

FOUR FUNDAMENTAL FOREGROUND POWER SPECTRUM SHAPES FOR 21 CM COSMOLOGY OBSERVATIONS

MIGUEL F. MORALES¹, BRYNA HAZELTON¹, IAN SULLIVAN¹, ADAM BEARDSLEY¹

Draft version April 27, 2012

ABSTRACT

Contamination from instrumental effects interacting with bright astrophysical sources is the primary impediment to measuring Epoch of Reionization and BAO 21 cm power spectra—an effect called mode-mixing. In this paper we identify four fundamental power spectrum shapes produced by mode-mixing that will affect all upcoming observations. We are able, for the first time, to explain the wedge-like structure seen in advanced simulations and to forecast the shape of an ‘EoR window’ that is mostly free of contamination. Understanding the origins of these contaminations also enables us to identify calibration and foreground subtraction errors *below* the imaging limit, providing a powerful new tool for precision observations.

1. INTRODUCTION

Observations of redshifted 21 cm emission from neutral hydrogen have the potential to reveal the process of reionization and provide an important new tool for observational cosmology (see Furlanetto et al. 2006 and Morales & Wyithe 2010 for recent reviews). A number of experiments are currently under construction to observe the 21 cm power spectrum, including LOFAR (LOW Frequency ARray²), PAPER (Precision Array for Probing the Epoch of Reionization³), the MWA (Murchison Widefield Array⁴), and CHIME (Canadian Hydrogen Intensity Mapping Experiment⁵). The primary challenge for these observations is removing contamination from the bright astrophysical foregrounds interacting with instrumental effects and foreground mis-subtractions.

The astrophysical foregrounds are up to five orders of magnitude stronger than the EoR signal, and initial studies of the 21 cm foreground concentrated on identifying all of the potentially offending sources (e.g., Matteo et al. 2002; Oh & Mack 2003; Matteo et al. 2004; Gnedin & Shaver 2004; Santos et al. 2005; McQuinn et al. 2007; reviewed by Morales & Wyithe 2010). The general consensus is that all known astrophysical foregrounds are either spectrally smooth or at known editable frequencies (e.g., galactic radio recombination lines), and none of them mimics the spherical symmetry of the EoR’s redshifted emission line. The spectrally smooth astrophysical foreground emission dominates at small k_{\parallel} (line-of-sight wave numbers), but quickly falls below the EoR signal strength at higher k_{\parallel} values. Conceptually, purely angular modes are dominated by the foregrounds, but the EoR signal can be observed in the line-of-sight (frequency) modes (Zaldarriaga et al. 2004; Morales & Hewitt 2004; Jelić et al. 2008; Wang et al. 2006; Harker et al. 2009).

The difficulty is that no instrument is perfect, and small instrumental and observational effects can throw the strong angular foregrounds into the frequency (k_{\parallel})

direction. Effects include chromatic side lobes from sources producing line-of-sight ripples (Bowman et al. 2009; Liu et al. 2009), small calibration errors leading to mis-subtraction of the chromatic PSFs (Datta et al. 2010), Faraday rotation of the galactic synchrotron beating with polarization mis-calibrations (Geil et al. 2011), chromatic primary beams, and numerous other effects. Collectively these are called mode-mixing foregrounds, and have been the primary focus of the EoR foreground subtraction community for the past few years.

Precision simulations by Datta et al. (2010) have shown a very distinctive wedge shape in the k_{\parallel} vs. k_{\perp} power spectrum for three types of mode-mixing contamination—small amplitude errors in bright source removal, small position errors in the bright source removal, and small calibration errors. However, the origin of the wedge shape is not clear. In this paper we expand on the work by Liu & Tegmark (2011) and Vedantham et al. (2011) to develop a mathematical framework for subtractive foreground removal and mode-mixing contamination in §2. In §3 & §4 we then explore four mode-mixing signatures and develop a quantitative and qualitative understanding of their origin and why they are fundamental to the measurement process. We conclude in §5 by discussing how our new understanding will allow us to calibrate and identify foreground errors *below* the imaging limit.

As the first installment in an informal series of papers on mode-mixing, this paper concentrates on fundamental contamination common to all observations with subsequent papers exploring array-layout dependent effects and new statistical methods for mitigating these instrumental contaminations.

