Title
Overcoming the Challenges of 21cm Cosmology

Author
Pober, Jonathan

Publication Date
2013

Peer reviewed|Thesis/dissertation

Overcoming the Challenges of 21cm Cosmology

By

Jonathan Pober

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Astrophysics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Aaron Parsons, Chair
Professor Carl Heiles
Professor Adrian Lee

Spring 2013

Overcoming the Challenges of 21cm Cosmology

Abstract

Overcoming the Challenges of 21cm Cosmology

by

Jonathan Pober

Doctor of Philosophy in Astrophysics

University of California, Berkeley

Professor Aaron Parsons, Chair

The highly-redshifted 21cm line of neutral hydrogen is one of the most promising and unique probes of cosmology for the next decade and beyond. The past few years have seen a number of dedicated experiments targeting the 21cm signal from the Epoch of Reionization (EoR) begin operation, including the LOw-Frequency ARray (LOFAR), the Murchison Widefield Array (MWA), and the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER). For these experiments to yield cosmological results, they require new calibration and analysis algorithms which will need to achieve unprecedented levels of separation between the 21cm signal and contaminating foreground emission. Although much work has been spent developing these algorithms over the past decade, their success or failure will ultimately depend on their ability to overcome the complications associated with real-world systems and their inherent complications.

The work in this dissertation is closely tied to the late-stage commissioning and early observations with PAPER. The first two chapters focus on developing calibration algorithms to overcome unique problems arising in the PAPER system. To test these algorithms, I rely on not only simulations, but on commissioning observations, ultimately tying the success of the algorithm to its performance on actual, celestial data. The first algorithm works to correct gain-drifts in the PAPER system caused by the heating and cooling of various components (the amplifiers and above ground co-axial cables, in particular). It is shown that a simple measurement of the ambient temperature can remove $\sim 10\%$ gain fluctuations in the observed brightness of calibrator sources. This result is highly encouraging for the ability of PAPER to remove a potentially dominant systematic in its power spectrum and cataloging measurements without resorting to a complicated system overhaul.

The second new algorithm developed in this dissertation solves a major calibration challenge not just for PAPER, but for nearly all of a large class of new wide-field, drift-scanning radio telescopes: primary beam calibration in the presence of a poorly measured sky. Since these telescopes lack the ability to steer their primary beams, while seeing nearly the entire sky at once, a large number of calibrator sources are necessary to probe the entire beam

response. However, the catalogs of radio sources at low-frequencies are not reliable enough to achieve the level of primary beam accuraccy needed for 21cm cosmology experiments. I develop, test, and apply a new technique which — using only the assumption of symmetry around a 180° rotation — simultaneously solves for the primary beam and the flux density of large number of sources.

In this dissertation, I also present the analysis of new observations from PAPER to test theoretical models which predict foreground emission is confined to a "wedge"-like region of cosmological Fourier space, leaving an "EoR window" free from contamination. For the first time in actual observations, these predictions are spectacularly confirmed. In many ways, this result shifts the burden for upcoming PAPER analysis from foreground removal to increased sensitivity. And although increasing sensitivity is no small feat in-and-of-itself, this result is highly encouraging for 21cm studies, as foreground removal was long-viewed as the principal challenge for this field.

The final result in this dissertation is the application of the all the lessons learned building PAPER and the MWA to design a new experiment for 21cm studies at $z \sim 1$ with the goal of measuring baryon acoustic oscillations (BAO). The design of the BAO Broadband and Broad-beam (BAOBAB) Array is described, and cosmological forecasts are presented. The bottom line is highly encouraging, suggesting that $z \sim 1$ 21cm observations can detect the neutral hydrogen power spectrum with a very modest ($16 - 32$ element) array, and that still reasonably sized ($128 - 256$ elements) arrays can produce significant advances in our knowledge of dark energy.

Contents

Acknowledgments

So much work has been done the past few years, and so many people have supported me — it's an impossible task to summarize it all in a few acknowledgments (especially when you're working on a deadline). My sincerest apologies to anyone left out.

The PAPER project on which so much of this dissertation is based is supported through the NSF-AST program (awards 0804508, 0901961, 1129258, and 1125558), the Mt. Cuba Astronomical Association, and by significant efforts by staff at NRAO's Green Bank and Charlottesville sites. I also would like to thank our collaborators at SKA-SA for ensuring the smooth running of PAPER, and for minimizing the number of (long) trips I've had to make to South Africa.

This research would not have happened without the support of the great PAPER team. My work is entirely inextricable from the enormous efforts everyone has put into designing an excellent instrument. In particular, I owe a huge debt of gratitude to my advisor Aaron Parsons, for not only steering the project through uncertain times, but also for making sure that my goals were always being met by the work I was doing. Thank you, Aaron, for always making me feel supported in my endeavors. I'd also like to thank Danny Jacobs for all his help over the years, answering countless questions over GChat that I was too embarrassed to ask anyone else. And, of course, we all owe our thanks to Don Backer, in whose memory all of our work is now dedicated. I am particularly grateful to Don for knowing exactly what to ask of (and what not to ask of) a new graduate student who, only months earlier, was on the brink of leaving graduate school. Without his help, I doubt I would be where I am today.

My friends, both inside and outside of astronomy, have meant so much to me. There have been countless escapades and adventures over the years; may there be many more. Within the astronomy department, I'd like to especially thank Chat Hull and Mike McCourt for rescuing me on a daily basis with coffee runs, and Peter Williams, not least of which for his ucastrothesis LaTeX template, and the countless hours of regulation cross-checking it has (hopefully) saved me.

No acknowledgments would be complete without thanking my family for their support for 28 years and counting. My parents' guidance has been beyond invaluable, and, whether intentionally or not, they have been some of the best role models I could dream of. My love to you both, and to my brother Jeremy, who I have been fortunate enough to have nearby in Berkeley the last several years.

And, finally, my thanks and love to my wife Claire. Beyond simply making coming home

and seeing you the highlight of my day, you have put up with all the joys of being married to a PhD student, when I am sure you could have happily never thought about graduate school again. And so, on to Seattle, and the next adventure!

Chapter 1

Introduction

Our understanding of cosmology – the history of the universe as a whole – has become remarkably complete in the past few decades. We have a concordance model for the time evolution of all the components of the universe: baryonic matter, radiation, (cold) dark matter (CDM), and dark energy (Λ). This ΛCDM cosmological model is an extraordinary achievement, in no small part because it is supported by a number of independent techniques at various periods in cosmic history, including: the production of light elements during Big Bang Nucleosynthesis ~ 3 minutes after the Big Bang (Beringer et al. 2012); the spectrum of fluctuations in the Cosmic Microwave Background when the universe was 380,000 years old (Larson et al. 2011; Planck Collaboration et al. 2013a) and the growth and distribution of structure (galaxies and galaxy clusters) in relatively recent times (Tegmark et al. 2004; Anderson et al. 2013). Despite its great success, the ΛCDM model poses obvious challenges to physicists and astronomers. Neither dark matter nor dark energy – which compose over 95% of the current energy density of the universe – are accounted for by the standard model of particle physics, nor is inflation, the preferred mechanism for generating the initial perturbations in an otherwise smooth universe. And, while the general picture of structure formation from gradual gravitational collapse of small initial perturbations accurately predicts the observed spatial distribution of galaxies, the detailed astrophysics of forming the galaxies we observe in the current epoch is still poorly understood. The work described in this dissertation has focused on developing and refining "21cm cosmology", a new technique with the potential to address many of these questions by probing a wide range of cosmic epochs.

1.1 Cosmic History

Before describing the techniques used in 21cm cosmology, it is worthwhile to first flush out our current picture of the universe in more detail, highlighting the gaps that remain to be filled. In this section, I divide the history of our universe into four epochs, each described sequentially below: the early universe, the "dark ages," the epoch of reionization, and the post-reionization or modern epoch. A more detailed overview of modern cosmology can be found in, e.g., Dodelson (2003).

1.1.1 The Early Universe

Although the furthest removed from our current epoch, our knowledge of the conditions in the early universe is exceedingly good, largely thanks to measurements of the light element abundances produced during primordial nucleosynthesis ~ 3 minutes after the Big Bang, and, especially, observations of the fluctuations in the cosmic microwave background 380,000 years after the Big Bang (or redshift $z \sim 1100$). In the standard model, these fluctuations are seeded by quantum fluctuations during a period of inflationary expansion, $\sim 10^{-30}$ seconds after the Big Bang (Dodelson 2003). Measurements of the temperature and polarization angular power spectra of the CMB are extremely well described by such a model. The current constraints from the CMB tend to favor the simplest models of inflation, which are in some ways the hardest to connect to current theories in particle physics (Planck Collaboration et al. 2013b). There are, however, several hints at exotic physics measured at $\sim 2-3\sigma$ by the Planck satellite, including a large scale "tilt" to the fluctuations, a low value for the measured quadrupole, and a overly-large "cold spot" (Planck Collaboration et al. 2013c). At the large scales in question, however, Planck is cosmic variance limited, so it is effectively impossible for CMB measurements alone to provide more insight into these phenomenon. Although they will be extremely challenging measurements to make, observations of 21cm emission from the Dark Ages (described in §1.2.3) may be the most promising tool for measuring more modes of the primordial power spectrum and reducing the cosmic variance limit.

