



THE IMPORTANCE OF WIDE-FIELD FOREGROUND REMOVAL FOR 21 cm COSMOLOGY: A DEMONSTRATION WITH EARLY MWA EPOCH OF REIONIZATION OBSERVATIONS

J. C. POBER^{1,2,25}, B. J. HAZELTON², A. P. BEARDSLEY^{2,3}, N. A. BARRY², Z. E. MARTINOT², I. S. SULLIVAN², M. F. MORALES², M. E. BELL⁴, G. BERNARDI^{5,6,7}, N. D. R. BHAT^{8,9}, J. D. BOWMAN³, F. BRIGGS¹⁰, R. J. CAPPALLO¹¹, P. CARROLL², B. E. COREY¹¹, A. DE OLIVEIRA-COSTA¹², A. A. DESHPANDE¹³, JOSHUA. S. DILLON^{14,15,16}, D. EMRICH⁸, A. M. EWALL-WICE¹⁶, L. FENG^{12,16}, R. GOEKE¹², L. J. GREENHILL⁷, J. N. HEWITT^{12,16}, L. HINDSON¹⁷, N. HURLEY-WALKER⁸, D. C. JACOBS^{3,25}, M. JOHNSTON-HOLLITT¹⁷, D. L. KAPLAN¹⁸, J. C. KASPER^{7,19}, HAN-SEEK KIM^{9,20}, P. KITTIWISIT³, E. KRATZENBERG¹¹, N. KUDRYAVTSEVA⁸, E. LENC^{4,9}, J. LINE^{9,20}, A. LOEB⁷, C. J. LONSDALE¹¹, M. J. LYNCH⁸, B. MCKINLEY¹⁰, S. R. MCWHIRTER¹¹, D. A. MITCHELL^{9,21}, E. MORGAN¹², A. R. NEBEN¹⁶, D. OBEROI²², A. R. OFFRINGA^{9,10}, S. M. ORD^{8,9}, SOURABH PAUL¹³, B. PINDOR^{9,20}, T. PRABU¹³, P. PROCOPIO²⁰, J. RIDING²⁰, A. E. E. ROGERS¹¹, A. ROSHI²³, SHIV K. SETHI¹³, N. UDAYA SHANKAR¹³, K. S. SRIVANI¹³, R. SUBRAHMANYAN^{9,13}, M. TEGMARK¹⁶, NITHYANANDAN THYAGARAJAN³, S. J. TINGAY^{8,9}, C. M. TROTT^{8,9}, M. WATERTON²⁴, R. B. WAYTH^{8,9}, R. L. WEBSTER^{9,20}, A. R. WHITNEY¹¹, A. WILLIAMS⁸, C. L. WILLIAMS¹⁶, AND J. S. B. WYITHE^{9,20}

¹Department of Physics, Brown University, Providence, RI 02912, USA

²Department of Physics, University of Washington, Seattle, WA 98195, USA

³School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA

⁴Sydney Institute for Astronomy, School of Physics, The University of Sydney, NSW 2006, Australia

⁵Square Kilometre Array South Africa (SKA SA), Pinelands 7405, South Africa

⁶Department of Physics and Electronics, Rhodes University, Grahamstown 6140, South Africa

⁷Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138, USA

⁸International Centre for Radio Astronomy Research, Curtin University, Bentley, WA 6102, Australia

⁹ARC Centre of Excellence for All-sky Astrophysics (CAASTRO), Australia

¹⁰Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia

¹¹MIT Haystack Observatory, Westford, MA 01886, USA

¹²Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

¹³Raman Research Institute, Bangalore 560080, India

¹⁴Department of Astronomy, University of California Berkeley, Berkeley, CA 94720, USA

¹⁵Berkeley Center for Cosmological Physics, University of California Berkeley, Berkeley, CA 94720, USA

¹⁶Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

¹⁷School of Chemical & Physical Sciences, Victoria University of Wellington, Wellington 6140, New Zealand

¹⁸Department of Physics, University of Wisconsin–Milwaukee, Milwaukee, WI 53201, USA

¹⁹Department of Atmospheric, Oceanic and Space Sciences, University of Michigan, Ann Arbor, MI 48109, USA

²⁰School of Physics, The University of Melbourne, Parkville, VIC 3010, Australia

²¹CSIRO Astronomy and Space Science (CASS), P.O. Box 76, Epping, NSW 1710, Australia

²²National Centre for Radio Astrophysics, Tata Institute for Fundamental Research, Pune 411007, India

²³National Radio Astronomy Observatory, Charlottesville and Greenbank, USA

²⁴SKA Organization, Jodrell Bank Observatory, Lower Withington, Macclesfield SK11 9DL, UK

Received 2015 July 24; accepted 2016 January 13; published 2016 February 23

ABSTRACT

In this paper we present observations, simulations, and analysis demonstrating the direct connection between the location of foreground emission on the sky and its location in cosmological power spectra from interferometric redshifted 21 cm experiments. We begin with a heuristic formalism for understanding the mapping of sky coordinates into the cylindrically averaged power spectra measurements used by 21 cm experiments, with a focus on the effects of the instrument beam response and the associated sidelobes. We then demonstrate this mapping by analyzing power spectra with both simulated and observed data from the Murchison Widefield Array. We find that removing a foreground model that includes sources in both the main field of view and the first sidelobes reduces the contamination in high k_{\parallel} modes by several per cent relative to a model that only includes sources in the main field of view, with the completeness of the foreground model setting the principal limitation on the amount of power removed. While small, a percent-level amount of foreground power is in itself more than enough to prevent recovery of any Epoch of Reionization signal from these modes. This result demonstrates that foreground subtraction for redshifted 21 cm experiments is truly a wide-field problem, and algorithms and simulations must extend beyond the instrument’s main field of view to potentially recover the full 21 cm power spectrum.

Key words: cosmology: observations – dark ages, reionization, first stars – techniques: interferometric

1. INTRODUCTION

A major goal of modern experimental cosmology is the detection of 21 cm emission from neutral hydrogen at high redshifts. Depending on the redshifts studied, these

observations can probe a wide range of physical and astrophysical phenomena. Observations at $\sim 100\text{--}200$ MHz ($z \sim 6\text{--}13$ in the 21 cm line) probe the Epoch of Reionization (EoR)—the reionization of the intergalactic medium (IGM) by ultraviolet photons emitted by the first stars and galaxies. Observations at higher frequencies (lower redshifts) trace the neutral hydrogen that remains in galactic halos, and provide a

²⁵ National Science Foundation Astronomy and Astrophysics Postdoctoral Fellow.

low-resolution “intensity map” of large-scale structure and, potentially, the baryon acoustic oscillation (BAO) features in the power spectrum. At lower frequencies (higher redshifts), one begins to trace the birth of the first stars during “Cosmic Dawn” and even the preceding Dark Ages. For reviews of the 21 cm cosmology technique and the associated science drivers, see Furlanetto et al. (2006), Morales & Wyithe (2010), Pritchard & Loeb (2012), and Zaroubi (2013).

A large number of experiments seeking to detect the power spectra of 21 cm fluctuations are already operational or being commissioned, including the LOW Frequency ARray (LOFAR; Yatawatta et al. 2013; Wise et al. 2013)²⁶, 21 CentiMeter Array (21CMA; Zheng et al. 2012)²⁷, the Giant Metrewave Radio Telescope EoR Experiment (GMRT; Paciga et al. 2013)²⁸, the MIT Epoch of Reionization Experiment (MITEoR; Zheng et al. 2014), the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010)²⁹, and the Murchison Widefield Array (MWA; Lonsdale et al. 2009; Bowman et al. 2013; Tingay et al. 2013)³⁰, all of which are targeting the signal from the EoR. A number of additional experiments are also under construction or planned, such as the low-frequency Square Kilometre Array (SKA-low; Mellema et al. 2013)³¹ and the Hydrogen Epoch of Reionization Array (HERA; Pober et al. 2014)³² at EoR and Cosmic Dawn redshifts, and BAOs from Integrated Neutral Gas Observations (BINGO; Battye et al. 2013), TianLai³³, BAORadio (Ansari et al. 2012a, 2012b), the Canadian Hydrogen Intensity Mapping Experiment (CHIME; Shaw et al. 2014)³⁴, and the BAO Broadband and Broad-beam experiment (BAOBAB; Pober et al. 2013b) at lower redshifts.

