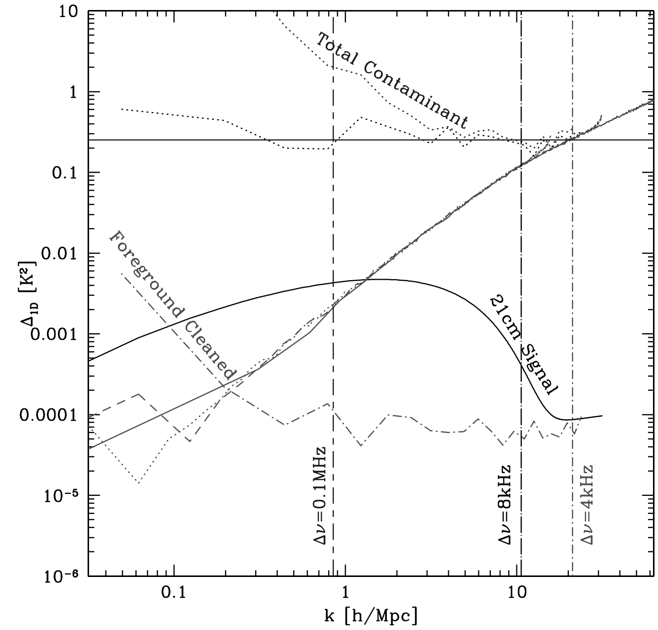

FIG. 6.—Same as Fig. 4, but with 200 times larger detector noise ($\sigma = 200$ mK). The coherent case residual contaminants (*dot-dashed gray curve*) are fast Fourier transformed from the coherently averaged residual shown in Fig. 5 (*gray curve*). The incoherent case residual contaminant (*gray solid, dashed, and dotted curves*) are computed from incoherently averaging 40,000 pixels. The gray solid, dashed, and dotted curves assume 1.7, 8.6, and 17.2 MHz bandwidths, respectively. [*See the electronic edition of the Journal for a color version of this figure*.]

experiments. For upcoming experiments, as we showed in equations (17) and (18), the detector noise is far above the 21 cm signal. In this section, we study the close-term case by assuming a detector noise of $\sigma = 200$ mK, which is 200 times larger than in the previous example, using the same baseline foreground model as in § 4.1.

One would expect that for each individual pixel, the huge level of white noise would destroy much of the information about the 21 cm signal and also obscure the frequency dependence of the foregrounds, making them harder to fit and remove, and that the residual would be noise-dominated. In this scenario, single-pixel cleaning is not enough for our cleaning purposes, and multiple pixels are needed to average down the noise and fit the foregrounds. However, complications arise when processing multiple pixels. Different pixels come from different lines of sight, so their 21 cm signals are either slightly different realizations or completely independent realizations of equation (7), depending on how far apart the pixels are from one another. Furthermore, for pixels that are close to one another, they have slightly different signals in general, but on large scales, the signals within these pixels are more or less identical, while on small scales, the signals are almost independent of signals in other pixels. The details of this complication would probably best be treated via a detailed 3D numerical simulation, where thousands of pixels can be simulated and the signals from them can be tested. Since this is beyond the scope of this paper, we will simply illustrate the basic effects by two extreme situations, which can also be applied to numerically generated signals. We will see from the following examples that our method for foreground cleaning still works reasonably well.

4.2.1. *Coherent Signal Approximation*

For closely separated pixels, we make the crude assumption that the line-of-sight 21 cm signals in these pixels are identical (same phase and amplitude as in eq. [7]). This approximation will simplify our calculation, yet the procedure of doing foreground removal is similar to generally incoherent signals, as we discuss in § 4.2.2.

Since the signals are coherent, the total signal for different pixels is the summation of the same signal and different foregrounds and noise. We use the same method as described in § 4.1 to remove the foreground from the total signal along the line of sight for each individual pixel. We then average all of them in real space and obtain the averaged cleaned signal.

Figure 5 shows the results before and after cleaning. The top and middle panels are plotted for a single pixel with 200 mK noise. The noise wiggles fast on top of the foreground and dominates the signal. It is impossible to tell the difference between situations with and without 21 cm signal (*top and middle panels*). The bottom panel shows the cleaned signal and residual by applying our method and combining 4000 such pixels. Although both foregrounds and noise are at a level that is orders of magnitude higher than the signal, the resulting cleaned signal still captures the main features of the "true" signal and the residual is well controlled.