1.1.2 The Dark Ages

The so-called cosmic "dark ages" begin after the formation of the CMB at $z \sim 1100$ and end with the formation of the first stars and galaxies at $z \sim 10-30$. They bear this name because there are no luminous sources in the universe during this period. (The one exception to this statement is the 21cm transition of neutral hydrogen, as discussed in §1.2.3). From a theoretical perspective, then, the dark ages are one of the best understood epochs because of their simplicity. Dark matter haloes continue to grow through gravitational collapse and accrete baryons freed from the photon-baryon fluid at last scattering. Eventually the structures go non-linear, and collapsed objects form and begin to emit photons, ending the dark ages.[1]

Although the dark ages are theoretically very well understood, the lack of luminous sources make them nearly inaccessible through observations. The James Webb Space Telescope (JWST) may be able to detect the first proto-galaxies from the end of the dark ages, but the bulk of this era may only be probable using highly redshifted 21cm observations. Because of the simple physics involved, the dark ages may be one of the best times to detect unexpected exotic physics, such as the signature of a decaying dark matter particle or deviations from the Gaussian perturbation spectrum predicted by inflation. I discuss the strengths of 21cm

[1] Tseliakhovich & Hirata (2010) first recognized that the high relative velocities between baryons and dark matter may complicate this picture; however, this effect seems to have little qualitative impact on the process of structure formation (O'Leary & McQuinn 2012; McQuinn & O'Leary 2012).

observations for making such measurements in §1.2.3.

1.1.3 The Epoch of Reionization

The Epoch of Reionization (EoR), which occurred roughly between $z \sim 6 - 13$, is one of the most active fields of research in astronomy today. The EoR — the "reionization" of the intergalactic medium (IGM) by the ultraviolet emission from early galaxies — represents the moment in cosmic history when structure formation has progressed far enough to profoundly impact the global state of the universe. In many ways, the EoR forms the bridge between the fundamental physics and cosmology of the early universe and the complicated astrophysics of galaxy formation. As such, measurements of the conditions during the EoR promise a wealth of information about the evolution of structure in the universe. Detailed reviews of the EoR can be found in Furlanetto et al. (2006) and Barkana & Loeb (2007).

The EoR also remains one of the last unexplored periods in the history of the universe, as traditional probes of cosmology at these redshifts — galaxy and quasar surveys — have been limited by the difficulties of NIR observing. Currently, $z \gtrsim 7$ galaxy surveys are only able to probe the brightest end of the galaxy population, while the interpretation of quasar spectra are limited by both their intrinsically low space-densities and the rapidly increasing optical depth of the intergalactic medium (IGM) as one pushes back into the reionization epoch. As such, many fundamental questions about the EoR remain: When did the universe reionize, and how long did it take? What were the sources that drove reionization? Were they early star-forming proto-galaxies (the theoretically preferred model), or did AGN/quasars play an important role? What role did metal-free Population III stars play relative to the next generation of stars (Pop II)? Did feedback from the ionizing background play an important role in the regulation of early galaxy formation? Much of what we do know about the EoR comes from large semi-analytic simulations (e.g. McQuinn et al. 2006, Trac & Cen 2007, Lidz et al. 2008, Iliev et al. 2012). These studies present a general qualitative picture of reionization driven by small star-forming galaxies, but any detailed predictions remain to be tested observationally.

Although these questions remain largely unanswered, a number of studies in recent years have offered tantalizing hints about the details of the reionization process. The two strongest constraints on reionization come from observations of the Gunn-Peterson effect in the spectra of high redshift quasars, and measurements of the optical depth in free electrons between us and the CMB, both of which I will describe in detail below. There are also numerous additional constraints which present some complexity in interpretation, but are worth briefly describing as well.

Observations of high-redshift quasars can reveal the presence of intervening neutral hydrogen along the line of sight, which will absorb cosmologically redshifted Lyman-α emission from the quasar. Quasars at redshifts above $z \gtrsim 5.7$ show a complete absence of transmission in their spectrum immediately redward of their intrinsic Lyα line (Fan et al. 2001, 2006) — the so-called Gunn-Peterson trough (Gunn & Peterson 1965). The rapid appearance of the Gunn-Peterson feature as one pushes to higher redshifts is interpreted

as a significant increase in the neutral hydrogen fraction of the universe above $z \sim 6$. In order to make this interpretation quantitative, however, one requires a model for the density distribution of the IGM, since neutral fractions of as little as 10^{-4} can result in complete absorption of Lyα photons. Therefore, these observations only probe the tail of the IGM density distribution. While most models are highly consistent with a rising neutral fraction above $z \sim 6$, McGreer et al. (2011) point out that a model free-constraint cannot rule out a reionization epoch which continues to $z \sim 5$.

Another strong constraint comes from measurements of the optical depth of free-electrons between us and the CMB. Since the CMB is essentially formed when the universe becomes neutral, any Thomson scattering of CMB photons will be from electrons released into the IGM during reionization. Since these scatterings occur relatively late in the history of the universe, they affect CMB power on large scales — roughly equivalent to the horizon size during reionization. In the temperature anisotropy angular power spectrum, these effects are largely degenerate with other effects. With measurements of the polarization power spectrum, many degeneracies can be broken. The best constraints from WMAP and other experiments yields a value of the optical depth $\tau = 0.086 \pm 0.013$ (Hinshaw et al. 2012). (The recent Planck results contain no polarization constraints, and so only marginally effect this measurement: $\tau = 0.089^{+0.012}_{-0.014}$ (Planck Collaboration et al. 2013a)). This constraint, however, is only an integral measurement, and leaves the ionization history relatively unconstrained. Interpreted as an (theoretically unmotivated) instantaneous reionization event, this value for the optical depth yields $z_{\rm reion} = 9.96$.

The combination of these two measurements provide a general picture for reionization: reionization began before $z = 10$ and was likely complete by $z = 6$. However, there have been a number of recent measurements which have the potential to tell us much more about the reionization process, but also have significant complications which make their interpretation less conclusive. They include:

- **Galaxy Counts** With the infrared-sensitive Wide Field Camera 3 (WFC3) installed on the Hubble space telescope in 2009, hundreds of $z \sim 6 - 9$ galaxies have been found through the Lyman break drop-out technique (see e.g. Finkelstein et al. (2012) and Bradley et al. (2012) for recent compilations). One can then ask if the observed galaxy population produces enough photons to reionize the universe. Unfortunately, answering this question requires both a model for converting broad-band photometric measurements into ionizing photon emissivities and a model for the distribution of dense neutral hydrogen clouds in the IGM. In general, only the brightest galaxies are observed, and the galaxy luminosity function must be extrapolated to significantly lower luminosities to produce enough ionizing photons (although see Finkelstein et al. (2012), who claim they observe enough galaxies at $z = 6$ to reionize the universe given reasonable models for the other unknown parameters).

- **Lyα Emitters** One can also study reionization by looking for Lyman-α emission from high-redshift galaxies (McQuinn et al. 2007). If the universe is sufficiently neutral, Lyα emission from these galaxies will be absorbed by the IGM. If the galaxies sit in large

enough ionized bubbles, then the Lyα photons can redshift out of resonance before being absorbed. With a large enough survey, one can study both the number density of these objects, and their clustering properties to learn about the state of the IGM. There are two methods currently being used to survey for Lyα emitters: searches using narrowband filters targeting emitters at specific redshifts (Tilvi et al. 2010; Ouchi et al. 2010; Hibon et al. 2011), and searches which spectroscopically target Lyman break galaxies (Pentericci et al. 2011; Ono et al. 2012; Schenker et al. 2012) to search for the presence of a Lyα emission line. In general, the narrowband photometric searches tend to find little evidence of redshift evolution in the Lyα-emitter luminosity function, while the spectroscopic searches see a significant decline the fraction of galaxies with visible Lyα emission between redshifts 6 and 7. If due entirely to absorption of Lyα photons by a neutral IGM, these observations are consistent with a neutral fraction of between 10 and 50% at $z = 7$, although it is difficult to disentangle the effects of potential evolution in the Lyα emitter population itself.

- **The Kinetic Sunyaev-Zel'dovich Effect** Measurements of small scale anisotropies in the CMB created by the Kinetic Sunyaev-Zel'dovich (KSZ) effect have the potential to constrain the duration of reionization. In particular, the "patchiness" of the reionization process creates a small spectral distortion when CMB photons scatter off moving electrons. The aggregate sum of these distortions adds small scale power to the CMB angular power spectrum which is proportional to the duration of the "patchy" period. Zahn et al. (2012) place the first constraints on this signal level, although the constraints on reionization are degenerate with a potential correlation between the thermal Sunyaev-Zel'dovich effect (which they must subtract out) and the cosmic infrared background (CIB). If there is no such correlation, they constrain the duration of reionization Δz to be ≤ 4.4; if they leave this correlation as a free parameter, the limit relaxes to $\Delta z \leq 7.9$. With future measurements from Planck, the South Pole or Atacama Cosmology Telescope, and Herschel, there are significant prospects for measuring this correlation and placing stringent constraints on the reionization epoch using this technique.