At all redshifts, however, 21 cm experiments are limited by both the inherent faintness of the cosmological signal and the presence of foregrounds, which can exceed the 21 cm emission by as much as five orders of magnitude in brightness temperature (Santos et al. 2005; Bernardi et al. 2013; Pober et al. 2013a; Yatawatta et al. 2013). As such, the only current detection of H I at cosmological distances comes from cross-correlation studies using maps from the Green Bank Telescope and optical galaxy surveys (Chang et al. 2010; Masui et al. 2013; Switzer et al. 2013). Analysis techniques for recovering the signal focus on the relative spectral smoothness of the foreground emission as an axis for distinguishing these contaminants from the 21 cm emission. Over the past decade, a large body of literature has worked to develop pipelines that can subtract foreground sources from 21 cm data sets (e.g., Morales et al. 2006; Bowman et al. 2009; Liu et al. 2009; Liu & Tegmark 2011; Chapman et al. 2012, 2013; Dillon et al. 2013; Wang et al. 2013). More recently, however, studies of the chromatic interaction of an interferometer with foreground emission have demonstrated that smooth-spectrum foregrounds will occupy an anisotropic wedge-like region of cylindrical $(k_{\perp}, k_{\parallel})$ Fourier space, leaving an “EoR window” above the wedge where the 21 cm signal can be cleanly observed (Datta

et al. 2010; Morales et al. 2012; Parsons et al. 2012b; Trott et al. 2012; Vedantham et al. 2012; Thyagarajan et al. 2013; Liu et al. 2014a, 2014b). These predictions have since been confirmed in data sets from PAPER and the MWA (Pober et al. 2013a; Dillon et al. 2014; Parsons et al. 2014; Ali et al. 2015; Jacobs et al. 2015; Thyagarajan et al. 2015a), although significantly more sensitive observations will be necessary to see if the window remains uncontaminated down to the level of the 21 cm signal. Pober et al. (2014) demonstrate that while current EoR observatories (PAPER, the MWA, and LOFAR) do not possess the sensitivity to detect the 21 cm signal with this pure “foreground avoidance” technique, next-generation experiments such as HERA and the SKA-low can yield high-fidelity power spectrum measurements using this approach, and begin to place constraints on the physics of reionization.³⁵ However, the cosmological signal strength peaks on large scales, so that k modes within the wedge can have significantly more 21 cm power than modes within the window. Pober et al. (2014) show that if foregrounds can be subtracted from 21 cm data sets, allowing the recovery of k modes from within the wedge, then the significance of any power spectrum measurement can be substantially boosted—enabling the current generation of 21 cm experiments to make a detection.

Continued research into foreground subtraction algorithms is therefore clearly well motivated. As yet, no technique—whether subtracting a model of the sky or using a parameterized fit in frequency—has demonstrated that foreground emission in actual observations can be removed to the thermal noise level of current instruments (although the EoR window has, to date, proven relatively free of foregrounds when care is taken to limit leakage from the wedge; Pober et al. 2013b; Parsons et al. 2014; Ali et al. 2015; Jacobs et al. 2015). The purpose of this work is to investigate some of the wide-field effects that complicate the removal of foreground emission using data from the MWA. In particular, we focus on the contribution of sources outside the main lobe of the instrument’s primary beam (in this work, we use the term primary beam to refer to the all-sky power pattern of the antenna or tile element, including sidelobes). Far from the pointing center, chromatic effects in the interferometer response become stronger; sources out in the sidelobes of the primary beam therefore create foreground contamination in higher k_{\parallel} modes than sources near the pointing center. Here, we explore this effect in more detail.

This paper is structured as follows. In Section 2, we lay out a heuristic derivation of how the instrument’s primary beam enters in measurements of the 21 cm power spectrum and how foregrounds are distributed throughout the $(k_{\perp}, k_{\parallel})$ plane. In Section 3, we briefly describe the MWA and the data analyzed in this study. In Section 4, we build on the pedagogical nature of the previous analysis through simulated MWA power spectra using a sky model containing a single point source of emission. By changing the location of this source, we demonstrate these primary-beam effects in a realistic but controlled fashion. In Section 5, we describe the calibration, preprocessing, and foreground subtraction applied to the observed data before making a power spectrum. The main result is presented in Section 6, where we compare two distinct power spectra made from our data: one where we have

²⁶ <http://www.lofar.org>

²⁷ <http://21cma.bao.ac.cn/index.html>

²⁸ <http://www.ncra.tifr.res.in/ncra/gmrt>

²⁹ <http://eor.berkeley.edu>

³⁰ <http://www.mwatelescope.org>

³¹ <http://www.skatelescope.org>

³² <http://reionization.org>

³³ <http://tianlai.bao.ac.cn>

³⁴ <http://chime.phas.ubc.ca>

³⁵ Although Pober et al. (2014) focused on results from EoR-frequency experiments, the breakdown into “wedge” and “EoR window” is generic for all 21 cm studies (Pober et al. 2013b).

subtracted a foreground model that includes sources in the beam sidelobes, and one where only sources in the main field of view are removed. We discuss the implications of these results for future foreground subtraction efforts in Section 7 and conclude in Section 8.

2. WIDE-FIELD EFFECTS IN THE EOR POWER SPECTRUM

Although many 21 cm experiments have wide fields of view, only recently have studies focused on how wide-field effects might complicate measurements of the 21 cm power spectrum. Theoretical work has identified the foreground wedge described above and provided a formalism for mapping the position of foreground emission on the sky to k modes of the 21 cm power spectrum (Morales et al. 2012; Parsons et al. 2012b; Trott et al. 2012; Vedantham et al. 2012; Thyagarajan et al. 2013; Liu et al. 2014a, 2014b). Broadly speaking, there are two flavors of 21 cm power spectrum analysis: a “delay spectrum” approach, where the line-of-sight Fourier transform is done on individual visibilities—and is therefore not strictly orthogonal to the transverse directions because of the frequency dependence of an individual visibility—and an “imaging” approach, where the line-of-sight Fourier transform spans multiple visibilities and is truly orthogonal to the transverse directions on the sky.³⁶ A full discussion of the differences between these two approaches is outside the scope of this work, but previous analyses have shown that the wedge and the mapping from foreground sky position to k_{\parallel} modes of the power spectrum remain valid for *both* frameworks. In the delay spectrum approach, the chromatic dependence of an individual baseline is completely preserved, so that all foreground emission at a given location maps to a given k_{\parallel} mode. An imaging approach, however, removes the mapping between delay and sky position by projecting out the frequency sine wave for a known geometric delay. In an imaging power spectrum, frequency structure is dominated by the intrinsic spectra of the sources, so that a significant amount of foreground emission maps to low k_{\parallel} modes, reflective of their inherent (smooth) frequency spectrum. However, the chromatic response of the interferometer still affects the observed emission, leading to a wedge feature analogous to that of the delay spectrum approach, but with more of the emission concentrated at low k_{\parallel} (Morales et al. 2012; Dillon et al. 2015). Because of the brightness of foreground emission, this wedge still dominates any 21 cm signal in the modes it occupies.