2. FRAMEWORK

Starting with the notation of Liu & Tegmark (2011) the band powers p_{α} from the measured data (\mathbf{m}) can be determined with a simplified version of their Equation 2

$$p_{\alpha} = \mathbf{m}^T \mathbf{E}_{\alpha} \mathbf{m}, \quad (1)$$

where \mathbf{E}_{α} is the power spectrum analysis and foreground subtraction, we have used the fact that interferometric data is naturally zero mean, and we have omitted the

¹ University of Washington, Seattle, 98195

² <http://www.lofar.org/>

³ <http://astro.berkeley.edu/~dbacker/eor/>

⁴ <http://www.mwatelescope.org/>

⁵ <http://www.physics.ubc.ca/chime/>

foreground bias term (in essence it is the bias term we are interested in calculating). For this paper it will be easier to split the matrix operator into two symmetric halves

$$\mathbf{E}_\alpha = \mathbf{D}_\alpha^T \mathbf{D}_\alpha, \quad (2)$$

each half describing the linear analysis and foreground subtraction pipeline. The processed data for each band $\mathbf{m}'(\mathbf{k}_\alpha)$ (the \mathbf{k} -space pixels for which $|\mathbf{k}|$ falls in the spherical shell of the band α) is then given by

$$\mathbf{m}'(\mathbf{k}_\alpha) = \mathbf{D}_\alpha(\mathbf{k}_\alpha, \mathbf{v}; \alpha) \mathbf{m}(\mathbf{v}) \quad (3)$$

and the band power by

$$p_\alpha = \mathbf{m}'_\alpha{}^T \mathbf{m}'_\alpha = (\mathbf{m}^T \mathbf{D}_\alpha^T)(\mathbf{D}_\alpha \mathbf{m}). \quad (4)$$

Throughout we will use the operator notation $\mathbf{A}(\mathbf{a}, \mathbf{b}; x)$ to signify a transformation from coordinate vector \mathbf{b} to coordinate vector \mathbf{a} that depends on parameters x , so the operator in Equation 3 is interpreted as transforming the data from the raw visibilities \mathbf{v} to the k -space pixels within a band \mathbf{k}_α , which depends on the band α we are interested in. A simplified version of a typical analysis pipeline would be

$$\mathbf{D}_\alpha(\mathbf{k}_\alpha, \mathbf{v}) = \mathbf{S}(\mathbf{k}_\alpha, \mathbf{k}) \mathbf{F}_{1D}(\mathbf{k}, \mathbf{u}) \tilde{\mathbf{B}}^T(\mathbf{u}, \mathbf{v}) \quad (5)$$

where the visibility data is gridded ($\tilde{\mathbf{B}}^T$) to the uv plane as a function of frequency (uv vs. frequency coordinates denoted by \mathbf{u}), the frequency direction is Fourier transformed along the line of sight (frequency) direction and the coordinates are mapped to cosmological wavenumbers \mathbf{k} , and the wavenumbers associated with the band power in question are selected (\mathbf{S}).

While most astrophysical foregrounds are smooth in frequency and thus separable from the nearly spherical EoR signal in principle (all emission near $k_\parallel = 0$, Morales et al. 2006; McQuinn et al. 2006), the effect of chromatic instrument responses, imperfect foreground models, and/or imperfect calibration conspire to throw this nearly purely angular foreground power into the line-of-sight k_\parallel direction, masking the faint EoR signal. Here we label the power spectra due to various mode mixing terms as $\mathbf{p}_e(\mathbf{k})$, which is defined as the square of the data residuals $\mathbf{r}(\mathbf{k})$ in k -space ($\mathbf{p}_e = |\mathbf{r}|^2$). We can formally calculate the shapes of these mode-mixing power spectra by considering the difference between the observed and subtracted foregrounds. Looking at the residuals for a simplified software holography chain we have

$$\begin{aligned} \mathbf{r}(\mathbf{k}) &= \mathbf{F}_{1D}(\mathbf{k}, \mathbf{u}) \tilde{\mathbf{B}}^T(\mathbf{u}, \mathbf{v}) \mathbf{M}(\mathbf{v}, \mathbf{u}) \mathbf{F}(\mathbf{u}, \theta) I(\theta) \\ &\quad - \mathbf{F}_{1D}(\mathbf{k}, \mathbf{u}) \tilde{\mathbf{B}}^T(\mathbf{u}, \mathbf{v}) \mathbf{B}(\mathbf{v}, \mathbf{u}) \mathbf{F}(\mathbf{u}, \theta) \hat{I}(\theta), \end{aligned} \quad (6)$$

or annotated

$$\begin{aligned} \mathbf{r}(\mathbf{k}) &= \overbrace{\mathbf{F}_{1D}(\mathbf{k}, \mathbf{u}) \tilde{\mathbf{B}}^T(\mathbf{u}, \mathbf{v})}^{\text{observed foreground}} \underbrace{\mathbf{M}(\mathbf{v}, \mathbf{u})}_{\text{instrument}} \underbrace{\mathbf{F}(\mathbf{u}, \theta) I(\theta)}_{\text{true foreground}} \\ &\quad - \overbrace{\mathbf{F}_{1D}(\mathbf{k}, \mathbf{u}) \tilde{\mathbf{B}}^T(\mathbf{u}, \mathbf{v})}^{\text{subtracted foreground}} \underbrace{\mathbf{B}(\mathbf{v}, \mathbf{u})}_{\text{model instrument}} \underbrace{\mathbf{F}(\mathbf{u}, \theta) \hat{I}(\theta)}_{\text{model foreground}}. \end{aligned}$$

The first line of Equation 6 shows the observed foreground as the true foreground I measured by the instrument \mathbf{M} and run through a holographic analysis $\tilde{\mathbf{B}}^T$. The second line is the subtracted foreground, with the model foreground \hat{I} and model instrument \mathbf{B} replacing the true foreground and instrument. The analysis portion of how visibilities are transformed into k -space measurements are usually the same for both the measurement and the model (they can differ for computational reasons), so we can introduce an analysis operator $\mathbf{A}(\mathbf{k}, \mathbf{v}) = \tilde{\mathbf{B}}^T$.