I have singled these three techniques out as those with significant potential to improve reionization constraints in the near future. There is a large range of additional techniques that have also been used to probe this epoch, including: direct measurements of the IGM temperature as a function of redshift (Bolton et al. 2012), measurements of the size of quasar "near-zones" (e.g. ionized bubbles around the quasar) at high redshift (Carilli et al. 2010; Bolton et al. 2011), measurements of the Lyα damping wing tails (to probe to smaller optical depths than in the saturated line centers) in high redshift gamma-ray bursts (Totani et al. 2006), and measurements of the Cosmic Infrared Background (CIB) to detect the integrated light from the reionizing galaxies directly (Kashlinsky et al. 2012; Cooray et al. 2012). Of these techniques, the quasar near zone measurements are suggestive of a significant ($\sim 10\%$) neutral fraction at $z \gtrsim 7$, while the measurements of the CIB still remain highly controversial. If anything, the broad range of techniques for probing reionization — and the broad range

of constraints — should illustrate both the difficulties in making measurements during the reionization epoch, but also the great interest in the EoR for improving our understanding of galaxy formation and cosmology. In §1.2.3, I will discuss the potential of 21cm cosmology to revolutionize our understanding of this epoch.

1.1.4 The Post-Reionization Epoch

In many ways, the epoch after reionization from $z \sim 0 - 6$ is one of the most complex, as a tremendous amount of detailed astrophysics takes place in the processes that turn small proto-galaxies into modern galaxies. This era is also the most accessible to the current generation of optical/NIR ground and space based telescopes, and in that sense, is fairly well understood. However one of the greatest mysteries in modern cosmology also first appears at $z \sim 2$: the accelerated expansion of the universe, presumably driven by some "dark energy". First measured by Riess et al. (1998) and Perlmutter et al. (1999), dark energy has now become a cornerstone of our standard cosmological model (Planck Collaboration et al. 2013a; Anderson et al. 2013). With essentially no theoretical model that can account for the observed strength of dark energy, current measurements are principally trying to answer whether the observed expansion is consistent with a cosmological constant.

One of the most exciting probes for revealing the nature of dark energy are the Baryon Acoustic Oscillation (BAO) features in the large-scale matter distribution. The BAO "wiggles" have recently drawn attention as a standard ruler by which the geometry of the universe can be directly measured (Eisenstein et al. 1998, 1999). These features in the cosmic microwave background (CMB) power spectrum and the matter power spectrum today are imprints from the acoustic oscillations in the primordial photon-baryon plasma that recombined at $z \approx 1100$. The features in the power spectrum appear at multiples of the sound horizon scale at recombination, making them effective standard rulers. Measuring the BAO wiggles at several redshifts yields geometric measurements of the universe — the Hubble parameter, $H(z)$, and the angular diameter distance, $d_A(z)$ — that constrain properties of the dark energy that dominates the cosmic energy content at $z = 0$ and is the current leading theory for the accelerated expansion of the universe. Since the first detection of the BAO signal (Eisenstein et al. 2005), several experiments have been undertaken to use these features for precision cosmology, including the Sloan Digital Sky Survey III Baryon Oscillation Spectroscopic Survey (SDSS-III BOSS; Schlegel et al. 2009)[2], WiggleZ (Drinkwater et al. 2010)[3], and the Hobby-Eberly Telescope Dark Energy Experiment (HETDEX; Adams et al. 2011)[4], as well as a number of planned future experiments, such as the Subaru Prime Focus Spectrograph (PFS; Ellis et al. 2012), Euclid (Amendola et al. 2012), BigBOSS (Schlegel et al. 2011), and the Wide-Field Infrared Survey Telescope (WFIRST)[5]. All of these experiments target individual galaxies with spectroscopic observations; in §1.2.3, I will discuss the prospects for a dark

[2] http://www.sdss3.org/surveys/boss.php/
[3] http://wigglez.swin.edu.au/site/
[4] http://hetdex.org/
[5] http://wfirst.gsfc.nasa.gov/

energy survey over unprecedented volumes of the universe with the redshifted 21cm line.

1.2 21cm Cosmology

In this section, I will very briefly review the fundamentals of observing the 21cm line at high-redshifts, often referred to as 21cm cosmology. Detailed reviews of this field can be found in Furlanetto et al. (2006), Morales & Wyithe (2010), and Pritchard & Loeb (2012), upon which much of this section is based.

1.2.1 21cm Emission at High Redshift

The so-called 21cm transition is the hyperfine spin-flip transition of neutral hydrogen; that is to say, the energy difference between the aligned and anti-aligned states of the proton and electron in a hydrogen atom is 5.9×10^{-6} eV, corresponding to a photon of wavelength 21 cm. It is common to parameterize the relative populations of these two states (n_1 and n_0) using a Boltzmann equation:

$$\frac{n_1}{n_0} = \left(\frac{g_1}{g_0}\right) e^{-T_*/T_S}, \tag{1.1}$$

where g_1 and g_0 are statistical weights, T_* is 0.068 K (the equivalent temperature of 5.9×10^{-6} eV), and T_S is the "spin temperature" which uniquely determines the occupation states. Although the lifetime of the excited state is exceedingly long ($\sim 10^7$ years), the high abundance of neutral hydrogen in the universe is more than enough to compensate for the long lifetime and make the transition observable. Moreover, the weakness of the line guarantees that it is always optically thin in all but the most extreme systems. Therefore, the principal idea of 21cm cosmology is relatively simple: observe all the neutral hydrogen in the universe. At cosmological distances, the observed frequency of the 21cm emission can easily be mapped into a distance, with the redshifting caused by cosmic expansion. The ultimate promise of 21cm cosmology is thus a 3-dimensional map of the distribution of neutral hydrogen in the universe.[6] In the remainder of this section, I will first discuss the evolution of the mean 21cm signal (the so-called "global" signal) and the physics governing the evolution of the spin temperature. I will then discuss the spatially varying components of the signal and how these terms evolve over the epochs outlined in §1.1. The focus of this discussion will be on the insights into fundamental physics and astrophysics that can be gleaned from these observations.

[6] In actuality, the observed frequency of the 21cm line does not perfectly map into the depth direction, but rather into "redshift space" where peculiar velocities of the emitting hydrogen also come into effect. For the first generation of 21cm experiments, redshift-space distortions have largely been ignored as a second-order effect. Ultimately, measuring redshift space distortions will offer rich science of its own, with the potential to probe the growth of structure and test theories of modified gravity.

1.2.2 The Global Signal

The fundamental observable of 21cm cosmology is the brightness temperature difference between neutral hydrogen and the CMB; in other words, if the spin temperature of the hydrogen is higher (lower) than the temperature of the CMB, the gas will emit (absorb) photons relative to the blackbody CMB spectrum at the 21cm frequency. Derivations of this brightness temperature difference can be found in the reviews cited at the beginning of this section. It is generally expressed as:

$$\delta T_b(z) \approx 28 x_{\rm HI}(1+\delta)\left[\frac{1+z}{10}\right]^{\frac{1}{2}}\left[1-\frac{T_{\rm CMB}(z)}{T_{\rm S}}\right]\left[\frac{H(z)/(1+z)}{{\rm dv}_\parallel/{\rm dr}_\parallel}\right]{\rm mK}, \tag{1.2}$$

where $x_{\rm HI}$ is the neutral hydrogen fraction, δ is a fractional overdensity relative to the cosmic mean, $T_{\rm CMB}(z)$ is the CMB temperature at redshift z, and the final term in brackets arises from redshift-space distortion effects, which will be neglected here. The global evolution of $\delta T_b(z)$ is directly connected to the global evolution of the spin temperature. In §1.2.3, I will discuss the parts of this equation that can spatially vary.

The global evolution of the spin temperature as a function of cosmic history is determined by the balance of different mechanisms which couple it to one of two temperatures in the universe: the CMB temperature and the kinetic temperature of the hydrogen gas. The redshift dependence of the signal arises from the effectiveness of the various coupling mechanisms as a function of cosmic time. Qualitatively, the evolution of the global signal is fairly straightforward. Up until $z \sim 200$ residual free electrons in the hydrogen are enough for Compton scattering of CMB photons to couple the gas temperature to the CMB temperature; since there is effectively only one temperature in the universe, the spin temperature also takes this value. Therefore, it is only below redshift ~ 200 that the 21cm signal becomes observable (recall from equation 1.2 that there is no signal when the spin temperature is equal to the CMB temperature).

At $z \sim 200$, the density of free electrons has dropped to a low enough level that Compton scattering is no longer an effective coupling mechanism. The gas temperature decouples from the CMB and begins to adiabatically cool at a rate proportional to $(1+z)^{-2}$ (compared with the CMB, which cools as $(1+z)^{-1}$); i.e., the kinetic gas temperature becomes lower than the CMB temperature. At this time, densities are still high enough that collisions between hydrogen atoms couple the spin temperature to the gas temperature. Therefore the 21cm signal can be seen in absorption relative to the CMB. However, the expansion of the universe continues to drive densities lower, and collisional coupling itself becomes ineffective around $z \sim 70$. The spin temperature therefore re-couples to the CMB temperature, and the 21cm signal effectively vanishes by $z \sim 30$. (Because the physics at this time is so well understood, measurements of the global 21cm signal during the dark ages are an extremely good effective thermometer. Any unusual heating signatures, e.g., the decay of dark matter or other relic particles can potentially be detected by these observations.)