Explorations of these wide-field effects in actual data have been more limited. Thyagarajan et al. (2015a, 2015b) studied both simulated and actual MWA observations using the delay spectrum technique and found an excellent match between the two, demonstrating a good understanding of both foreground emission and the primary beam of the MWA. They also found that the foreshortening of baseline lengths when projected toward the horizon creates sensitivity to diffuse emission normally resolved out on longer baselines. Diffuse foregrounds are bright enough that they can be detected despite the small

(but non-zero) response of the MWA element toward the horizon. This led to what they dubbed the “pitchfork” effect, a foreground signature in delay space where bright emission from within the main field of view appeared at low delays, and emission from the horizon at high delays.

This work studies similar effects using an imaging power spectrum approach and will confirm that the sky-position to k_{\parallel} mapping still holds. We will also focus on the ability to *subtract* foreground emission away from the main field of view to lower the contamination in high k_{\parallel} modes. In this section, however, we use the delay-spectrum formalism (Parsons & Backer 2009; Parsons et al. 2012b) to provide a general framework for understanding these effects. We stress that the delay spectrum provides a straightforward, pedagogical way to interpret power spectrum results, since the wide-field chromatic effects appear at first order. As argued in Morales et al. (2012), Trott et al. (2012), and Liu et al. (2014a, 2014b), and as will be confirmed with data below, these wide-field effects are generic to all interferometric 21 cm experiments.

The basic premise of the delay-spectrum technique presented in Parsons et al. (2012b) is that the square of the frequency Fourier transform of a single baseline’s visibility spectrum (i.e., the delay spectrum, $\tilde{V}_b(\tau)$) approximates a measurement of the cosmological power spectrum (to within a proportionality factor):

$$|\tilde{V}_b(\tau)|^2 \propto P(k_{\perp}, k_{\parallel}), \quad (1)$$

where

$$\tilde{V}_b(\tau) = \int d\nu V_b(\nu) e^{2\pi i\nu\tau} \quad (2)$$

is the delay spectrum, τ is delay, ν is frequency, V is a visibility, and the subscript b indicates that the visibilities are from a single baseline.

Intuitively, this relation is well motivated. To a good approximation, a single baseline b probes a single transverse scale, and thus a single k_{\perp} mode. And, since cosmological redshifting of the 21 cm line maps observed frequencies into line-of-sight distances, the Fourier transform of the frequency spectrum approximates a range of k_{\parallel} modes. Put more succinctly, for an interferometer, baseline length b maps to cosmological k_{\perp} and delay τ maps to k_{\parallel} .

The power of this simple formalism is that we can now map the effects of the primary beam, which enter into a visibility measurement in a well-known way, to cosmological Fourier space and the power spectrum $P(k_{\perp}, k_{\parallel})$. We begin with the form of a visibility in the flat-sky approximation (Thompson et al. 2001)³⁷:

$$V_b(\nu) = \int dl dm A(l, m, \nu) I(l, m, \nu) e^{-2\pi i(ul+vm)}, \quad (3)$$

where A is the primary beam, I is the sky brightness distribution, l and m are direction cosines on the sky, ν is frequency, and u and v are the projected baseline lengths on the ground plane measured in wavelengths. We can rewrite this expression in terms of the geometric delay τ_g (Parsons &

³⁶ The terminology of an “imaging” power spectrum is potentially misleading, but it has become somewhat standard in the community. The key feature is not that an image of the sky is made, but rather that visibilities are gridded into the uv plane and the frequency Fourier transform is taken in a direction truly orthogonal to u and v . The nomenclature of an “imaging” power spectrum arises because the gridded uv data are only a 2D spatial Fourier transform away from an image.

³⁷ Although use of the flat-sky approximation to derive a wide-field interferometric effect may seem ill-motivated, it greatly simplifies the math in this pedagogical treatment. See Parsons et al. (2012a, 2012b) and Thyagarajan et al. (2015a, 2015b) for a discussion of the subtleties introduced by the curved sky in the delay formalism.

Backer 2009):

$$V_b(\nu) = \int dl dm A(l, m, \nu) I(l, m, \nu) e^{-2\pi i \nu \tau_g}, \quad (4)$$

where

$$\tau_g = \frac{\mathbf{b} \cdot \hat{\mathbf{s}}}{c} = \frac{1}{c}(b_x l + b_y m), \quad (5)$$

$\mathbf{b} \equiv (b_x, b_y)$ is the baseline vector measured in meters (i.e., $\mathbf{u} \equiv (u, v) = \nu \mathbf{b}/c$), and $\hat{\mathbf{s}} \equiv (l, m)$. Performing the delay transform given by Equation (2) gives us a delay spectrum:

$$\tilde{V}_b(\tau) = \int dl dm d\nu A(l, m, \nu) I(l, m, \nu) e^{-2\pi i \nu (\tau_g - \tau)}. \quad (6)$$

If we make the pedagogical assumption that both A and I are independent of frequency, we can straightforwardly perform the delay transform integral³⁸:

$$\tilde{V}_b(\tau) = \int dl dm A(l, m) I(l, m) \delta(\tau_g - \tau). \quad (7)$$

Since Equation (5) relates the geometric delay τ_g to a specific set of sky direction cosines (l, m) , the delta function selects a subset of sky positions that contribute to each τ mode in the delay spectrum, albeit with a baseline-dependent non-trivial mapping between sky position and τ . It is always true, however, that sources that appear at high delays are those that are far from the pointing center of the instrument (hence the name ‘‘horizon limit’’ given to the maximum delay a source can appear at in Parsons et al. 2012b). Following Equation (1), we can say

$$P(k_{\perp}, k_{\parallel}) \propto \left[\int dl dm A(l, m) I(l, m) \delta(\tau_g - \tau) \right]^2, \quad (8)$$

where the length of baseline b sets k_{\perp} and $\tau \propto k_{\parallel}$. This analysis therefore implies that sources at large delays (i.e., sources near the edges of the field of view, by Equation (5)) contaminate the highest k_{\parallel} modes of the wedge $(k_{\perp}, k_{\parallel})$ space. Although not always stated as directly, this result was also found in Morales et al. (2012), Vedantham et al. (2012), Thyagarajan et al. (2013), and Liu et al. (2014a) using entirely independent formalisms.

Equation (8) also shows the main result we wished to derive in this section: the (smooth-spectrum) sky emission $I(l, m)$ that appears in each delay mode is multiplicatively attenuated by the primary beam of the instrument. Therefore, the foreground emission that contaminates those k_{\parallel} modes measured by a single baseline will itself be attenuated by a (distorted) slice through the square of the primary beam of the instrument. This result is illustrated schematically in Figure 1. Note that, because delay space is a one-dimensional projection of the sky coordinates (see Equation (5)), the attenuating beam in a k_{\parallel} spectrum will vary depending on the orientation of the baseline. On an east/west baseline, for example, the delay axis probes the relative east/west position of the source and is insensitive to north/south translations in source positions. Such a baseline will therefore clearly show the effects of the eastern and western sidelobes of the primary beam in its k_{\parallel} spectrum.

³⁸ Parsons et al. (2012b) showed that the frequency dependence of both A and I creates a convolving kernel, broadening the footprint of each delay mode. The ramifications of this effect are discussed below, but they only complicate the pedagogical nature of the current analysis.




Figure 1. Schematic diagram of the effects discussed here. The primary beam attenuates foregrounds in the k_{\parallel} direction.

Similar logic applies to a north/south baseline and the northern and southern sidelobes of the primary beam. Delays on a northeast/southwest baseline, however, probe northeast/southwest sky position, and thus the east/west translation of a source through the eastern and western sidelobes does not cause as rapid a change in k_{\parallel} . The net effect is that when all baselines of the same magnitude are averaged into a k_{\perp} bin, these different k_{\parallel} sidelobe patterns add up and smear out the location of the sidelobes.