In general the residual foreground contamination is due to an admixture of model errors ($\hat{I} \neq I$) and calibration errors ($\mathbf{M} \neq \mathbf{B}$), but it is instructive to look at these sources of errors independently

$$\begin{aligned} \mathbf{r}_M &= \mathbf{A}\mathbf{B}(I - \hat{I}) \quad \text{model error, calibration correct} \quad (7) \\ \mathbf{r}_C &= \mathbf{A}(\mathbf{M} - \mathbf{B})I \quad \text{calibration error, model correct.} \quad (8) \end{aligned}$$

Formally our work is now done. We can parametrize the kinds of errors we make, e.g. amplitude and or position model errors in Equation 7, and calculate the associated power spectra $\mathbf{p}_e(\mathbf{k})$ for our instrument.

However, it is useful to qualitatively identify the origin of these shapes. In §3 & 4 we identify four fundamental residual power spectrum shapes that will be seen by all upcoming instruments and describe their origins. In this paper we concentrate specifically on the contamination that will be seen by any upcoming observation, independent of the array layout. In subsequent papers we explore the effects of array layout and non-uniform sampling of the uv plane (Hazelton et al. in preparation), and more advanced statistical techniques for removing the instrumental effects. For this paper we limit ourselves to effects related to fundamental information loss on the antenna scale.

3. MODEL ERROR SHAPE AND THE FOREGROUND WEDGE

First we consider the effect of foreground model errors. In this case we assume the calibration of the instrument is perfect—both the complex gain and beam shape of each antenna is known exactly—but the subtracted foreground model is not equal to the true sky (Equation 7). Typical model errors include mistakes in the amplitude and location of compact sources, errors in the amplitude and shape of diffuse emission, and polarization errors for both compact and diffuse sources. Both thermal noise and systematics contribute to foreground model errors, and the PSF sidelobes often make it impossible to disentangle the true brightness and locations of the sources in the field. Accurately determining the foregrounds is particularly difficult in the confusion limit, and Liu et al. (2009) and Bowman et al. (2009) have done excellent work on improving the accuracy of the foreground models in the confusion limit.

In Figure 10 of Datta et al. (2010) random errors in the foreground model generated a very distinctive 2D k -space “wedge” which has now been seen in observations with the MWA prototype (Bernardi, et al. in preparation). As we will explain in this section, this wedge is due to the chromatic instrument response and inherent information loss.

Any interferometric measurement suffers information loss on the scale of one antenna. An antenna sums the electric field across its surface into a single voltage signal which is cross-correlated with the integrated electric field from a second antenna to form a visibility. In this process the variations in the electric field correlation across the antenna surface (e.g. from a source far from field center) are lost. Equivalently, in the uv plane one visibility is formed by integrating the true uv correlation with the beam pattern of that antenna pair $\mathbf{B}(\mathbf{v}, \mathbf{u})$ over a small frequency channel as shown in Equation 6. Modern array simulators use this integral of the uv plane to predict the visibilities including direction dependent gain patterns. (For a conceptual introduction to interferometric measurement, please see Chapter 3 of Morales & Wyithe 2010). The first step of any analysis is to reconstruct an estimate of the uv plane $\hat{I}(\mathbf{u})$ by gridding the visibilities ($\tilde{\mathbf{B}}^T$ in Equation 6). However, \mathbf{B} is not invertible and the operation $\tilde{\mathbf{B}}^T \mathbf{B}$ does not recreate the true uv plane distribution of the sky.

Various statistical priors can be used to mitigate this information loss, and deconvolution is a powerful technique for reconstructing the contributions to each visibility below the antenna scale. Effectively the statistical prior that most of the flux is due to a list of sources (or CLEAN components) allows their contribution to each visibility to be determined accurately despite the smearing due to the antenna integration in the uv plane.

However, the EoR signal is below the imaging noise floor for the first generation observatories (S/N less than one per per visibility, or equivalently, phase noise greater than one radian). Thus deconvolving the EoR signal is not possible, and we must inherently work with a ‘dirty’ residual map (this is also true for all CMB measurements and other diffuse power spectrum measurements). Foreground model errors appear as small spurious signals in the EoR dirty map. First we will explore the effect of amplitude errors which exhibit the basic features common to several of our model and calibration error signatures, then we will describe source locations errors in the subsequent sub-section.