Three different effects produce radical changes in the 21cm signal at $z \sim 20$, all related to the formation of the first luminous structures. First, the creation of a Lyman-α radiation field

re-couples the spin temperature to the gas temperature (which has been cooling adiabatically since $z \sim 200$), creating a strong absorption feature in the 21cm signal which in some models can reach values as high as -200 mK.[7] The second effect of luminous sources is to create X-ray radiation, which is effective at heating large volumes of the gas. Eventually, the temperature of the gas grows well in excess of the CMB temperature, leading to 21cm emission. The third effect of luminous sources is to emit ionizing photons, leading to the epoch of reionization, and destroying the 21cm signal from the IGM (although see 1.2.3 for a description of post-reionization 21cm emission). In fiducial models, these three effects proceed in order presented above: Lyman-α coupling creates an absorption signal, X-ray heating causes the signal to transition to emission, and, finally, ionizing photons destroy the signal.

It must be noted that there is significant uncertainty in these predictions, and one of the main goals of global 21cm experiments is to constrain the timing of these transitions. The underlying physics is clearly rich, giving such experiments the potential to inform us as to the formation times of the first stars and collapsed objects (and, equally, importantly, their spectral hardness, which can disentangle the roles of Pop II and Pop III stars, as well as the contributions of X-ray binaries). Finally, global 21cm experiments should be able to measure the disappearance of the 21cm signal coinciding with the reionization epoch although see §1.3.3 for a discussion of the inherent difficulties in all these measurements.

1.2.3 The Spatially Fluctuating Signal

The previous section detailed the evolution of the mean of the 21cm signal over cosmic history. However, there are terms in equation 1.2 that can fluctuate spatially with regard to the background value. Of particular interest are fluctuations in $x_{\rm HI}$, the neutral fraction of hydrogen, δ, the local fractional overdensity, and $T_{\rm S}$, the spin temperature. Which of these terms contribute to the overall observable signal is a function of cosmic history and is discussed in the sections below.

Ultimately, the clearest way to understand the spatial fluctuations in the 21cm signal is to image them. The first generation of experiments aiming to measure this signal, however, will not have adequate sensitivity (or, in most cases, resolution) to image individual fluctuations. Therefore, a statistical measure of the fluctuations — in particular, the power spectrum — is the principal observable for the immediate future. In what follows, the "21cm power spectrum" will be used roughly interchangeably with the "spatially fluctuating 21cm signal."

The Dark Ages

The dark ages present one of most exciting 21cm signals for fundamental cosmology (and, unfortunately, one of the least accessible; see §1.3.2). As stated above, there is no

[7] This Lyman-α coupling mechanism, known as the Wouthuysen-Field effect (Wouthuysen 1952; Field 1958), actually couples the spin temperature to the color temperature of the Lyman-α radiation field. The color temperature of the radiation is itself strongly coupled with the kinetic gas temperature (Furlanetto et al. 2006; Pritchard & Loeb 2012).

observable 21cm signal above $z \sim 200$, but once the gas temperature decouples from the CMB temperature, the 21cm signal becomes a very good tracer of matter density of the universe. There are no spatial fluctuations in the ionization fraction (the gas is all neutral) nor the spin temperature (the gas is all adiabatically cooling); rather the fluctuations traced by the 21cm signal are in δ, the matter density. In the standard cosmological picture, these matter fluctuations were seeded by primordial quantum fluctuations during inflation. Therefore, 21cm observations during the dark ages have the potential to constrain models of inflation by precisely measuring the primordial power spectrum. In fact, these observations have two significant advantages over the CMB power spectrum measurements described in §1.1.1. Firstly, photon diffusion damps small scale power in the CMB power spectrum. The actual matter itself clusters on much smaller scales, permitting accurate measurements of these modes in the power spectrum. Secondly, and perhaps more importantly, 21cm dark ages measurements probe a 3D volume, allowing many more samples of the largest scale modes, reducing the effects of cosmic variance. It may be that 21cm dark ages measurements will be necessary to conclusively identify the large scale CMB anomalies described in §1.1.1. As will be discussed in §1.3.2, the challenges for making these observations are even more substantial than 21cm observations at any other epoch; however, their scientific potential is unique enough that it should not be altogether ignored.

At this point, it should also be clear that there is no well-defined transition between the end of the dark ages and the beginning of the Epoch of Reionization. The formation of the first stars and galaxies leads to complex changes in the 21cm signal before significant ionization takes place in the IGM (§1.2.2). The rich astrophysics underlying this period make it one of interest, although interpretation of any signal will be complicated by the fact that significant spatial fluctuations likely occur in both δ and $T_{\rm S}$.

The Epoch of Reionization

Once significant ionized regions appear in the IGM, these fluctuations will dominate the observed power spectrum of 21cm emission, with peak amplitude occurring at 50% ionization (Furlanetto et al. 2006). In terms of absolute brightness, the 21cm power spectrum during the EoR is the strongest of the spatially fluctuating signals, peaking at $\sim 10 - 100$ mK^2. Relative to sky noise (§1.3.1) and foreground emission (§1.3.2), the EoR power spectrum also represents one of the most detectable 21cm signal, and has been the target of the most experimental efforts in the past decade.

The scientific impact of measuring and characterizing the 21cm signal from the EoR is significant, with the potential to answer most of the outstanding questions described at the beginning of §1.1.3. Some of the earliest observables will come from measuring the power spectrum evolution as a function of redshift. The redshift of peak power spectrum brightness corresponds to the time at which the universe was 50% ionized, and the rapidity of power spectrum evolution around this redshift can constrain the duration of the EoR. Measuring the amplitude of the power spectrum can also constrain the temperature contrast between the CMB and the 21cm spin temperature, thus measuring the effects of X-ray heating. Detailed

characterizations of the power spectrum can one day hope to constrain the clustering — and therefore, mass — of the reionizing galaxies, putting direct constraints on the first galaxies in the universe.

The Post-Reionization Epoch

After reionization, the power spectrum of 21cm fluctuations is expected to be a biased tracer of the matter power-spectrum, since the remaining neutral hydrogen resides in high-density, self-shielded regions such as in galaxies and other collapsed halos (Barkana & Loeb 2007; Madau et al. 1997). Rather than targeting individual objects, a 21cm experiment can detect fluctuations in neutral hydrogen emission on large scales (Chang et al. 2008; Wyithe et al. 2008; Morales & Wyithe 2010; Pritchard & Loeb 2012), with two dimensions corresponding to angles on the sky, and the third line-of-sight dimension arising from the differential redshifting of 21cm line emission as a function of distance. Because they smear out individual galaxies and only detect fluctuations on the largest scales, 21cm experiments at $z \sim 1-3$ are often referred to as "intensity mapping" experiments. Although this poor resolution makes 21cm experiments ill-suited for detailed galaxy evolution studies, the ability of intensity mapping to survey large volumes extremely quickly makes it a promising complement to spectroscopic galaxy surveys for constraining dark energy with baryon acoustic oscillations, as described in §1.1.4.

1.3 Challenges of Observing the 21cm Signal

In light of the great scientific impact achievable with observations of the highly redshifted 21cm signal, considerable effort has been spent in the last decade designing and commissioning experiments targeting its detection. The challenges for these experiments are sizeable, however. In this section I will describe three inter-related challenges for any experiment looking to measure the power spectrum of 21cm emission: the inherent faintness of the signal, the extreme brightness of foreground emission, and the unprecedented degree of calibration necessary to separate the two. All of these problems are also germane to experiments targeting the global 21cm signal, but the focus in this section will be on the measuring the spatially fluctuating signal during the Epoch of Reionization. A discussion of detecting the 21cm power spectrum at $z \sim 1$ can be found in Chapter 5.

1.3.1 The Faintness of the 21cm Signal

The brightness temperature of the 21cm signal was discussed in §1.2.2 and §1.2.3. The characteristic scale of global signal at $\sim$ 10 to 100 mK, although faint, is not a significant challenge in-and-of-itself. We can use the radiometer equation to estimate the amount of observing time needed for a single element to reach 10 mK sensitivity:

$$T_{\rm rms} = \frac{T_{\rm sky}}{\sqrt{Bt}} \tag{1.3}$$

Using fiducial values at $z = 10$ of bandwidth $B = 1$ MHz and $T_{\rm sky} = 1000$ K yields:

$$t_{\rm obs} = \left(\frac{T_{\rm sky}}{T_{\rm rms}}\right)^2 \times \frac{1}{B} = 10,000 \text{ s} \approx 3 \text{ hours}. \tag{1.4}$$

The challenge for observing the global signal, then, is not the faintness of the signal but the difficulty in separating it from the foregrounds, discussed below.