An important but subtle point is that the above derivation for mapping sky coordinates into k space was strictly for flat-spectrum emission. As shown in Parsons et al. (2012b), any spectral structure—whether intrinsic to the source or the instrumental response—introduces a convolving kernel that broadens the footprint of each k_{\parallel} mode in cosmological Fourier space. While this kernel is narrow for smooth-spectrum foregrounds, spectral structure in the 21 cm signal spreads the 21 cm power across a wide range of k_{\parallel} modes. This is equivalent to saying that the 21 cm signal intrinsically has power on these cosmological scales. The situation for foreground emission is different, however. Although power spectrum plots are labeled with axes of $(k_{\perp}, k_{\parallel})$ with units of $h\text{Mpc}^{-1}$, these cosmological scalings apply *only* to the 21 cm signal. The analysis above shows how foregrounds map into this space, and how the primary beam affects this mapping. The primary beam of the instrument does still act as a window function and can affect high k_{\parallel} modes of the cosmological signal; however, the cosmological signal has been shown to be relatively featureless on the scale of this kernel (see Parsons et al. 2012b), rendering this effect very small. Nevertheless, the 21 cm signal is an all-sky signal with real intrinsic k_{\parallel} structure. There is therefore always a 21 cm signal at the peak beam response, so there will always be power at all k_{\parallel} modes truly intrinsic to the cosmological signal. This point will be discussed further in Section 7, where we consider the possibility of detecting 21 cm emission at k_{\parallel} modes where the foregrounds fall in the nulls of the primary beam.

The very wide and relatively smooth primary beam of the PAPER instrument makes the predicted foreground attenuation difficult to see in the analysis of Pober et al. (2013a). However, for instruments such as the MWA and LOFAR, which use tiles




Figure 2. Configuration of the MWA-128 array; each square represents one tile of 16 dipoles.

of dipoles to increase the system gain and narrow the size of the primary beam, there should be two clear effects visible in the power spectra. First, there should be significant attenuation of the wedge foreground emission before the horizon limit, since the instrument’s field of view is significantly smaller than 2π steradians, as is seen in Dillon et al. (2014). Second, at higher k_{\parallel} values than those corresponding to the main beam of the instrument, foreground emission should appear coming from the sidelobes of the primary beam. These two effects can be seen in the delay-space simulations of different antenna elements presented in Thyagarajan et al. (2015b). For an imaging power spectrum technique that averages baselines together, the second effect will be less clear for an instrument such as the MWA, in which all the dipoles and tiles are oriented in the same direction. In this case, the sidelobes are always oriented north/south and east/west; as explained above, however, the beam footprint in k_{\parallel} will differ from baseline to baseline, depending on each baseline’s orientation relative to the sidelobe pattern. This will have the effect of smearing out the sidelobe across a wider range of k_{\parallel} modes than would be seen in an instrument with circularly symmetric sidelobes, but as we will show, the feature is still quite visible in the power spectrum.

The structure of the remainder of this paper is as follows. First, in Section 3, we describe the MWA instrument and observations in more detail. With this context provided, we provide the results of two principal analyses. In Section 4, we present simulated MWA power spectra made from a sky consisting of a single point source. By moving the position of this source from simulation to simulation, we can see the primary-beam effects described above in a controlled fashion. In Section 5, we use observations from the MWA to analyze these primary-beam features and present the power spectra of these data in Section 6. In particular, we focus on the effect of subtracting sources from sidelobes outside the primary field of view.




Figure 3. Three MWA tiles, each consisting of 16 dual-polarization dipole elements in a 4×4 grid.

3. OBSERVATIONS WITH THE MWA

The MWA in Australia consists of 128 tile elements, and each tile is composed of 16 dual-polarization dipole antennas; the array configuration is shown in Figure 2. The tile element has the effect of significantly narrowing the MWA’s field of view over that from a single dipole, but also introduces significant regular sidelobes in the primary beam. Figure 3 shows three MWA tiles; every tile is aligned north/south, so the sidelobes from each tile appear with nearly the same orientation.

The data used in this work were taken with the MWA on 2013 August 23 (Julian Date 2456528) over the course of approximately three hours from 16:47:27 to 19:58:24 UTC. The observations were taken over a frequency band centered on 182.415 MHz, with a total bandwidth of 30.72 MHz divided into 24 1.28 MHz coarse channels, which are each further divided into 768 40 kHz fine channels.

The data used in this analysis span a total of six pointings each 30 minutes long, where an analog beamformer steers the main lobe of the primary beam to nearly the same sky coordinates for each pointing. The sky is then allowed to drift overhead for 30 minutes before repointing. The data within each pointing are saved as individual “snapshot” observations, each lasting 112 s, with individual integrations of 0.5 s. Figure 4 shows the tile primary beam at three different beamformer pointings: the beginning of the observation, a zenith-phased pointing, and the end of the observation. Since each pointing changes the overall primary-beam response of the instrument, the sidelobe patterns in the final integrated power spectrum will be smeared. As will be shown below, however, the effects of the sidelobes are still quite visible despite the changing shape of the primary beam.

4. PEDAGOGICAL SIMULATIONS

Before presenting the full analysis of this data set, we will first investigate the effects of the location of celestial emission on the cosmological power spectrum and the wedge in particular. In this section, we will simulate visibilities for a single point source and calculate the dependence of the power spectrum on the source’s location. Visibilities are simulated using the Fast Holographic Deconvolution (FHD) software package.³⁹ Visibility simulation is one of several functions in FHD; as described below, FHD also performs calibration and source subtraction on our actual data. As a simulator, FHD

³⁹ Source code publicly available at <https://github.com/miguelfmorales/FHD>.

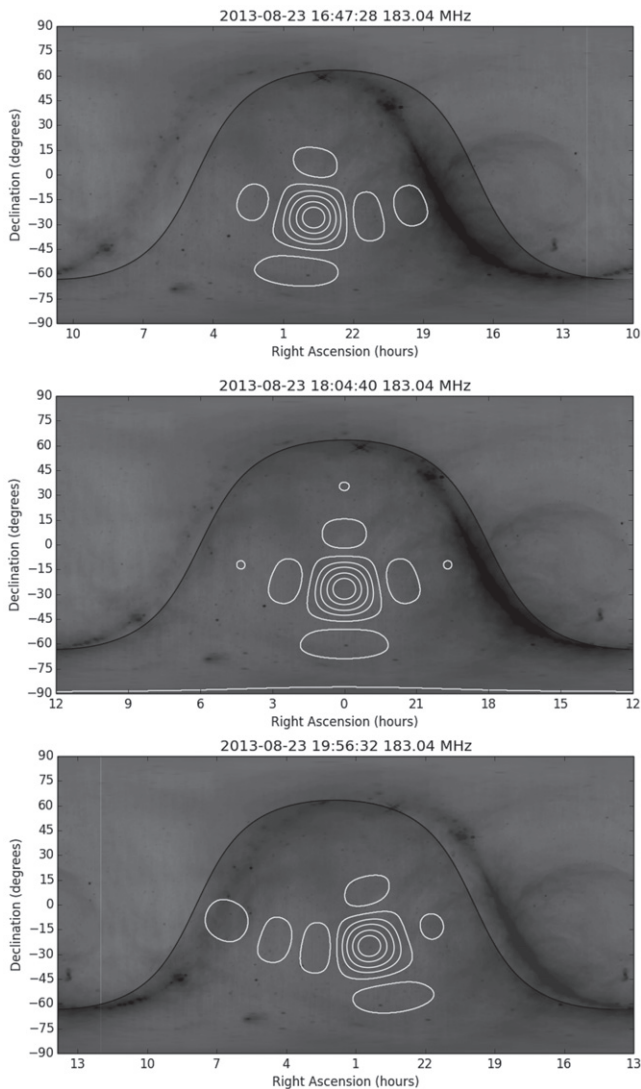


Figure 4. Primary-beam responses of the MWA tiles at several pointings. White contours show the beam response overlotted on the all-sky map of Haslam et al. (1982); contour levels are 0.01, 0.1, 0.25, 0.5, and 0.75 of peak beam response. Although the sidelobes move over the course of the observation, the main field of view remains relatively constant. Top: the first (earliest) pointing in the 3 hr data analyzed here. Center: the zenith-phased pointing near the center of the three hours. Bottom: the last (latest) pointing of the data set.

constructs a model of the sky in uv space and integrates small regions of the uv plane using the holographic beam kernel (Morales & Matejek 2009) to create model visibilities.