3.1. Amplitude errors

An error in source brightness will leave a small residual positive or negative source at the location of the true source, and is conceptually the simplest of the contaminations. Figure 1 shows the true correlation of a flat spectrum source far from the field center in the \mathbf{u}, f space (or equivalently k_{\perp} and line-of-sight Mpc) with a few baselines overlaid. The oscillation in k_{\perp} is given by the distance from the field center (quick oscillations for sources near field edge) and for a flat spectrum source amplitude is constant in frequency. The baselines all migrate to larger distances (in wavelengths) with higher frequency, with the slope of the chromatic migration increasing with baseline length.

Figure 2 then zooms in on a small region of Figure 1 to explore the effects of measurement (lefthand panel) and reconstruction (righthand panel) for a single baseline. In the left hand panel the true correlation for the residual source (vertical corrugations) is integrated by the antenna beam (grey rectangles) to form the residual visibility. The measured visibility traces the true corre-




FIG. 1.— This cartoon represents a slice through the three dimensional \mathbf{u}, f space, where u is proportional to k_{\perp} and f is proportional to line-of-sight distance in Mpc. The vertical corrugation shows the spatial-frequency corrugation of a flat spectrum source away from the center of the field (real part shown). The diagonal stripes show the paths of individual antenna baselines as a function of frequency. While the distance between two antennas in wavelengths (or k_{\perp}) increases with frequency for all baselines, the mis-alignment angle is larger for longer baselines. It is this increasing mis-alignment of the baselines with the corrugations of a residual source that create the foreground wedge.

lation as shown by the size of the hexagons and the line width. (Mathematically this is a sampled convolution in \mathbf{u} of the beam and the true correlation—or the true correlation times the direction dependent gain of that antenna pair.) As expected, the measured visibility oscillates as a function of frequency due to the chromatic increase in baseline length.

The difficulty comes in the reconstruction of an unbiased estimate of the sky. As previously mentioned we have lost information both where there are no baselines (missing spatial frequencies) and on scales smaller than the antenna in the process of measuring the signal. Because the signal to noise on the EoR signal is much less than one for a single baseline, we must resort to gridding the visibilities to form a dirty map. Gridding with the holographic antenna pattern guarantees an unbiased power spectrum measurement (Morales & Matejcek 2009; Tegmark 1997), but the exact form of the gridding is not important for this discussion.

The righthand panel of Figure 2 shows the effect of gridding for one baseline. The sinusoidal variations in the visibility are spread out by the gridding kernel (boxes), resulting in regions of high and low amplitude repre-




FIG. 2.— These figures zoom in two a small region of the \mathbf{u}, f plane to study the effect of baseline chromaticity. In the left hand panel, the vertical corrugations from a residual source are shown, with a line showing the measured visibility. The measured visibility is the integral of the true correlation function (grey corrugations) times the antenna beam (grey boxes in left hand figure). The resulting measured visibility wraps as a function of frequency, as indicated by the line width and the size of the hexagons. The right hand figure then shows the result of the gridding process, where the measured visibilities are gridded to the *reconstructed* \mathbf{u}, f plane. Effectively the measured visibility is smeared over a region of \mathbf{u} by the gridding kernel. Light grey indicates a small (real component) reconstructed in that region of \mathbf{u}, f , while dark grey indicates a large value. When we Fourier transform in the line-of-sight direction to form the power spectrum cube (Figure 4), we will see oscillations in the line-of-sight direction as indicated by the oscillating grey values along the vertical dashed line.

sented by the grey scale. This produces oscillations in the line-of-sight direction as seen by the varying amplitude along the dashed line. This oscillation is not in the original source (lefthand panel), but because we lost information on the scale of the antenna unavoidably shows up in our reconstructed signal (right hand panel). When we Fourier transform in the line-of-sight direction to get into the three dimensional \mathbf{k} -space this reconstruction error throws power that was originally purely in the k_{\perp} direction into the k_{\parallel} dimension—the hallmark of mode-mixing.

Figure 3 shows a precision simulation similar to the simulations by Datta et al. (2010) of a single mis-subtracted off-axis point source as seen by a simplified array. In addition to the phase wrap predicted in Figure 2b, the longer baselines produce higher frequency contamination due to their higher chromatic migration angles. Mathematically, the wavelength of the residual source in the \mathbf{u} direction λ_{\perp} is related to the wavelength of the contamination in the line-of-sight λ_{LoS} direction by the baseline to wavelength conversion $\mathbf{u}(\lambda) = \mathbf{u}(m)f/c$, giving the relationship

$$\frac{\lambda_{\text{LoS}}}{\lambda_{\perp}} = \frac{c}{\mathbf{u}(m)} = \frac{f}{\mathbf{u}(\lambda)}. \quad (9)$$

Fourier transforming in the line-of-sight direction and converting to cosmological coordinates (Morales & Hewitt 2004) one obtains, after a little algebra,

$$\mathbf{k}_{\parallel} \approx \theta'(\text{rad}) k_{\perp} \left(\frac{D_{\text{M}}(z)}{D_{\text{H}}} \frac{E(z)}{(1+z)} \right). \quad (10)$$

$\theta'(\text{rad})$ is the vector distance of the residual source from the field center, k_{\perp} is the baseline length in cosmological coordinates, and the cosmological terms in parenthesis range from 0.63 at a redshift of 1 to 3.4 at a redshift of 8 ($\Omega_{\Lambda} = 0.728$, $h = 0.704$, flat universe following notation of Hogg 1999).