For experiments targeting the 21cm power spectrum, a characteristic brightness (at least at EoR frequencies) of $\Delta^2(k) \sim 10$ mK^2 sounds similar to that of the global signal. However, unlike the simple application of the radiometer equation to the global signal case, it is no trivial task to compute the power spectrum sensitivity of an interferometer. Such a derivation was performed in Morales (2005), McQuinn et al. (2006), Pen et al. (2009), and Parsons et al. (2012a). Parsons et al. (2012a) focused on incorporating array redundancy as a way to build sensitivity to the 21cm power spectrum, a technique now used by the PAPER experiment described in §1.5 and discussed in depth in Chapters 2, 3 and 4 of this dissertation. One of the main results of Parsons et al. (2012a) is their equation (27) which gives the thermal noise limits on a measured power spectrum:

$$\begin{aligned}\Delta^2_{\rm N}(k) \approx 60 &\left[\frac{k}{0.1\ h{\rm Mpc}^{-1}}\right]^{\frac{5}{2}} \left[\frac{6{\rm MHz}}{B}\right]^{\frac{1}{2}} \left[\frac{1}{\Delta{\rm lnk}}\right]^{\frac{1}{2}} \left[\frac{\Omega}{1.69 \text{ sr}}\right] \\ &\times \left[\frac{T_{\rm sys}}{500 \text{ K}}\right]^2 \left[\frac{6 \text{ hrs}}{t_{\rm per\_day}}\right]^{\frac{1}{2}} \left[\frac{120 \text{ days}}{t_{\rm days}}\right] \left[\frac{32}{N}\right] \left[\frac{10^4 f_0}{f}\right]^{\frac{1}{2}} {\rm mK}^2. \end{aligned} \tag{1.5}$$

The terms in this equation are defined as follows: B is the observing bandwidth, which has a fiducial value of 6 MHz as the maximum volume at $z \sim 10$ over which the universe can be considered co-eval (Furlanetto et al. 2006); Ω is the angular sized of the observed volume — note that it has been modified from the value given in (Parsons et al. 2012a) to reflect the new understanding given in footnote 4 of Chapter 4[8]; $T_{\rm sys}$ is the system temperature of the instrument, with a fiducial value of 500 K for PAPER; $T_{\rm per\_day}$ is the size of a foreground free cold-patch in hours of right ascension; $t_{\rm days}$ is the length of the observing campaign; N is the number of antennas in the array; and $\frac{f}{f_0}$ is a metric defined in Parsons et al. (2012a) which incorporates the redundancy of the baselines in the array, which has a fiducial value of $\sim 10^4$ for the configurations used by PAPER.

Although this is clearly a complicated relation (and admittedly somewhat non-transparent, given the redundancy metric), there are several useful lessons to be gleaned. Firstly, the large primary beam (i.e. small collecting area) of the PAPER instrument causes a significant hit in sensitivity; some more positive aspects of this element choice are discussed in §1.5. Secondly, it gives a sense of a fiducial EoR observation with a PAPER-like array: 120 days with 128

[8] The astute observer may note that the pre-factor of of 60 mK^2 also appeared in Parsons et al. (2012a). That publication, however, omitted a factor of 2 in the radiometer equation for a cross-correlation interferometer (as opposed to an auto-correlation total power system; c.f., Kraus 1966). This factor of 2 roughly offsets the factor of 2 error in beam area, leaving the bottom line essentially unchanged.

elements gets 1σ error-bars in the characteristic signal range. This result is the quantitative form of the statement that the signal is "faint." Long observations with large numbers of elements and/or large collecting areas will be necessary to detect the 21cm EoR power spectrum. See Chapter 5 for a derivation of equation 1.5 applicable for $z \sim 1$ experiments.

1.3.2 The Brightness of Foreground Emission

In contrast to the cosmological signal, foreground emission — primarily from our own Galaxy and extragalactic radio point sources — is extremely bright. The dominant emission process for these sources is synchrotron radiation, which leads to generally smooth, featureless power-law spectra:

$$T_b(\nu) = T_b(150\text{ MHz})\left(\frac{\nu}{150\text{ MHz}}\right)^{-\alpha}, \tag{1.6}$$

where the choice to normalize the power-law at 150 MHz is arbitrary. In general, Galactic synchrotron emission is well-described by a spectral index $\alpha = 2.5$, whereas extragalactic point sources have spectral indices closer to $\alpha \sim 1$, although with considerable spread around this value. At 150 MHz, Galactic synchrotron dominates the perceived emission; typical values of $T_b(150\text{ MHz})$ have already appeared in the text. For global experiments, $T_b(150\text{ MHz}) \sim 1000$ K; for interferometers targeting power spectrum measurements, the brightest emission can be resolved out, yielding measurements of $T_b(150\text{ MHz}) \sim 300$ K. These values lead to the often stated fact that foreground emission dominates the ~ 10 mK EoR signal by up to 5 orders of magnitude. Actual measurements of foregrounds in the power spectrum domain are presented in Chapter 4 of this dissertation. In Chapter 5, this calculation will be revisited for $z \sim 1$ experiments, concluding that the foreground-to-signal ratio is nearly identical to that at EoR frequencies. For dark ages measurements, the situation is far worse. Galactic synchrotron grows significantly brighter at low frequencies due to its steep spectral index, while the cosmological power spectrum is fainter than at $z \sim 10$; the rising brightness of Galactic emission also dominates the system temperature in equation 1.5, making these experiments starved for sensitivity as well. These statements form the basis for some of the pessimism regarding dark ages experiments in previous sections.

1.3.3 Calibration and Foreground Removal

Fortunately, the fact that foreground emission generally follows a smooth, featureless power law in frequency presents an extremely useful axis for separating this emission from the 21cm signal. (See Moore et al. 2013 for a discussion of Faraday rotated extragalactic synchrotron emission, a potentially problematic foreground that will not posses a smooth power law in frequency.) The redshift (i.e. frequency) evolution of the global 21cm signal was discussed in §1.2.2; although there will clearly be frequency structure in this signal, it is still relatively smooth. Separating this signal from the foreground emission on spectral smoothness alone will be a difficult, but potentially tractable, problem (Harker et al. 2012; Liu et al. 2013).

The fluctuating signal, however, will possess strong frequency structure which will be a powerful piece of information for discriminating the cosmological signal from foreground emission. As such, considerable effort has been devoted to developing schemes to remove or isolate foregrounds from 21cm power spectrum observations (e.g. Morales et al. 2006, Bowman et al. 2009, Liu et al. 2009, Liu & Tegmark 2011, Parsons et al. 2012b, Dillon et al. 2012). The simplest techniques involve subtracting polynomials in frequency to remove the spectrally smooth components, although considerably more detailed schemes have been developed in recent years to take better account of the effect of foreground subtraction on the cosmological signal. In §1.5.1, I will discuss the delay spectrum technique of Parsons et al. (2012b) in some more detail, as it forms the principal foreground removal tool used in Chapters 4 and 5 of this dissertation.

One of the most important aspects for employing these foreground removal techniques is to ensure that the instrument itself does not introduce spectral structure into otherwise spectrally-smooth foreground emission (sometimes referred to as "mode-mixing" in the literature). Foreground emission is bright enough that a corruption of the spectral-smoothness of foregrounds at even the fraction of a percent level can completely swamp the cosmological signal. Therefore, the utmost importance has been put on calibrating 21cm EoR experiments at an unprecedented level, so that instrumental effects can be corrected for and removed from the data. (The delay spectrum technique of Parsons et al. 2012b attempts to reduce the level of calibration necessary for a cosmological measurement, but even in this approach a good calibration is very important.) Many new techniques have been invented to calibrate 21cm experiments at the levels needed, including real-time calibration and imaging systems (Ord et al. 2010) and pulsar-gating (Pen et al. 2009). I will report on two new calibration techniques in Chapters 2 and 3 of this dissertation.

1.4 21cm Experiments

In this section, I will acknowledge the wide range of approaches towards tackling the obstacles associated with 21cm cosmology at all redshifts, before describing in the next section the PAPER experiment on which Chapters 2-4 of this dissertation are based. There is little room to do justice the considerable efforts that have been invested in the last decade towards developing new experiments; however, it is worthwhile to briefly consider the diversity within what is sometimes considered a niche field.

There are, in general, fewer global 21cm experiments than those targeting a power spectrum measurement. Because sensitivity is less of an issue for a global experiment, the first generation of instruments have been small, single dipoles, such as the Experiment to Detect the Global EoR Step (EDGES; Bowman et al. 2008)[9], which released the first 21cm constraints on the EoR in 2010, ruling out extremely rapid reionization with $\Delta z < 0.06$ (Bowman & Rogers 2010). A new generation of global 21cm experiments are targeting the 21cm absorption signal from the first stars at $z \sim 20$. The Large Aperture-experiment to

[9] http://www.haystack.mit.edu/ast/arrays/Edges/

Detect the Dark Ages (LEDA)[10] experiment on the Long Wavelength Array (LWA; Ellingson et al. 2009))[11] is looking to improve on the calibration of the total power dipoles by using interferometric cross-correlations with other elements in the LWA. Whereas the Dark Ages Radio Explorer (DARE; Burns et al. 2012)[12] will orbit the moon, observing only when behind the far side, to avoid both radio frequency interference (RFI) from Earth and the Earth's ionosphere, which can introduce significant fluctuations into radio signals at sub-100 MHz frequencies.