For this analysis, we simulate visibilities for all the baselines in the MWA in 768 fine frequency channels spanning the observed 30.72 MHz frequency band. We simulate only one 112 s snapshot when the primary beam is pointed at zenith (i.e., the snapshot shown in the middle panel of Figure 4). In addition to reducing the computational demand of the simulations, using only one snapshot allows us to see the sidelobe patterns most clearly, since integrating over a longer time means including data when the array had a different pointing and primary beam.

We conduct four simulations, each consisting of one radio point source at a different location on the sky; the locations simulated are shown in Figure 5. For each simulation, the inherent flux density of the source is increased relative to

source D (located at zenith) by the inverse of the primary beam response at its location. In other words, each source simulated has the same *apparent* flux density. This choice places all the final power spectra on the same scale, allowing for easier comparison.

In Figure 6, we show the 2D $(k_{\perp}, k_{\parallel})$ power spectra for each of the four simulations described above. We show the power spectrum from only one of the two linearly polarized dipoles of the MWA; the power spectrum for the other polarization is quite similar. Letters correspond to the source labels in Figure 5. To make the power spectrum, the simulated visibilities are imaged by FHD and then analyzed by the ϵ pipeline described in B. J. Hazelton et al. (2015, in preparation).⁴⁰ For more information on the data products transferred between FHD and ϵ , see D. C. Jacobs et al. (2015, in preparation).

The effect of source position on the power spectrum is clear and agrees with the intuition developed in Section 2. Source D is located directly at zenith, with the subsequent sources offset to higher declination (with right ascension held fixed). In Figure 6, source D exhibits no wedge feature. (The power at high k_{\perp} values is due to poor uv coverage on these scales and is described in more detail below.) Sources C and B show a clear wedge feature arise as the source is moved away from zenith, and the power spectrum of source A—where the source is located in the sidelobe of the primary beam—shows a concentration of power outside the main field of view (indicated by the dashed black line) but inside the horizon limit (solid black line). This feature is in exact accord with our predictions. Simulations using sources offset in right ascension (instead of declination) show the same effect, as do sources with offsets in both right ascension and declination: power moves to higher k_{\parallel} as the source moves further from the field center.

5. DATA ANALYSIS

In this section we present the full analysis of the three hours of MWA data described in Section 3. The data are processed through the same imaging and power spectrum analyses (done by FHD and ϵ , respectively) applied to the simulations. However, there are initial preprocessing, calibration, and foreground subtraction steps applied to the data, which we describe here.

5.1. Preprocessing

Preprocessing of the data uses the custom-built Cotter pipeline, which performs time averaging of the integrations to 2 s and frequency averaging of the narrow-band channels to 80 kHz (Offringa et al. 2015). Cotter also uses the AOFLAGGER code to flag and remove radio-frequency interference (Offringa et al. 2010, 2012). Cotter also performs a bandpass correction, removing the spectral shape within each coarse channel as well as correcting for variations in digital gain between the coarse channels. Finally, the data are converted from an MWA-specific data format to `uvfits` files.

⁴⁰ Source code publicly available at <https://github.com/miguelfmorales/eppsilon>.




Figure 5. Positions of the four sources simulated. Source locations are in red; black contours show the 1% primary-beam levels. Note that there are four independent simulations, each consisting of one point source only. Letters correspond to the power spectra in Figure 6.

5.2. Calibration and Imaging

After the preprocessing, data are further calibrated and imaged using the FHD software package. FHD was designed for interferometers with wide fields of view and direction-dependent gains such as the MWA, and it uses the holographic beam pattern to grid visibilities to the uv plane. FHD also keeps track of the gridding statistics in the uv plane to allow for full propagation of errors through the image and into the power spectrum. In this analysis, we do not use FHD to perform a deconvolution and construct a source model from the data themselves, as was described in Sullivan et al. (2012); rather, we input a catalog of point sources and use FHD to calculate model visibilities. In all calculations, FHD uses a simulated primary-beam model including the effects of mutual coupling between dipoles in a tile (Sutinjo et al. 2015).

FHD also applies a calibration to the data, using the source model provided to solve for frequency-dependent complex gain parameters per tile and per polarization. Using an iterative approach, we reduce the number of free parameters by averaging the calibration solutions into a bandpass model that is updated on a per-pointing (i.e., 30 minute) basis. Depending on the position of a tile in the array, one of five cables of different lengths is used to return the signal for central processing; we find it necessary to calculate a different bandpass model for each type of cable in the system. We also fit and remove a per-antenna polynomial (quadratic in amplitude, linear in phase) that varies on a per-snapshot (112 s) timescale, as well as a fit for a known ripple caused by a reflection within a 150 m cable. This particular cable is not present in all tiles, so the ripple is only removed from those that contain this cable; reflections from cables of other lengths on




Figure 6. $(k_{\perp}, k_{\parallel})$ power spectra of the simulated point sources. Letters correspond to source positions in Figure 5. The solid black line shows the horizon limit; the dashed black line indicates the main field of view. The wedge feature is absent for source D, located exactly at zenith, and power moves higher in k_{\parallel} as the source moves further from the center of the field of view. Note that the schematic Figure 1 is plotted with linear axes, whereas this figure uses logarithmic axes, which cause the horizon and field-of-view lines to be parallel.

other tiles appear to have much smaller amplitude, although work is in progress to remove these effects as well.

For the present analysis, we image each snapshot at each frequency channel and make 3D image cubes in HEALPix (Górski et al. 2005). Each snapshot cube is then summed in image space to make a final integrated cube for power spectrum analysis.

5.3. Foreground Subtraction

It is through FHD that model visibilities are also subtracted from the data. We use two sets of model visibilities generated from a custom-made point source catalog. In the main field of view, the catalog contains sources generated from FHD deconvolution outputs and an advanced machine-learning source identifier designed to reject spurious sources (P. A. Carroll et al. 2015, in preparation). Outside the main field of view, our catalog combines sources from the MWA Commissioning Survey (MWACS; Hurley-Walker et al. 2014), the Culgoora catalog (Slee 1995), and the Molongolo Reference Catalog (Large et al. 1981). In one model we include only ~ 4600 sources that fall within the primary lobe of the MWA beam; in the other, we include all sources up to and including the first sidelobe (~ 8500 sources). An image of all the sources included during the zenith-phased snapshot is shown in Figure 7. There are two effects that serve to limit the number of sources included in our model. First, we use a primary-beam threshold cut: any sources that fall where the beam response is less than 1% of the peak response are not included in the model. Second, because it is a composite of several surveys, the completeness of our catalog is not uniform over the sky. In particular, MWACS does not cover the full declination range of the observations here; the effect is that fewer sources are removed from the lower declinations of the southern sidelobe, and very few are in the northern sidelobe. This has the effect of introducing a small time dependence in the number of sources included in our model, since the declination coverage of the primary beam does change with pointing (see Figure 4). MWACS also avoids the Galactic plane, which reduces the number of sources in the model at the early and late pointings to ~ 7000 .