The contamination in Equation 10 is remarkably simple: a power law of 1 relationship between k_{\perp} and k_{\parallel} , depending only on the angle of the source from the field center and an order one constant based on the cosmology. Figure 4 shows how this contamination will appear in the k_{\perp} vs. k_{\parallel} space. This is an approximate relationship, as only a small length of the line-of-sight oscillation is contributed by one baseline, and the contributions of different baselines (at same k_{\perp}) are incoherent. The sum of these short incoherent oscillations at the same wavelength is a strong peak at the associated line-of-sight wavenumber k_{\parallel} , but will contain power at other




FIG. 3.— This figure represents a uf slice through a precision simulation of a simple array, and illustrates the contamination described in Figures 1–4. On the right half of the figure (positive u) there is a single baseline in a reconstructed uf plane (analogous to Figure 2b), and on the left side there are numerous baselines of different lengths. The color scale shows the reconstructed phase for a single residual source 14° from the field center (black = $-\pi$, white = $+\pi$, grey = 0 phase or no baseline contribution). The phase wrap is clearly seen with increasing frequency in the single baseline on the right, and the increasing rate of phase wrap is seen in the longer baselines on the left.

wavelengths and the magnitude of the mode-mixing will decrease as more baselines are added (smoother PSF).

There is, however, a soft limit to the line-of-sight contamination imposed by the instrument field-of-view. The corrugations of the residual sources are band limited by the instrument field-of-view, giving a maximum line-of-sight contamination of

$$\mathbf{k}_{\parallel \text{Max}} \approx \frac{1}{2} \text{FoV}(\text{radians}) k_{\perp} \left(\frac{D_M(z)}{D_H} \frac{E(z)}{(1+z)} \right). \quad (11)$$

This identifies a convenient EoR window above $\mathbf{k}_{\parallel \text{Max}}$ where foreground model errors will not contaminate the EoR signal (after Vedantham et al. 2011). As we will see in later sections, the $\mathbf{k}_{\parallel \text{Max}}$ line plays an important role in determining the region where the EoR signal can be measured and the relative importance of different contaminations.

The small residual sources from amplitude errors have a constant magnitude as a function of k_{\perp} (FT of a δ -function), so the amplitude along any dashed line in Figure 4 is constant. However, different lines in the wedge correspond to sources at different (projected) distances from the phase center so the amplitude can vary perpendicular to the $\mathbf{k}_{\parallel \text{Max}}$ line. The characteristic shape of amplitude errors is a wedge of power below $\mathbf{k}_{\parallel \text{Max}}$ that is constant parallel to the $k_{\parallel} \propto k_{\perp}$ diagonal but varies in the orthogonal $k_{\parallel} \propto -k_{\perp}$ direction, with errors closer to the field edge appearing higher in the wedge.




FIG. 4.— This figure shows a slice through the k_{\parallel} vs. k_{\perp} space on logarithmic coordinates. The residual contamination described in Figures 1 & 2 appears as a line of contamination along $k_{\parallel} \propto k_{\perp}$, with the intercept depending on distance from the phase center. The instrumental field-of-view provides a soft cutoff in the contamination, leaving an EoR window that is free of foreground model errors. The amplitude of the contamination depends on the PSF of the instrument (smoother PSF leads to lower contamination), but in general it will be difficult to remove foreground contamination from within the mode-mixing wedge. For amplitude errors the contamination is uniform in amplitude along lines of $k_{\parallel} \propto k_{\perp}$, while the contamination from location errors increases \sim linearly with $|k|$.

In the simulations by Datta et al. (2010), mis-subtracted sources were scattered uniformly across the image, creating many such diagonal lines for sources at different projected distances from the image center. These many diagonal lines in the power spectrum added together to produce the wedge seen in that work.

WHY DON'T WE JUST...

The origin of the foreground wedge is the inherent information loss below the antenna scale, coupled with our inability to deconvolve when the signal-to-noise is less than one. However, this is not at all self apparent and it is natural to suggest a number of alternatives. Common ideas include:

Why don't we just grid in meters instead of wavelengths? Gridding in meters skews Figures 1 & 2 so the baselines are vertical but the corrugations of the residual source and the line-of-sight Fourier transform are at an

angle (which increases with baseline length). The mode-mixing is unchanged by this coordinate transformation—the same effect occurs, just in a skewed coordinate frame. While gridding in meters may be more computationally efficient, the effect described here is independent of the gridding coordinates.