Numerous facilities and experiments targeting the signal from the EoR are already online or under construction, including the Giant Metre-Wave Radio Telescope (GMRT; Swarup et al. 1991)[13], the LOw Frequency ARray (LOFAR; Yatawatta et al. 2013)[14], the Murchison Widefield Array (MWA; Tingay et al. 2012)[15], and the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010)[16]. The GMRT is an all-purpose low-frequency radio observatory in Pune, India. As such, it has considerable collecting area (thirty 40 m dishes), but lacks many of the other particular adaptations developed for 21cm experiments. LOFAR is also a multi-purpose radio facility located in the Netherlands (although long baselines span much of Europe). Rather than single dishes, LOFAR employs phased-array tiles of dipoles, giving it an extremely wide field of view. The MWA is an EoR experiment located in the radio quiet SKA reserve in Western Australia; as of the writing of this dissertation, the MWA has deployed 128 tiles of 16-dipoles each. Early efforts with the MWA focused on developing a real-time imaging and calibration system, although more recent results have focused on developing observing strategies for measuring the EOR power spectrum (Dillon et al. 2013). PAPER will be described below in §1.5.

Although somewhat less attention has been paid to $z \sim 1-3$ intensity mapping, there has been considerable effort towards designing new experiments in the last few years, in part thanks to the recent detections of the 21cm signal through cross-correlations of Green Bank Telescope (GBT) observations with the DEEP-II and WiggleZ galaxy catalogs (Chang et al. 2010; Masui et al. 2013; see Switzer et al. 2013 for a constraint on the 21cm signal using GBT observations alone). Several dedicated 21cm intensity mapping experiments have been proposed, including the Canadian Hydrogen Intensity Mapping Experiment (CHIME)[17], based off the prototype Cylindrical Radio Telescope (CRT; Peterson et al. 2006; Seo et al. 2010)[18] and the BAO from Integrated Neutral Gas Observations experiment (BINGO; Battye et al. 2012). In Chapter 5, I will present the BAO Broadband and Broad-beam Array (BAOBAB), a design for a new $z \sim 1$ 21cm intensity mapping experiment.

[10] http://www.cfa.harvard.edu/LEDA/
[11] http://www.phys.unm.edu/~lwa/
[12] http://lunar.colorado.edu/dare/
[13] http://gmrt.ncra.tifr.res.in/
[14] http://www.lofar.org/
[15] http://www.mwatelescope.org/
[16] http://eor.berkeley.edu/
[17] http://www.mcgillcosmology.ca/chime
[18] http://cmb.physics.wisc.edu/people/lewis/webpage/

1.5 The Precision Array for Probing the Epoch of Reionization

Most of the work in this dissertation will focus on calibrating and observing with the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER). A fuller discussion of the PAPER architecture can be found in Parsons et al. (2010). This section will briefly review the salient details of the PAPER system design, before discussing some of the key elements of the PAPER "approach" to detecting the 21cm EoR signal.

PAPER is an interferometric array of single, dual-polarization dipole elements. There are two PAPER sites: one in Green Bank, West Virginia, and one in the radio quiet SKA-SA reserve in the Karoo desert of South Africa. Green Bank is primarily an engineering testbed site, whereas the pristine RFI environment in South Africa makes it our science site. 32 dipoles are current deployed at the Green Bank site; South Africa currently has 64 dipoles, with a planned upgrade to 128 in summer 2013. These single dipoles are a hallmark of the PAPER design. Although they are somewhat limited in collecting area compared with other 21cm experiments, they have been extensively engineered to reduce both spatial and frequency structure in their responses, allowing for an extremely clean separation of foregrounds from the cosmological signal. The PAPER system possesses a large instantaneous bandwidth spanning $100 - 200$ MHz. As will be discussed in Chapter 4, this wide bandwidth has proven invaluable for foreground removal inside of a smaller, cosmological band. Further details of the PAPER system will be discussed when relevant in subsequent chapters.

1.5.1 The PAPER Approach

Two unique features of the PAPER approach are its use of "maximum-redundancy" arrays for building sensitivity, and the "delay-spectrum" technique for isolating foreground emission, published in in Parsons et al. (2012a,b), respectively. It is worthwhile to briefly summarize the results of these two papers, as they will be referenced extensively in the remaining chapters.

In Parsons et al. (2012a), we developed a formalism for computing the sensitivity of various array configurations to the 21cm power spectrum. We divided arrays configurations into two main classes: traditional interferometric configurations designed to sample as many distinct Fourier modes of the sky as possible — hereafter, "minimum redundancy" arrays — and arrays with multiple copies of the same baseline designed to achieve high SNR measurements of those particular Fourier modes — hereafter, "maximum redundancy" arrays (see Figure 1.1). Because of the modest size of its dipole elements, PAPER has been able to reconfigure and utilize both types of array configurations, to achieve both high-fidelity images characterizing foreground emission and high-sensitivity observations for putting constraints on the 21cm power spectrum.

In our companion paper, Parsons et al. (2012b), we combined the sensitivity calculations of Parsons et al. (2012a) with a formalism for mapping an understanding of foreground emission in interferometric "delay space" into predictions for the effects of foreground contamination on recovering the 21cm power spectrum. The conclusions were that foregrounds contaminate

Figure 1.1: *Left*: Aerial photograph of a 32-element PAPER array deployed in a minimum redundancy configuration in May 2010. *Right*: Photograph of a 64-element PAPER array deployed in a maximum redundancy configuration in September 2012.

cosmological Fourier space (k-space) in a baseline-length dependent manner — shorter baselines can access more of the 21cm power spectrum, free from contamination by foregrounds. Significantly more detailed discussions of the delay spectrum technique appear in both Chapter 4 — where it is applied to actual data for the first time — and Chapter 5 — where it it used to make predictions for the effect of foreground on $z \sim 1$ intensity mapping observations.

1.6 Outline of This Dissertation

The outline of this dissertation is as follows. In Chapters 2 and 3, I will present the results of two new calibration techniques I have designed and applied to observations with PAPER. Chapter 2 studies the effect of temperature-dependent gains on PAPER observations. Chapter 3 designs a new technique for recovering a wide-field primary beam model when both the beam and the sky are unknown, and applies this technique to recover the PAPER primary beam. In Chapter 4, I use the delay-spectrum technique to analyze observations from a 64-element PAPER South Africa array to measure — for the first time — the power spectrum properties of foreground emission in the EoR band. In Chapter 5, I will present a design and cosmological forecasts for a $z \sim 1$ 21cm intensity mapping experiment — BAOBAB — which borrows design elements from both PAPER and the MWA. Finally, I will conclude in Chapter 6.

Chapter 2

Temperature Dependent Gains in the PAPER System

This chapter describes the testing of the temperature dependence of the PAPER system gain using astronomical observations. This temperature dependence had been measured in previous laboratory experiments, but correction schemes has received only minimal testing in real data. I use observations of Cygnus A from the PAPER Green Bank array to independently derive the magnitude of the temperature dependent effects and compare with laboratory measurements. The overall effect is well described using all methods, although attempts to separate the effects of various system components are more difficult. Finally, I describe the scheme ultimately used to correct for these effects in the PAPER South Africa deployments.

2.1 Introduction

Given the faintness of the 21cm signal, instruments — especially those like PAPER with relatively small collecting areas — will need to undertake very long integrations to build sensitivity. A fiducial PAPER observing campaign will span 4-5 months (Parsons et al. 2012a). Ensuring system stability over these long time periods is essential for properly combining data from throughout the campaign. One known source of such instability are temperature dependent gain variations in the PAPER balun/psuedo-differential amplifier (PDA) as well as in the above-ground co-axial cables connecting the antennas to the correlator container (Parashare & Bradley 2009; Parsons et al. 2010). If uncorrected, these gain variations can break the symmetry around zenith within a single day (i.e., the gain of the system when observing an EoR field when 30° east of zenith will likely be different than when the same field is 30° west of zenith, due to heating over the course of the day). Potentially even more problematic, the PAPER observing seasons spans runs from September to April; given the intense heat in the summers of the Karoo, temperature dependent gains can introduce significant fluctuations between different days of the observing campaign.

The purpose of the investigations in this chapter is two-fold: first, we want to use

astronomical observations to test the validity of the laboratory-measured fitting formulae for correcting these temperature dependent effects, and, second, we want to determine the design of a system for temperature monitoring that can correct for these effects in our South African array.

2.2 Laboratory Measurements

In Parashare & Bradley (2009), the gains of various components of the PAPER system as a function of temperature were measured in the laboratory. The balun/PDA was found to have a temperature dependent gain with with a coefficient of -0.024 dB K^{-1}. A similar coefficient of -0.018 dB K^{-1} was also found for the RG-6 coaxial cables used to connect the PDAs to the receiver cards.[1] Further investigations have refined these models into 3^{rd} and 4^{th} order polynomials, the exact form of which are:

$$g_{balun} = 30.3573 - 0.02485T_b + 0.00010256T_b^2 - 0.000001979T_b^3 \tag{2.1}$$

$$g_{cable} = 15\times(-0.72167-0.0032929T_c+0.000078251T_c^2-0.0000013723T_c^3+0.0000000098601T_c^4) \tag{2.2}$$

Note that the temperatures in these fits are in Celsius. In order to conduct field studies of the temperature dependence, a system for actively monitoring the temperature of a representative PDA and cable with a thermister was deployed in Green Bank. This system forms the basis of the measurements used in this analysis.