It is also important to note that our sky model assumes a fixed spectral index of -0.8 for each source. Although the actual sources on the sky will have some spectral structure, the fact that we include minimal frequency dependence in the model serves to strengthen the arguments below: subtracting a nearly achromatic foreground model removes power from chromatic (i.e., high k_{\parallel}) modes of the power spectrum. This is a clear demonstration of the inherent chromaticity of the interferometer response pattern.

6. POWER SPECTRA

We now present the power spectra of these data generated by the ϵ code. With observational data, ϵ empirically calculates the noise level in the visibilities and fully propagates errors in the visibilities through to the 3D power spectrum. The important results here are the cylindrically averaged 2D power spectra, shown in Figure 8. In each row in this figure, the left-hand panel shows the power spectrum with only sources in the primary lobe removed, while the center panel shows the power spectrum where sources are also subtracted from the sidelobes. In order to enhance the subtle difference between the two

panels, we subtract the power spectrum where sidelobe sources were removed from the power spectrum where only main-lobe sources were removed (i.e., we subtract the center panel from the left-hand panel). Note that we perform the subtraction in full 3D (k_x, k_y, k_z) space before binning into 2D (k_{\perp}, k_{\parallel}) space. We plot the result of this subtraction in the right-hand panel in each row of Figure 8. Most of the difference fluctuates randomly between positive (blue) and negative (red) values, showing no systematic change of the power spectrum in these regions. However, the consistently blue region shows that subtracting sources from the sidelobes removes a non-trivial amount of power (as much as 10% compared to the power spectrum with no subtraction of sidelobe sources, although typical values are $\sim 1\%$) from the region where the sidelobe is expected—outside the main lobe (dashed black line) but within the horizon (solid line). Since the size of the main lobe is dependent on frequency and pointing, the dashed black line is only an approximate marker; the power that is removed from k_{\parallel} modes below this line is consistent with being sidelobe power from a range of frequencies and pointing centers.

Although not the primary goal of this paper, there are a few additional features in the power spectra that warrant explanation.

1. The horizontal lines running across the EoR window are the effect of the coarse channelization used by the MWA. Between any two 1.28 MHz coarse channels are two 80 kHz channels that are flagged due to low signal response and potential aliasing concerns. This flagging in frequency has the effect of introducing covariance into the line-of-sight k_{\parallel} modes, which are effectively produced by a Fourier transform of the frequency axis. This additional covariance has the effect of coupling power from the wedge into higher k_{\parallel} modes. Because the flagging is at regular intervals, this additional power also appears at regular intervals in k_{\parallel} (the appearance of irregular spacing comes from the logarithmic scale on the y axis). Work is underway on algorithms that can reduce this covariance using priors on the fact that the power comes from k_{\parallel} modes within the wedge.
2. The vertical lines, which are especially prevalent at high k_{\perp} modes, come from the uv coverage of the MWA. The MWA has exceptionally dense coverage at low k_{\perp} due to its large number of short baselines. However, at higher k_{\perp} ($\gtrsim 10^{-1} h\text{Mpc}^{-1}$) there are gaps in the coverage, which results in particularly noisy measurements of certain modes. Therefore, while these modes appear to have very high power, they also have very large associated error bars. A plot of the errors calculated by ϵ for the XX polarization is shown in Figure 9.
3. There are blue/purple regions outside the wedge that are negative. This is because ϵ cross-multiplies the even time samples in the data set with the odd time samples (with the samples interleaved on a timescale of 2 s); this has the effect of removing the positive-definite noise bias that would result from squaring the entire data set. Alternating positive and negative values correspond to noise-dominated regions.
4. Most obviously, a large amount of foreground power remains in the power spectra. This is not surprising, because our analysis subtracted only a few thousand point sources, ignoring diffuse emission from both the Galaxy and unresolved point sources. Subtracting models of this




Figure 7. Sources used for calibration and subtraction. This image shows the source positions during the zenith-phased pointing. Any sources where the beam response is greater than 1% of the peak value during the zenith pointing are included in our model. The sidelobes are clearly distinguishable from the main beam. The declination range of the MWACS survey is -10° to -55° , which accounts for the drop in source density outside this interval.

emission will clearly be necessary for any possibility of recovering the 21 cm signal from inside the wedge. The effects concerning sidelobes presented in this work, however, are still quite important: the additional fraction of emission removed when including sidelobe sources is more than enough to swamp the EoR signal, which might have a peak power spectrum brightness of the order of $10^6 \text{ mK}^2 h^{-3} \text{ Mpc}^3$.

7. DISCUSSION

Through the advances in our understanding of EoR foregrounds (i.e., the “wedge” and “EoR window” paradigm),

we now have a model for the detailed impact of sources far from pointing center on the recovery of the 21 cm power spectrum. This work demonstrates that sources outside the main field of view are a significant contaminant of the modes of interest in the 21 cm power spectrum, even for an “imaging” power spectrum analysis. It is therefore worthwhile to heuristically consider the detailed pattern that sources far from pointing center leave in cylindrically averaged (k_\perp, k_\parallel) space. While not all of the conclusions below directly follow from the empirical power spectra analyzed here, the formalism presented in Section 2, the delay spectrum analyses in Thyagarajan et al. (2015a, 2015b), combined with the results of our sidelobe source subtraction from MWA observations suggest several interesting lines of reasoning. To guide this discussion, we




Figure 8. $(k_{\perp}, k_{\parallel})$ power spectra of the data. XX linear polarization is on the top row, YY on the bottom. The solid black line shows the horizon limit; the dashed black line indicates the main field of view. Left: power spectra where only sources in the main lobe of the beam are used for calibration and then subtracted from the data. Center: power spectra where sources in both the main lobe and the sidelobes are used for calibration and then subtracted. Right: the difference between the left and center plots. (Note that the data are differenced in 3D (k_x, k_y, k_z) space and then averaged in k_{\perp} annuli.) Although the left and center panels appear indistinguishable, subtracting one from the other reveals a significant difference outside the first null of the primary beam. The consistently blue region shows that removing sources in the sidelobes has removed power at high k_{\parallel} outside the main field of view.

divide the sky into three rough categories—the primary field of view, the sidelobes, and the nulls—and discuss the effects of sources appearing in each regime.

1. *Primary Field of View:* These are the sources that are traditionally considered when treating foreground removal from 21 cm experiments. Although the primary field of view may not be at zenith for phased array telescopes (e.g., MWA and LOFAR), phase rotation still places these sources at low delays, and therefore low k_{\parallel} in power spectrum measurements. The dashed line in Figure 8 roughly indicates the edge of the main lobe of the MWA, and the brightness of emission can be seen clearly to fall as one moves to higher k_{\parallel} modes. Since this emission is located where the beam response is at a maximum, it appears as the brightest contaminant in the

21 cm power spectrum. Because they are detected at high signal-to-noise ratio, point sources in the primary field of view are often used to simultaneously calibrate the instrument response while they are subtracted from a visibility model (e.g., Yatawatta et al. 2013). Diffuse emission and unresolved point sources generally dominate the total foreground power, requiring additional models or parametric methods for removal (e.g., Chapman et al. 2012).