Why don't we Direct Fourier Transform the visibilities without gridding? This seems like the obvious solution as it eliminates the gridding step entirely, but there are two reasons this does not solve our problem. First, the antenna field-of-view changes significantly with frequency (even if all the antennas are identical to each other), and we must be able to insert this into our analysis. One way of doing this is to form maps at different frequencies via DFT, which are then tapered with the primary beam before assembling into an image cube (this is what NVSS did to assemble mosaics). However, multiplying by an image space taper is identical to convolving in the uv plane. This smooths out our visibilities in the uv plane, reproducing the effect seen in the righthand panel of Figure 2. Gridding with the antenna beam is identical to tapering in the image (this duality is the basis of software holography). Any image tapering reproduces this effect.

One can consider performing a full sky DFT and not applying a taper (as one might be tempted to do if the antenna beam was frequency independent), but this turns out to be worse for a power spectrum analysis. Effectively, contributions from a visibility are reconstructed across the full sky, even in areas where the telescope had no sensitivity (e.g. far outside the primary beam). This introduces *bias* into the power spectrum measurement (Tegmark 1997). Fundamentally we know the contributions to a single visibility came from a region of the uv plane, not a point, and we need to reconstruct this effect. We can do this either by gridding in uv or tapering the image, but not doing either actually biases our result by over emphasizing areas of the sky.

What if we only include visibilities which exactly overlap in \mathbf{u} when performing the line-of-sight FT? The first problem is that if the antenna field-of-view changes, the integral of the uv plane that went into making those visibilities differs. This reproduces the difficulty discussed in the previous question and tapering/gridding is needed to calibrate the data. Even if the antennas are not chromatic, this represents a drastic decrease in the effective collecting area. One cannot design a two dimensional antenna layout that provides the necessary uv coverage without discarding the majority of the baselines (even a regular grid of antennas only has good redundancy along the two grid directions).

What about deconvolving the EoR signal itself? This could be done for the SKA or other instruments which will image the EoR and BAO fluctuations. However, deconvolution requires a signal to noise per visibility greater than one (phase noise \lesssim one radian) and the first generations of instruments will not have the sensitivity.

What if we measured every baseline (at the scale of $\lambda/2$ or better)? This would eliminate the contamination we have identified, not by removing the wedge but by driving its amplitude to zero. With dense enough baseline coverage the incoherent summing of many line-of-sight oscillations in the righthand panel of Figure 2 would push the amplitude to zero. However, this criteria is equiva-

lent to building an instrument with a δ -function PSF. Desirable but impractical.

We believe that the foreground wedge identified here is fundamental. It is a product of the information loss at the antenna scale suffered in making the interferometric measurement, and the need to reconstruct a faithful ‘dirty’ reconstruction of the sky.

3.2. Location errors

Location errors assume that the amplitude of the model source is correct, but the location is slightly wrong—typically by much less than one beamwidth. In the absence of calibration errors, this is equivalent to having two δ -functions with opposite sign very close to one another. After transforming to the \mathbf{u}, f frame of Figures 1 & 2, the Fourier shift theorem shows the contamination to be the subtraction of two sinusoidal corrugations of slightly different frequency. These subtract perfectly at zero k_{\perp} with the amplitude growing with k_{\perp} as the corrugations slowly creep out of phase. In the limit of offsets much smaller than one beam (\ll one radian of phase at longest baseline), this gives a sinusoidal corrugation whose magnitude increases approximately linearly with k_{\perp} .

We can then repeat the argument of the previous section and Figures 1–4, the only difference being that the source of the contamination is a linearly increasing function of k_{\perp} instead of a constant. This limits the contamination to the same region as the amplitude errors, with the amplitude of the contamination proportional to $|k|$ along lines of $k_{\parallel} \propto k_{\perp}$. Thus location errors preserve the same EoR window above $\mathbf{k}_{\parallel \text{Max}}$ but with a distinctive linear increase in amplitude along lines of $k_{\parallel} \propto k_{\perp}$. This is the origin of the distinctive wedge shape seen in Figure 10 of Datta et al. (2010).

4. CALIBRATION ERROR SHAPES

In this section we consider the effect of calibration errors. We now assume that the foreground model is equal to the true sky, but that the model instrumental response is not equal to the true response ($\mathbf{B} \neq \mathbf{M}$, Equation 8). There are a number of ways the instrumental calibration can be wrong, from simple gain errors (visibility amplitude error) to errors in the antenna FoV or gain profile (wrong integral of the true uv plane) to assuming all antennas are identical (not accounting for antenna-to-antenna variation) to antenna location or survey errors (error in location of a visibility) or full fledged direction and frequency calibration errors.