This confirms that the foregrounds can still be removed effectively when the noise is orders of magnitude higher than the signal. Our fitting method does not introduce additional contamination to the signal, *even when foregrounds with many different spectral shapes are averaged together*.

Figure 6 shows the k-space signal power spectrum compared with foregrounds and noise power spectra for single-pixel and residual power spectra from averaging 4000 pixels. Before cleaning, for each individual pixel, the signal is completely buried under huge noise and foregrounds. After cleaning, the residual (*gray dot-dashed curve*) is orders of magnitude below both the signal and the original contaminants, except on very large scales. Similarly to Figure 4, the plot suggests that we need a frequency resolution of substantially smaller than 0.1 MHz to take advantage of foreground removal sensitivity.

The number of pixels needed to average down the noise varies sharply with the actual noise level. To achieve a similar level of accuracy shown in Figure 5, for a noise level of $\sim$500 mK, we would need approximately 20,000–30,000 pixels. For a noise level of $\sim$100 mK, we would need around 1000 pixels. If the noise is about the same level as the signal, $\sigma = 20$ mK, we only need 50–100 pixels to adequately lower the noise.

When larger numbers of pixels get combined, several side effects appear, such as the variation of spectral indices among different pixels, signal and foreground angular correlation, overestimation of the sensitivity at small scales, etc. We will discuss some of these issues a bit more in § 4.2.2. The best approach, however, is to combine this method with angular approach and remove the contaminants in three dimensions (McQuinn et al. 2006; Morales et al. 2006). Although this is beyond the scope of the present paper, we hope to address this further in future work.

4.2.2. *Incoherent Signals*

For pixels that are far apart, the line-of-sight 21 cm signals in these pixels are no longer the same. They have different phases and amplitudes and therefore are more or less independent of one another. Here we study the case in which all signals are completely independent. Compared to the case studied in § 4.2.1, this case is the opposite extreme. Signals obtained from real observations will probably have a behavior intermediate between these two extremes. By simulating such signals numerically, our cleaning technique is likely to give residual contaminant levels that are intermediate between those we find here for these two extreme cases.

When the signals are incoherent, we still use the same method to remove foregrounds from the total signal for each individual pixel. However, instead of averaging them in real space, we FFT the signal in each pixel to Fourier space to obtain the cleaned power spectrum for each pixel. Then we average all individual power from different pixels and get the final average cleaned power spectrum.

Compared with averaging coherent signals in real space, averaging incoherent signals in Fourier space requires a larger number of pixels to remove the foregrounds effectively. Figure 6 shows the true 21 cm power spectrum compared with residual foregrounds and noise power spectra (*gray solid, dashed, and dotted curves*), defined as the difference between the average cleaned power spectrum and the true 21 cm power spectrum, from incoherently averaging 40,000 pixels. The noise and foreground levels are kept the same as in the previous coherent example. Although we average 10 times more pixels here than for the coherent case, the residual contaminant is at a level higher than in the previous case. However, the residual is still reasonable. On most of the scales, the residual contaminant is less than 10% of the signal. Also note in this case that the 21 cm power spectrum is best measured for scales around $k = 0.1$ h Mpc^{-1}, another consequence from incoherence averaging. In the previous two examples, namely, low-noise and high-noise coherence signal examples, we recovered the signal instead of the power spectrum.

The three different residual curves in the plot are computed assuming 1.7 MHz (*gray solid curve*), 8.6 MHz (*gray dashed curve*), and 17.2 MHz (*gray dotted curve*) bandwidths, respectively. (We used the 1.7 MHz bandwidth for all previous calculations and figures.) As the bandwidth increases, the residual power decreased, especially for smaller values of k. That is to say, a larger bandwidth will help with foreground removal at large scales.

So the bottom line is that, for incoherent signals, our method for foreground cleaning still works, yet its efficiency is reduced due to the fact that the signals are independent. In this case, we could consider increasing bandwidth, combining with angular direction measurements, etc., to improve efficiency and remove foregrounds effectively.

5. DISCUSSION

We have explored how well foreground contamination can be removed from a 21 cm tomography data cube by using frequency dependence alone. We found that with realistic experimental sensitivities, 21 cm tomography is limited mainly by foregrounds for scales of $k \ll 1$ h Mpc^{-1} and limited mainly by noise for $k \gg 1$ h Mpc^{-1}, a result that is rather robust to changing the foreground assumptions. In optimizing the design of upcoming experiments, a useful rule of thumb is therefore to make the channel width substantially smaller than 0.1 MHz, allowing one to take full advantage of the detector sensitivity by pushing residual foregrounds down to the noise level. Fortunately, attaining such a narrow channel width is realistic for upcoming experiments, where the analysis is all done by dedicated high-speed electronics.