It should also be noted that past analyses on PAPER data have shown the gain corrections using these temperature measurements lower the variance in the total power levels seen by the PAPER system between integrations (Parsons et al. 2010). This fact has been confirmed using both sky data and data from a "gain-o-meter": a balun terminated with a matched load, rather than the dipole arms of the antenna. However, there has been no study of the effects of the temperature dependent corrections on the perceived flux densities from single, specific celestial sources. To produce a reliable measurements of source flux densities, and, ultimately, to properly measure the brightness scale of the EoR signal, it will be necessary to correct for these temperature dependent gains.

2.3 Analysis

12 days of data from a 16 element array in Green Bank, West Virginia were taken in summer 2009 were used for this analysis. During the campaign, 3 of the antennas were not

[1] The receiver cards used to amplify the signal after travelling the length of the cable were found to have the strongest temperature dependent gain; however, the receiver cards are kept in a climate controlled room, drastically reducing the amount of temperature fluctuations. In the work here, we assume the receiver cards maintain a stable temperature.

functioning, while a 4th served as a “gain-o-meter," leaving 12 operating antennas in the array. (The thermisters for temperature monitoring were connected to the gain-o-meter.) The data were reduced using AIPY with relatively minimal prodcedures. First, the gain linearization and RFI excision algorithms described in Parsons et al. (2010) were applied to the raw data.

To extract the perceived flux densities of individual sources, an Markov Chain Monte Carlo (MCMC) application of delay/delay rate (DDR) filters was applied (Parsons & Backer 2009). This MCMC implementation differs from that described in Parsons & Backer (2009), in that it fits all sources simultaneously, using MCMC to explore parameter space and isolate each source from the effects of sidelobes from the others. The output of the algorithm is a model of each source's perceived flux density as a function of time and frequency; the final model is that which accounts for the majority of the flux density in the raw visibilities. These filters also reconstuct source data with antenna-based solutions, allowing each antenna's gain to vary independently.

In this analysis, we extracted the brightest sources in the northern hemisphere: the Sun, Cygnus A, Casseopeia A, the Crab Nebula and Virgo A. These 5 sources account for the vast majority of the flux density in the observations. Two hours data from each day near the transit of Cyg A were selected for detailed analysis, over 480 frequency channels ranging from 123.4 to 170.2 MHz, a relatively clean window in Green Bank. The magnitude of each reconstructed visibility was corrected for the modelled response of the PAPER primary beam to Cyg A at that time and freqeuncy, using the model from Parsons et al. (2010). The median value across the band was then taken in order to compress each 480-channel integration into a single data point. All visibilities from each antenna were then divided by their median in order to reduce the effects of antenna-to-antenna gain differences. The resulting data are plotted in Figure 2.1 vs. the measured cable/PDA temperature. The y-axis is a ratio of the perceived flux with the predicted flux from the beam model. The slope of the best fit line is -0.047 dB K^{-1} and -0.036 dB K^{-1} when plotted against the balun and cable temperatures, respectively. Issues concering the uncertainties in these measured values are addressed in Section 2.4 below.

Although these fits are illustrative of the general trend, there is, of course, strong covariance between the PDA and cable temperatures, as they are both exposed to the same ambient temperature. The perceived power cannot be fit for each gain coefficient individually. Fitting for both parameters simulataneously yields values of -0.030 dB K^{-1} and -0.013 dB K^{-1} for the PDA and cable coefficients, respectively. These values are still highly covariant, however. Without an estimate for off-diagonal terms in the covariance amatrix, standard error calculations would underestimate the uncertainty in these figures. Instead, attempts to assess the significance of these figures compared with the models are presented in Section 2.4.

Running the same analysis with data near the transit of Cas A yields coefficients of -0.078 dB K^{-1} for the PDA and +0.036 dB K^{-1} for the cable. This result further illustrates the degeneracy between these two measurements. The positive slope for the cable gain is in direct contradiction with all laboratory measurements. However, the sum of the two coefficients yields a *total* temperature dependent gain that is consistent with all previous

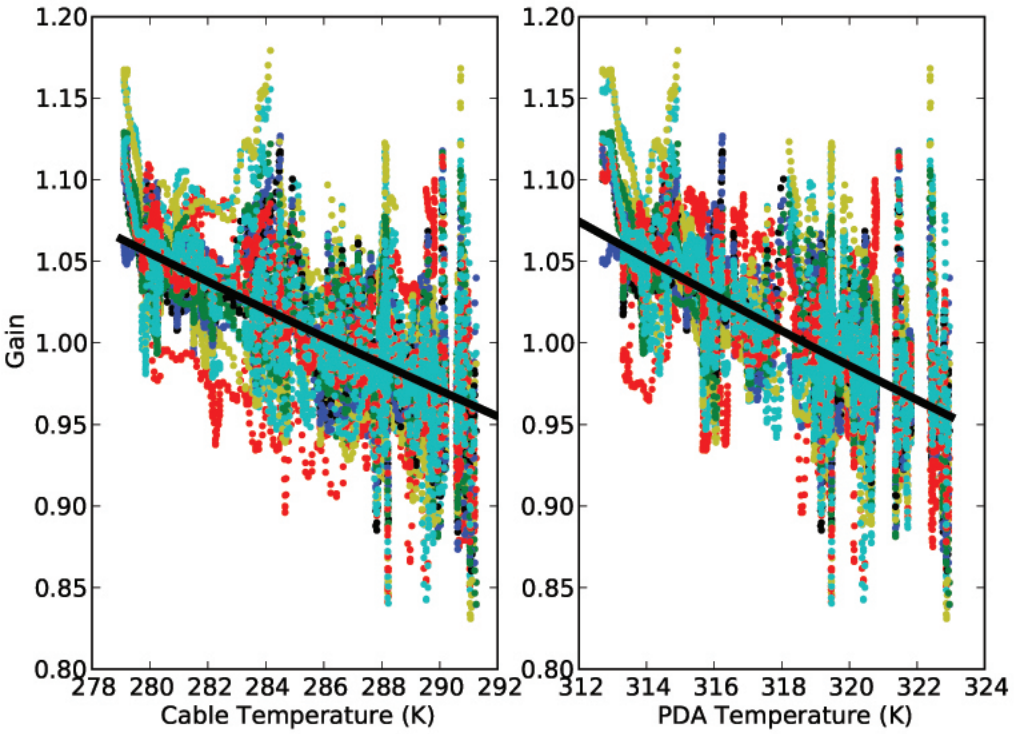

Figure 2.1: Temperature dependendent gain effects vs. a single measured temperature (PDA or cable). Each color represents data from one antenna.

	Balun	Cable	Residuals
Fit	-0.030	-0.013	0.001207
Linear Model	-0.024	-0.018	0.001208
Polynomial Model	—	—	0.001333

Table 2.1: Residuals from fit and both models. Residuals are calculated using Equation 2.3.

analyses. It should also be noted that the scatter around the fit is significantly larger for Cas A than Cyg A, and so the result from the Cyg A fit will be taken as the best estimate from this technique.

2.4 Significance

In this section, we will compare the fit presented above to the two laboratory derived models. One model assumes a linear relationship between temperature and gain, and has been the model used to correct these effects in data analysis up to now. As mentioned above, this model uses coefficients of $H_{balun} = -0.024$dB K^{-1} and $H_{cable} = -0.018$dB K^{-1}. The other model uses 3rd and 4th order polynomial fits to represent the gain flucatuations of the PDA and cable, respectively. These formulae are given in Equations 2.1 and 2.2.

The fit derived from the approach outlined above is compared with both models in Table 2.1. The "residuals" are the sum of the squares of the residuals, given by:

$$\chi^2 = \frac{\sum_i (V_i - m_i)^2}{\nu}, \tag{2.3}$$

where V_i is a reconstructed visibility with an associated T_{cable} and T_{balun}, m_i is the model or fit prediction for that same T_{cable} and T_{balun}, and ν is the number of degrees of freedom in the fit. It should be noted that χ^2 as defined is not a true χ^2, as no estimate of the variance in the measurements has been used; instead we have only divided by the number of degrees of freedom.

Since these residuals are not true χ^2, additional analysis is necessary to assess the significance of the differences in residuals. In order to get a characteristic scale for the reported residuals, we fit the null hypothesis, i.e., that there are no temperature dependent gains. This yields a χ^2 of 0.002077. The fit and linear model are therefore both 42% improvements over the null model, whereas the polynomial fit is a 36% improvement. The idea that the linear models are favored over the polynomial fit with statistical signifance is supported by the data shown in Figure 2.2. This figure shows the *total* gain change for a given (T_{cable}, T_{balun}) pair; on the x-axis is the gain change predicted by the fit, whereas the y-axis is the gain change predicted by the two models. The fit appears indistinguishable from the linear model. The polynomial model, however, appears to overpredict the amount of gain change at a given temperature.