To begin with we will concentrate on small frequency-independent errors on a subset of the baselines. Simple gain errors and differences between antennas will be reflected in a few visibilities with the wrong complex amplitude. If most of the visibility calibrations are correct (and the model is correct), the correct visibilities will subtract perfectly leaving only the erroneous component of the miscalibrated baselines. These calibration errors lead to two residual power spectrum shapes—a diffuse component that contaminates nearly all \mathbf{k} modes and a third wedge shape that is similar to the foreground model residuals of §3 but with a distinct functional form.

Figure 5 shows the effect of this calibration residual, where the dotted visibilities have subtracted perfectly




FIG. 5.— This cartoon shows the effect of calibration errors in a portion of the \mathbf{u}, f (or k_{\perp} , line-of-sight Mpc) space. The diagonal lines show the paths of a number of visibilities (unequally spaced). The dotted visibilities are correctly calibrated and have subtracted perfectly, but one visibility had a small calibration error so left a residual visibility (solid line). This residual visibility is gridded to the u, f plane by the grey shaded region. After Fourier transforming along the frequency direction (dashed line), this residual visibility produces contamination at all k_{\parallel} modes. Effectively it appears as the Fourier transform of a windowed δ -function (see text for details).

but the black visibility has not. An error in the gain amplitude leaves a residual visibility that is in phase with the underlying foreground signal, while a phase error leaves a residual visibility that is $\pm\pi/2$ rad out of phase. The residual complex visibility is then gridded to the u, f plane as shown by the grey region. When taking a Fourier transform along the line-of-sight direction (indicated by the dashed line) this gain error appears as a windowed δ -function. Effectively the calibration error for a visibility produces a sharp feature in the frequency direction, producing contamination over a wide range of k_{\parallel} modes.

The contamination in k_{\parallel} vs. k_{\perp} space appears as a flat contamination at all k_{\parallel} (δ -function FT) times a window function that falls with increasing k_{\parallel} , as shown in Figure 6. If the visibilities are gridded with the beam pattern $\mathbf{B}(\mathbf{v}, \mathbf{u})$, the windowing function is a scaled version of the antenna beam pattern—the dashed line in Figure 5 traverses the uv beam pattern and is Fourier transformed to produce a scaled version of the angular beam pattern in the k_{\parallel} direction. This is the most insidious of the foregrounds detailed in this paper because the contamination extends into the ‘EoR window.’




FIG. 6.— This cartoon shows the effect of calibration errors in the k_{\parallel} vs. k_{\perp} space. The first component reaches to high k_{\parallel} , contaminating the EoR window. This contamination is windowed (represented by grey scale), and the effectiveness of this windowing depends on the details of the analysis. There is also a wedge component due to calibration errors that is similar to the two foreground wedges, but depends on the total baseline error with a $\sim k_{\perp}$ functional form.

The calibration errors also produce another wedge shape. The visibilities with calibration errors see sources from across the sky and oscillate with frequency just like in Figure 2. This again produces the same kind of \mathbf{k} space contamination, but because the sky is filled with sources it naturally contains sources from all distances from phase center up to the field-of-view cutoff. The difference is that the contamination is only over the range of k_{\perp} (or \mathbf{u}) impacted by that baseline. For the random calibration errors on all baselines in Figure 11 of Datta et al. (2010), this produces a wedge of contamination below $k_{\parallel \text{Max}}$ with a functional form that follows the baseline distribution. Unlike the foreground model errors which are functions of the diagonal $k_{\parallel} = k_{\perp}$ and $k_{\parallel} = -k_{\perp}$ axes, the calibration error is a function of the baseline distribution and is \sim parallel to k_{\perp} .

The two simple shapes presented here are for frequency independent calibration errors. Frequency dependent calibration errors produce residuals that oscillate along the visibilities’ lengths in f and map almost directly into the k_{\parallel} direction. It is hoped that many of these calibration errors will be slow and map to small k_{\parallel} , but they will generate additional foreground components that are not qualitatively described here. It is, however, straightforward

ward to calculate their shapes. The kinds of errors that appear in the calibration \mathbf{B} can be parametrized, and for each kind of error the residual power spectrum $\mathbf{p}(\mathbf{k})$ can be explicitly calculated using Equation 8. This will be an important step for all 21 cm power spectrum measurements, but also is quite instrument dependent. Here we have concentrated on four power spectrum residuals that will be important for all upcoming observations.

5. DISCUSSION

Using the results of §2–§4 we can quantitatively and qualitatively understand the origin of some of the different wedge-shaped structures seen in advanced simulations. This understanding also allows us to invert the process and associate observed power spectrum shapes with specific calibration and foreground subtraction errors. In practice, we feel it is this ability to identify the origin of observed power spectrum contaminations that will be the most influential result of this work.