We used a simple "blind" removal technique using no prior information about the nature of the foregrounds, merely fitting out a quadratic polynomial in log-log space for the frequency dependence separately for each pixel in the sky. The basic reason that this works so well is that the foregrounds have much smoother frequency spectra than the 21 cm signal.

Although highly effective, this frequency-based cleaning should be viewed as merely one of three complementary foreground countermeasures. First, bright point sources can be identified as strong positive outliers, and the corresponding sky pixels can be discarded, since they constitute only a small fraction of the total survey area. Second, after the frequency-based cleaning step, noise and signal can be further distinguished by their different angular correlations, as described in Santos et al. (2005). This angular approach will be particularly helpful for early 21 cm experiments in which the signal-to-noise ratio is limited. The angular and frequency-based approaches are therefore complementary, and the combination of the two will give the best cleaned 21 cm signal with which to study the "dark" epoch of reionization.

Although our results are quite encouraging for the prospects of doing cosmology with 21 cm tomography, much work remains to be done on the foreground problem, and we conclude by mentioning a few examples.

A key assumption of this paper is that the foregrounds are dominated by emission mechanisms producing fairly smooth spectra. The basis for this assumption is that typical atomic and molecular transitions that can produce spectral lines correspond to much higher frequencies than those relevant to 21 cm tomography. One loophole that needs to be checked quantitatively is the possible contribution of recombination lines from hydrogen cascading down through very large energy quantum numbers $n \sim 10^7$, although early estimates suggest that this is not a significant contaminant (e.g., Oh & Mack 2003).

We performed our analysis on simulated data over a small redshift range, limited by our linearization approximation. It is clearly worthwhile to repeat our analysis with a proper hydrodynamic simulation of the 21 cm signal over the full relevant redshift range. In this case, our three-parameter foreground fit should be generalized to one that assumes that the foregrounds are simple only locally in log-frequency space. An obvious generalization of our method would be to increase the order of the log-log polynomial beyond 2. However, we effectively want to high-pass filter the observed frequency spectrum to clean out foregrounds, and high-order polynomials can in principle spoil this by having sharp localized features. A better generalization of our method to long frequency baselines may therefore be either a cubic spline in log-log space or a Fourier series expansion. Such end-to-end simulations will also be a valuable tool for quantifying how redshift space distortions (whereby the peculiar velocity of the gas breaks the one-to-one correspondence between redshift and frequency) can be exploited to separate the effects of the matter power spectrum from the "gastrophysics" (Barkana & Loeb 2005a). This becomes important especially for channel widths of $\lesssim 0.1$ MHz (Desjacques & Nusser 2004; Iliev et al. 2003).

In summary, the potential scientific return from 21 cm tomography is enormous, both for understanding the reionization epoch and for probing inflation and dark matter with precision measurements of the small-scale power spectrum. Our calculations strengthen the conclusion that foreground contamination will not be a showstopper. The current situation is similar to the quest for the cosmic microwave background in the 1980s in that the cosmological signal has not yet been detected, but is better in the sense that the amplitude of both signals and foregrounds are approximately known, guaranteeing success if the engineering challenges can be overcome.

The authors wish to thank Angelica de Oliveira-Costa, David Rusin, Gary Bernstein, Jose Diego, Lincoln Greenhill, Jacqueline Hewitt, Wayne Hu, Avi Loeb, Marijke Haverkorn, Miguel Morales, and Matias Zaldarriaga for helpful discussions and comments. The authors are especially grateful to Miguel Morales for very informative explanations and discussions of experimental issues. The authors wish to thank the anonymous referee for invaluable comments that improved the paper substantially. X. W. was supported in part by the Kavli Institute for Cosmological Physics through the grant NSF PHY-0114422. X. W. and M. T. were supported by NASA grant NAG5-11099, NSF CAREER grant AST 01-34999, and fellowships from the David and Lucile Packard Foundation and the Research Corporation. L. K. and M. G. S. were supported by NSF grant 0307961 and NASA grant NAG5-11098. M. G. S. was also supported in part by FCT-Portugal under grant BPD/17068/2004/Y6F6.