2. *Sidelobes:* As seen in the present work, emission in the sidelobes appears at higher k_{\parallel} values than emission from inside the primary field of view. Therefore, a model of the sidelobes and the emission that falls within them must be subtracted from the data in order to recover these k_{\parallel} modes closer to the EoR window. The primary beam




Figure 9. Errors on the XX power spectrum shown in the upper left panel of Figure 8 calculated by ϵ . uv coverage is worse at high k_{\perp} , leading to higher errors. These errors downweight the vertical streaks seen at high k_{\perp} in Figure 8 when estimating a 1D power spectrum.

attenuation has the effect of reducing the fractional accuracy required in modeling these sources, since only their residual apparent flux density contaminates the power spectrum. Since emission in sidelobes contaminates higher k_{\parallel} modes than emission in the main beam, however, removing wide-field emission may be more effective at reducing leakage into the EoR window than removing emission from the primary field of view.

3. *Nulls*: Lastly, one might assume that the nulls in the primary beam might serve as good “EoR windows”, just like the regions of k space outside the horizon. (Recall the discussion in Section 2: while nulls on the sky clearly attenuate all emission from that sky position, when we refer to beam nulls in k_{\parallel} , however, these nulls are confined to the foreground emission. The EoR signal from other positions on the sky still produces unattenuated power at these k_{\parallel} modes.) And while the present analysis does indeed suggest that areas of k space corresponding to sky positions of exceedingly low beam response will be free from foreground contamination, some caveats must be issued. First, the nulls between sidelobes are likely not as deep as analytic models suggest, due to effects such as mutual coupling between the dipoles in a station and group delay errors (Neben et al. 2015). Second, as derived in Parsons et al. (2012b), there is a non-negligible k -space point-spread function convolving each source of emission. Therefore, while there may be narrow nulls between sidelobes, emission within the sidelobes can contaminate these nulls due to this spillover effect. Finally, it is worth remembering that the mapping from zenith angle to delay (which, recall, maps to k_{\parallel}) is nonlinear. A delay bin near the horizon corresponds to much lower elevation than a delay bin near zenith. This means that while emission in bins far from the main lobe of the beam is strongly attenuated by the beam response, the total aggregate sum of emission in that k_{\parallel} bin can still be large, since it corresponds to a large area of sky (Thyagarajan et al. 2015b). It may still be that even with instruments such as SKA and HERA, which have narrower fields of view and less sensitivity to foreground emission away from pointing center

(Thyagarajan et al. 2015b estimate SKA and HERA will suppress emission near the horizon 40 dB more than MWA), the only safe place for foreground “avoidance” is beyond the horizon.

These arguments have important ramifications for experiments looking to subtract foreground emission and recover k modes from inside the wedge. In particular, they suggest that experiments looking to “enlarge” the EoR window and remove foreground emission from modes near the horizon will benefit most from subtracting emission outside the main lobe. For this subtraction to be effective, accurate wide-field primary beam calibration is necessary to properly characterize the sidelobe patterns as a function of frequency. Such wide-field calibration may require new techniques, e.g., Pober et al. (2012), Yatawatta et al. (2013), and Neben et al. (2015). Additionally, these arguments motivate the need for low-frequency, wide-field sky surveys, especially in the Southern Hemisphere, where the vast majority of EoR-frequency 21 cm experiments are being constructed. Experiments such as HERA do not have a steerable beam, and thus will have difficulty measuring the flux densities of sources in their sidelobes. An accurate catalog produced by another survey covering a larger area (e.g., Jacobs et al. 2011; Williams et al. 2012; Hurley-Walker et al. 2014; Wayth et al. 2015) will be highly valuable for subtracting sources outside the main field of view. Northern Hemisphere experiments such as LOFAR and GMRT may also be valuable for characterizing the foregrounds at higher declinations.

Another important conclusion of this work is the implication that foreground subtraction cannot simply target the removal of some total amount of flux density independently of the position of that emission on the sky. Even if all emission from inside the primary field of view could be perfectly removed, sources in the sidelobes would continue to dominate higher k_{\parallel} modes. To recover all modes of the 21 cm power spectrum, foreground models must extend into any primary-beam sidelobes where the level of beam attenuation does not reduce the foreground power below that of the EoR signal. While attention has been paid to the removal of bright off-axis point sources (Offringa et al. 2012), the remaining diffuse emission and confused source background will still have significant spectral structure from the instrumental effects we have described. Given the extremely high foreground-to-signal ratio in 21 cm experiments, emission far from the pointing center cannot be neglected even if it is largely attenuated by the instrument’s primary beam.

In practice, wide-field source subtraction at the level needed to recover the EoR signal may require more than an accurate foreground model, especially for experiments with very wide fields of view such as PAPER. First, curved-sky effects become important near the horizon; an imaging-based analysis that does not correctly handle the curved sky (e.g., with w -projection; Cornwell et al. 2008) could wash out the input source model and reduce the amount of power subtracted. Second, ionospheric effects may become important for the level of accuracy needed in subtraction (Mitchell et al. 2008; Bernardi et al. 2009; Intema et al. 2009). The ionosphere may introduce frequency-dependent, time-dependent, and direction-dependent gains, all of which could lead to errors in model-based source subtraction if not corrected for. None of these effects alleviates the need for wide-field foreground subtraction, however; rather, they increase the difficulty of implementing a scheme that

potentially removes foregrounds to a level below the EoR signal.

8. CONCLUSIONS

In this work, we have presented a heuristic description of the imprint of the primary beam in power spectrum measurements from 21 cm interferometers. In particular, we find that wide-field effects—especially primary-beam sidelobes—leave a highly chromatic imprint, so that even smooth-spectrum emission that falls within the sidelobes corrupts high k_{\parallel} modes of the power spectrum. We further demonstrate this effect both with pedagogical simulations using single point sources and by removing a source model from MWA observations. When the model includes sources out to the first primary-beam sidelobes, it produces a significant (percent-level) reduction in power at high k_{\parallel} values.

This result has significant implications for experiments looking to measure the power spectrum of 21 cm emission from the EoR or any other epoch. In particular, it shows that foregrounds *must* be considered as a wide-field contaminant. Removing foregrounds from just the primary field of view will not reduce power in high k_{\parallel} modes corresponding to the primary-beam sidelobes. As a corollary, pipeline simulations that include only foregrounds within the primary field of view are missing a major contaminant of the EoR signal. Experiments looking to use foreground subtraction to enlarge the EoR window must also pay particular attention to emission in the sidelobes, since it is this emission that corrupts modes closest to the EoR window.

The authors wish to thank Adrian Liu for helpful conversations and our referee for a number of helpful suggestions that improved the paper. J.C.P. is supported by an NSF Astronomy and Astrophysics Fellowship under award AST-1302774. This scientific work makes use of the Murchison Radio-astronomy Observatory, operated by CSIRO. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the MWA comes from the U.S. National Science Foundation (grants AST-0457585, PHY-0835713, CAREER-0847753, and AST-0908884), the Australian Research Council (LIEF grants LE0775621 and LE0882938), the U.S. Air Force Office of Scientific Research (grant FA9550-0510247), and the Centre for All-sky Astrophysics (an Australian Research Council Centre of Excellence funded by grant CE110001020). Support is also provided by the Smithsonian Astrophysical Observatory, the MIT School of Science, the Raman Research Institute, the Australian National University, and the Victoria University of Wellington (via grant MED-E1799 from the New Zealand Ministry of Economic Development and an IBM Shared University Research Grant). The Australian Federal government provides additional support via the Commonwealth Scientific and Industrial Research Organization (CSIRO), National Collaborative Research Infrastructure Strategy, Education Investment Fund, and the Australia India Strategic Research Fund, and Astronomy Australia Limited, under contract to Curtin University. We acknowledge the iVEC Petabyte Data Store, the Initiative in Innovative Computing and the CUDA Center for Excellence sponsored by NVIDIA at Harvard University, and the International Centre for Radio Astronomy Research (ICRAR), a Joint Venture of Curtin University and The University of

Western Australia, funded by the Western Australian State government.