One of the key problems facing EoR observations is identifying calibration and foreground subtraction errors below the imaging limit. When you are working several orders of magnitude below the confusion limit, different subtraction or calibration algorithms can produce nearly identical images with very different power spectra. Further, when looking at measured power spectra in k_{\parallel} vs. k_{\perp} , to date it has been impossible to identify the causes of observed contamination. Using the developments in this article, not only can we compare the performance of these analyses, we can identify the specific sources of observed contamination. A few examples might include:

- Two calibration algorithms that both produce good images can be compared, and their contribution to the power spectrum in the EoR window can be used

to pick the superior algorithm (§4).

- Symmetric bands of power as seen in Figure 4 (accounting for projection), can be used to identify higher source subtraction errors near the field edge (§3.1).
- Linearly increasing bands of power can be used to identify pointing offsets in particular regions of the survey (e.g., far from a calibrator or near the ionospheric equator, §3.2).
- Systematic errors can be estimated by observing the contributions of specific foregrounds and their covariance with the EoR signal (§2, Morales et al. 2006).

In this paper we have described four types of mode-mixing contamination any EoR instrument will see. The next paper in this series will explore the array-dependent effects and the influence of inhomogeneities in the baseline distribution (Hazelton et al. in preparation). Future papers will use these results to identify and quantify the mode-mixing contamination seen in 32 antenna MWA observations, and develop more advanced statistical methods to mitigate the contaminations due to calibration errors and array inhomogeneities.

ACKNOWLEDGEMENTS

This work has been supported by the National Science Foundation Astronomy Division through CAREER award #0847753 and NSF Postdoctoral Fellowship #1003314, and by the University of Washington. We'd particularly like to thank Harish Vedantham, Udaya Shankar, Ravi Subrahmanyam, Judd Bowman for the discussions that initiated this work.

REFERENCES

- Bowman, J. D., Morales, M. F., & Hewitt, J. N. 2009, *The Astrophysical Journal*, 695, 183
- Datta, A., Bowman, J., & Carilli, C. 2010, *The Astrophysical Journal*, 724, 526
- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Physics Reports*, 433, 181, Elsevier B.V.
- Geil, P., Gaensler, B., & Wyithe, J. 2011, *Monthly Notices of the Royal Astronomical Society*, 1, 1416
- Gnedin, N. Y., & Shaver, P. A. 2004, *The Astrophysical Journal*, 608, 611 (c) 2004: The American Astronomical Society
- Harker, G., et al. 2009, *Monthly Notices of the Royal Astronomical Society*, 397, 1138 (c) Journal compilation © 2009 RAS
- Hogg, D. W. 1999, arXiv, astro-ph
- Jelić, V., et al. 2008, *Monthly Notices of the Royal Astronomical Society*, 389, 1319
- Liu, A., & Tegmark, M. 2011, *Physical Review D*, 83, 103006
- Liu, A., Tegmark, M., & Zaldarriaga, M. 2009, *Monthly Notices of the Royal Astronomical Society*, 394, 1575
- Matteo, T. D., Ciardi, B., & Miniati, F. 2004, *Monthly Notices of the Royal Astronomical Society*, 355, 1053 (c) 2004 RAS
- Matteo, T. D., Perna, R., Abel, T., & Rees, M. J. 2002, *The Astrophysical Journal*, 564, 576 (c) 2002: The American Astronomical Society
- McQuinn, M., Lidz, A., Zahn, O., Dutta, S., Hernquist, L., & Zaldarriaga, M. 2007, *Monthly Notices of the Royal Astronomical Society*, 377, 1043
- McQuinn, M., Zahn, O., Zaldarriaga, M., Hernquist, L., & Furlanetto, S. R. 2006, *The Astrophysical Journal*, 653, 815 (c) 2006: The American Astronomical Society
- Morales, M., & Wyithe, J. 2010, *Annual Review of Astronomy and Astrophysics*, 48, 127
- Morales, M. F., Bowman, J. D., & Hewitt, J. N. 2006, *The Astrophysical Journal*, 648, 767 (c) 2006: The American Astronomical Society
- Morales, M. F., & Hewitt, J. N. 2004, *The Astrophysical Journal*, 615, 7 (c) 2004: The American Astronomical Society
- Morales, M. F., & Matejek, M. 2009, *Monthly Notices of the Royal Astronomical Society*, 400, 1814
- Oh, S. P., & Mack, K. J. 2003, *Monthly Notices of the Royal Astronomical Society*, 346, 871
- Santos, M. G., Cooray, A., & Knox, L. 2005, *The Astrophysical Journal*, 625, 575 (c) 2005: The American Astronomical Society
- Tegmark, M. 1997, *Physical Review D*, 55, 5895 (c) 1997: The American Physical Society
- Vedantham, H., Shankar, N. U., & Subrahmanyam, R. 2011, arXiv, astro-ph.IM
- Wang, X., Tegmark, M., Santos, M. G., & Knox, L. 2006, *The Astrophysical Journal*, 650, 529 (c) 2006: The American Astronomical Society
- Zaldarriaga, M., Furlanetto, S. R., & Hernquist, L. 2004, *The Astrophysical Journal*, 608, 622 (c) 2004: The American Astronomical Society