REFERENCES

Barkana, R., & Loeb, A. 2005a, ApJ, 624, L65
———. 2005b, ApJ, 626, 1
Bennett, C., et al. 2003, ApJS, 148, 97
Bowman, J. D., Morales, M. F., & Hewitt, J. N. 2006, ApJ, 638, 20
Carilli, C. L., Gnedin, N., Furlanetto, S., & Owen, F. 2004, NewA Rev., 48, 1053
Carilli, C. L., Gnedin, N. Y., & Owen, F. 2002, ApJ, 577, 22
Ciardi, B., & Madau, P. 2003, ApJ, 596, 1
Desjacques, V., & Nusser, A. 2004, MNRAS, 351, 1395
Di Matteo, T., Ciardi, B., & Miniati, F. 2004, MNRAS, 355, 1053
Di Matteo, T., Perna, R., Abel, T., & Rees, M. J. 2002, ApJ, 564, 576
Furlanetto, S. R., Sokasian, A., & Hernquist, L. 2004a, MNRAS, 347, 187
Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. 2004b, ApJ, 613, 16
Gnedin, N. Y., & Shaver, P. A. 2004, ApJ, 608, 611
Haiman, Z., & Holder, G. P. 2003, ApJ, 595, 1
Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, A&AS, 47, 1
Haverkorn, M., Katgert, P., & de Bruyn, G. 2003, A&A, 403, 1031
Holder, G., Haiman, Z., Kaplinghat, M., & Knox, L. 2003, ApJ, 595, 13
Hui, L., Stebbins, A., & Burles, S. 1999, ApJ, 511, L5
Iliev, I. T., Scannapieco, E., Martel, H., & Shapiro, P. R. 2003, MNRAS, 341, 81
Iliev, I. T., Shapiro, P. R., Ferrara, A., & Martel, H. 2002, ApJ, 572, L123
Knox, L. 2003, NewA Rev., 47, 883
Kogut, A., et al. 2003, ApJS, 148, 161
Loeb, A., & Zaldarriaga, M. 2004, Phys. Rev. Lett., 92, 211301
Madau, P., Meiksin, A., & Rees, M. J. 1997, ApJ, 475, 429
McQuinn, M., Zahn, O., Zaldarriaga, M., Hernquist, L., & Furlanetto, S. R. 2006, ApJ, in press (astro-ph/0512263)
Morales, M. F. 2005, ApJ, 619, 678
Morales, M. F., Bowman, J. D., & Hewitt, J. 2006, ApJ, 648, 767
Morales, M. F., & Hewitt, J. 2004, ApJ, 615, 7
Oh, S. P., & Mack, K. J. 2003, MNRAS, 346, 871
Peacock, J. A. 1999, Cosmological Physics (Cambridge: Cambridge Univ. Press)
Pen, U.-L., Wu, X.-P., & Peterson, J. 2004, preprint (astro-ph/0404083)
Platania, P., Burigana, C., Maino, D., Caserini, E., Bersanelli, M., Cappellini, B., & Mennella, A. 2003, A&A, 410, 847
Röttgering, H., de Bruyn, A. G., Fender, R. P., Kuijpers, J., van Haarlem, M. P., Johnston-Hollitt, M., & Miley, G. K. 2003, in Texas in Tuscany: XXI Symposium on Relativistic Astrophysics, ed. R. Bandiera, R. Maiolino, & F. Mannucci (Singapore: World Scientific), 69
Santos, M. G., Cooray, A., Haiman, Z., Knox, L., & Ma, C.-P. 2003, ApJ, 598, 756
Santos, M. G., Cooray, A., & Knox, L. 2005, ApJ, 625, 575
Seljak, U., et al. 2005, Phys. Rev. D, 71, 103515
Shaver, P. A., Windhorst, R. A., Madau, P., & de Bruyn, A. G. 1999, A&A, 345, 380
Spergel, D. N., et al. 2003, ApJS, 148, 175
———. 2006, ApJ, submitted (astro-ph/0603449)
Tegmark, M. 1997, ApJ, 480, L87
———. 1998, ApJ, 502, 1
Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. J. 2003, Phys. Rev. D, 68, 123523
Tegmark, M., Eisenstein, D. J., Hu, W., & de Oliveira-Costa, A. 2000, ApJ, 530, 133
Tegmark, M., et al. 2004, Phys. Rev. D, 69, 103501
Tozzi, P., Madau, P., Meiksin, A., & Rees, M. J. 2000, ApJ, 528, 597
Wyithe, J. S. B., & Loeb, A. 2004a, ApJ, 610, 117
———. 2004b, Nature, 432, 194
Zaldarriaga, M., Furlanetto, S. R., & Hernquist, L. 2004, ApJ, 608, 622