REFERENCES

- Ali, Z. S., Parsons, A. R., Zheng, H., et al. 2015, *ApJ*, 809, 61
 Ansari, R., Campagne, J.-E., Colom, P., et al. 2012a, *CRPhy*, 13, 46
 Ansari, R., Campagne, J. E., Colom, P., et al. 2012b, *A&A*, 540, A129
 Battye, R. A., Browne, I. W. A., Dickinson, C., et al. 2013, *MNRAS*, 434, 1239
 Bernardi, G., de Bruyn, A. G., Brentjens, M. A., et al. 2009, *A&A*, 500, 965
 Bernardi, G., Greenhill, L. J., Mitchell, D. A., et al. 2013, *ApJ*, 771, 105
 Bowman, J. D., Cairns, I., Kaplan, D. L., et al. 2013, *PASA*, 30, 31
 Bowman, J. D., Morales, M. F., & Hewitt, J. N. 2009, *ApJ*, 695, 183
 Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, *Natur*, 466, 463
 Chapman, E., Abdalla, F. B., Bobin, J., et al. 2013, *MNRAS*, 429, 165
 Chapman, E., Abdalla, F. B., Harker, G., et al. 2012, *MNRAS*, 423, 2518
 Cornwell, T. J., Golap, K., & Bhatnagar, S. 2008, *ISTSP*, 2, 647
 Datta, A., Bowman, J. D., & Carilli, C. L. 2010, *ApJ*, 724, 526
 Dillon, J. S., Liu, A., & Tegmark, M. 2013, *PhRvD*, 87, 043005
 Dillon, J. S., Liu, A., Williams, C. L., et al. 2014, *PhRvD*, 89, 023002
 Dillon, J. S., Tegmark, M., Liu, A., et al. 2015, *PhRvD*, 91, 023002
 Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *PhR*, 433, 181
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
 Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, 47, 1
 Hurley-Walker, N., Morgan, J., Wayth, R. B., et al. 2014, *PASA*, 31, 45
 Intema, H. T., van der Tol, S., Cotton, W. D., et al. 2009, *A&A*, 501, 1185
 Jacobs, D. C., Aguirre, J. E., Parsons, A. R., et al. 2011, *ApJL*, 734, L34
 Jacobs, D. C., Pober, J. C., Parsons, A. R., et al. 2015, *ApJ*, 801, 51
 Large, M. I., Mills, B. Y., Little, A. G., Crawford, D. F., & Sutton, J. M. 1981, *MNRAS*, 194, 693
 Liu, A., Parsons, A. R., & Trott, C. M. 2014a, *PhRvD*, 90, 023018
 Liu, A., Parsons, A. R., & Trott, C. M. 2014b, *PhRvD*, 90, 023019
 Liu, A., & Tegmark, M. 2011, *PhRvD*, 83, 103006
 Liu, A., Tegmark, M., Bowman, J., Hewitt, J., & Zaldarriaga, M. 2009, *MNRAS*, 398, 401
 Lonsdale, C. J., Cappallo, R. J., Morales, M. F., et al. 2009, *IEEEP*, 97, 1497
 Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, *ApJL*, 763, L20
 Mellema, G., Koopmans, L. V. E., Abdalla, F. A., et al. 2013, *ExA*, 36, 235
 Mitchell, D. A., Greenhill, L. J., Wayth, R. B., et al. 2008, *ISTSP*, 2, 707
 Morales, M. F., Bowman, J. D., & Hewitt, J. N. 2006, *ApJ*, 648, 767
 Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, 752, 137
 Morales, M. F., & Matejek, M. 2009, *MNRAS*, 400, 1814
 Morales, M. F., & Wyithe, J. S. B. 2010, *ARA&A*, 48, 127
 Neben, A. R., Bradley, R. F., Hewitt, J. N., et al. 2015, *RaSc*, 50, 614
 Offringa, A. R., de Bruyn, A. G., Biehl, M., et al. 2010, *MNRAS*, 405, 155
 Offringa, A. R., van de Gronde, J. J., & Roerdink, J. B. T. M. 2012, *A&A*, 539, A95
 Offringa, A. R., Wayth, R. B., Hurley-Walker, N., et al. 2015, *PASA*, 32, 8
 Paciga, G., Albert, J. G., Bandura, K., et al. 2013, *MNRAS*, 433, 639
 Parsons, A., Pober, J., McQuinn, M., Jacobs, D., & Aguirre, J. 2012a, *ApJ*, 753, 81
 Parsons, A. R., & Backer, D. C. 2009, *AJ*, 138, 219
 Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, *AJ*, 139, 1468
 Parsons, A. R., Liu, A., Aguirre, J. E., et al. 2014, *ApJ*, 788, 106
 Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. 2012b, *ApJ*, 756, 165
 Pober, J. C., Liu, A., Dillon, J. S., et al. 2014, *ApJ*, 782, 66
 Pober, J. C., Parsons, A. R., Aguirre, J. E., et al. 2013a, *ApJL*, 768, L36
 Pober, J. C., Parsons, A. R., DeBoer, D. R., et al. 2013b, *AJ*, 145, 65
 Pober, J. C., Parsons, A. R., Jacobs, D. C., et al. 2012, *AJ*, 143, 53
 Pritchard, J. R., & Loeb, A. 2012, *RPPH*, 75, 086901
 Santos, M. G., Cooray, A., & Knox, L. 2005, *ApJ*, 625, 575
 Shaw, J. R., Sigurdson, K., Pen, U.-L., Stebbins, A., & Sitwell, M. 2014, *ApJ*, 781, 57
 Slee, O. B. 1995, *AuJPh*, 48, 143
 Sullivan, I. S., Morales, M. F., Hazelton, B. J., et al. 2012, *ApJ*, 759, 17
 Sutinjo, A., O'Sullivan, J., Lenc, E., et al. 2015, *RaSc*, 50, 52
 Switzer, E. R., Masui, K. W., Bandura, K., et al. 2013, *MNRAS*, 434, L46
 Thompson, A. R., Moran, J. M., & Swenson, G. W., Jr. 2001, *Interferometry and Synthesis in Radio Astronomy* (2nd ed.; New York: Wiley-Interscience)
 Thyagarajan, N., Jacobs, D. C., Bowman, J. D., et al. 2015a, *ApJL*, 807, L28
 Thyagarajan, N., Jacobs, D. C., Bowman, J. D., et al. 2015b, *ApJ*, 804, 14
 Thyagarajan, N., Udaya Shankar, N., Subrahmanyan, R., et al. 2013, *ApJ*, 776, 6
 Tingay, S. J., Goetze, R., Bowman, J. D., et al. 2013, *PASA*, 30, 7

- Trott, C. M., Wayth, R. B., & Tingay, S. J. 2012, [ApJ](#), **757**, 101
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, [A&A](#), **556**, A2
- Vedantham, H., Udaya Shankar, N., & Subrahmanyam, R. 2012, [ApJ](#), **745**, 176
- Wang, J., Xu, H., An, T., et al. 2013, [ApJ](#), **763**, 90
- Wayth, R. B., Lenc, E., Bell, M. E., et al. 2015, [PASA](#), **32**, 25
- Williams, C. L., Hewitt, J. N., Levine, A. M., et al. 2012, [ApJ](#), **755**, 47
- Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. 2013, [A&A](#), **550**, A136
- Zaroubi, S. 2013, in *Astrophysics and Space Science Library*, Vol. 396, *The First Galaxies*, ed. T. Wiklind, B. Mobasher, & V. Bromm (Berlin: Springer), 45
- Zheng, H., Tegmark, M., Buza, V., et al. 2014, [MNRAS](#), **445**, 1084
- Zheng, Q., Wu, X.-P., Gu, J.-H., Wang, J., & Xu, H. 2012, [ApJL](#), **758**, L24