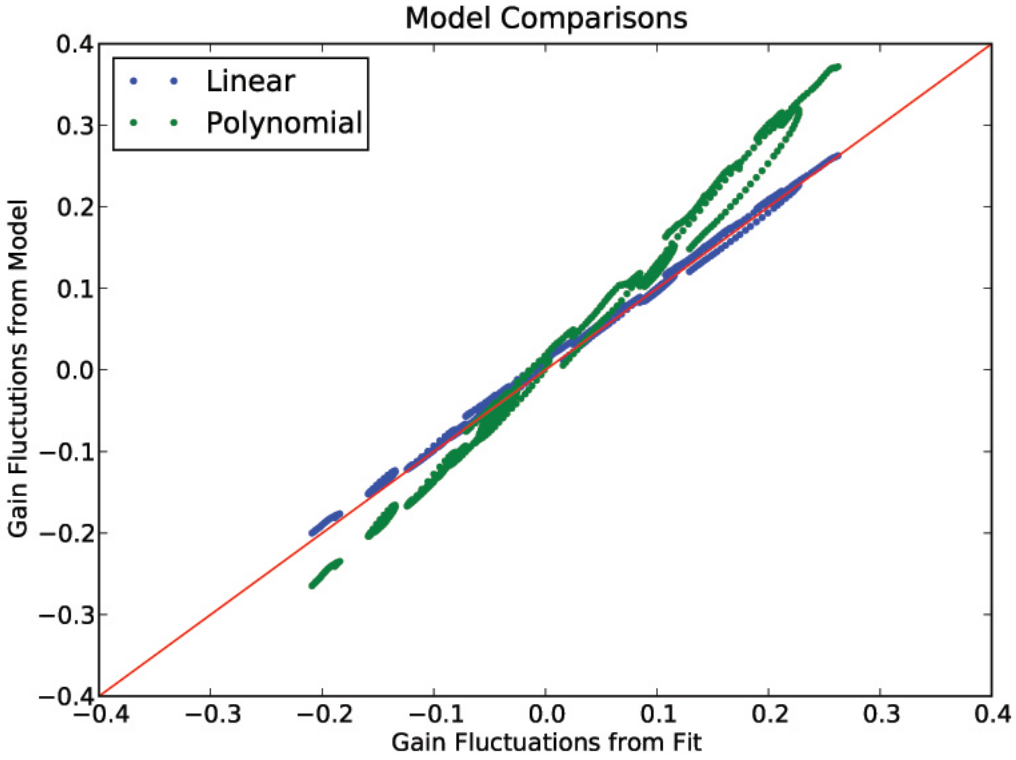

Figure 2.2: Comparison of total gain predictions: linear and polynomial models versus the fit from §4.3. The red line is $y = x$, i.e., it represents indistinguishable models.

2.5 Discussion

The conclusion suggested by this analysis is that the linear model of $H_{balun} = -0.024\text{dB K}^{-1}$ and $H_{cable} = -0.018\text{dB K}^{-1}$ should continue being used to correct the data. Although this model appears indistinguishable from the fit produced by this analysis, it is also derived from laboratory measurements, and therefore has more credibility. One should also be cautious of the results of the fit due to the strong degeneracy between the two coefficients, as can be seen from analysis of Cas A.

It is worth offering some speculation about the polynomial model, as it is certainly the product of careful and concerted efforts. Perhaps it actually is the more accurate model when the true balun and cable temperatures are known, as they were in these laboratory experiments. However, data taken in the field uses values reported by thermisters near the balun and cable as proxies for the balun and cable temperatures themselves. If the balun or cable were hotter than the value reported by the thermister (a realistic possibility for the balun, which produces its own heat), then this model would *underpredict* the gain decrement. This hypothesis therefore predicts the opposite result as is seen in the data; at the present time, there is still no compelling explanation for the failure of the polynomial model.

2.6 Implementation in South Africa

The temperature monitoring system described in §2.2 poses challenges for deployment in South Africa. The Green Bank system consists of a specially design balun/PDA and cable with built in thermisters; these themisters are then monitored with a dedicated PC using a specialized card PCI card too large for a laptop machine. It is clear from the measurements described above, however, that measuring one temperature and using a gain correction coefficient of $H_{\text{system}} = -0.042\text{dB K}^{-1}$ (the sum of the cable and PDA gain coefficicents) accounts for nearly all of the variance in the data. For the South African site, we therefore use a small LabJack[2] system to monitor a thermister on one of the dipoles in the array. Since 2010, this single ambient temperature measurement has been collected throughout the day, and is used to correct the visibilities from the array.

Acknowledgements

A version of this chapter originally appeared as PAPER Memo #16, available at http://eor.berkeley.edu/.

[2]http://labjack.com

Chapter 3

A Technique for Primary Beam Calibration of Drift-Scanning, Wide-Field Antenna Elements

In this chapter, I present a new technique for calibrating the primary beam of a wide-field, drift-scanning antenna element. Drift-scan observing is not compatible with standard beam calibration routines, and the situation is further complicated by difficult-to-parametrize beam shapes and, at low frequencies, the sparsity of accurate source spectra to use as calibrators. I overcome these challenges by building up an interrelated network of source "crossing points" — locations where the primary beam is sampled by multiple sources. Using the single assumption that a beam has 180° rotational symmetry, one can achieve significant beam coverage with only a few tens of sources. The resulting network of crossing points allows one to solve for both a beam model and source flux densities referenced to a single calibrator source, circumventing the need for a large sample of well-characterized calibrators. I illustrate the method with actual and simulated observations from the Precision Array for Probing the Epoch of Reionization (PAPER).

3.1 Introduction

The past decade has seen a renewed interest in low frequency radio astronomy with a strong focus on cosmology with the highly redshifted 21cm line of neutral hydrogen. As discussed in Chapter 1, numerous facilities and experiments are already online or under construction, including the Giant Metre-Wave Radio Telescope (GMRT; Swarup et al. 1991)[1], the LOw Frequency ARray (LOFAR; Röttgering 2003)[2], the Long Wavelength Array (LWA; Ellingson et al. 2009)[3] and the associated Large Aperture-experiment to Detect the Dark

[1] http://gmrt.ncra.tifr.res.in/
[2] http://www.lofar.org/
[3] http://www.phys.unm.edu/~lwa/index.html

Ages (LEDA) experiment, the Cylindrical Radio Telescope (CRT/BAORadio, formerly HSHS, Peterson et al. 2006; Seo et al. 2010)[4], the Experiment to Detect the Global EoR Step (EDGES; Bowman et al. 2008)[5], the Murchison Widefield Array (MWA; Lonsdale et al. 2009)[6], and the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010)[7]. 21cm cosmology experiments will need to separate bright galactic and extragalactic foregrounds from the neutral hydrogen signal, which can be fainter by as much as 5 orders of magnitude or more (see, e.g., Furlanetto et al. 2006 and Santos et al. 2005). As such, an unprecedented level of instrumental calibration will be necessary for the detection and characterization of the 21cm signal.

Achieving this level of calibration accuracy is complicated by the design choice of many experiments to employ non-tracking antenna elements (e.g. LWA, MWA, LOFAR and PAPER). Non-tracking elements can provide significant reductions in cost compared to traditional dishes, while also offering increased system stability and smooth beam responses. However, non-tracking elements also present many calibration challenges beyond those of traditional radio telescope dishes. Most prominently, the usual approach towards primary beam calibration — pointing at and dithering across a well-characterized calibrator source — is not possible. Instead, each calibrator can only be used to characterize the small portion of the primary beam it traces out as it passes overhead. Additionally, the wide fields of view of many elements make it non-trivial to extract individual calibrator sources from the data (see, e.g., Parsons & Backer 2009 and Paciga et al. 2011 for approaches to isolate calibrator sources). Finally, many of these arrays use dipole and tile elements, the response of which are not easily described by simple analytic functions.

In this chapter, I present a method for calibrating the primary beams of non-tracking, wide-field antenna elements using astronomical sources. I illustrate the technique using both simulated and observed PAPER data from a 12 antenna array at the NRAO site in Green Bank, WV. As discussed in Chapter 1, PAPER is an interferometer operating between 100 and 200 MHz, targeted towards the highly-redshifted 21cm signal from the epoch of reionization. Although the cosmological signal comes from every direction on the sky, an accurate primary beam model will be necessary to separate the faint signal from bright foreground emission. In this work, I use a subset of the brightest extragalactic sources to calibrate the primary beam. Because interferometers like PAPER are insensitive to smooth emission on large scales, these extragalactic sources are easily detectable despite the strongly increasing brightness of Galactic synchrotron emission at low radio frequencies. Only the measured relative flux densities of each source are needed to create a beam model, but to facilitate comparision with other catalogs, we use the absolute spectrum of Cygnus A from Baars et al. (1977) to place our source measurements on an absolute scale.

The structure of this chapter is as follows: in §3.2 I motivate the problem and the need for a new approach to primary beam calibration for wide-field, drift-scanning elements. In

[4] http://cmb.physics.wisc.edu/people/lewis/webpage/index.html
[5] http://www.haystack.mit.edu/ast/arrays/Edges/
[6] http://www.mwatelescope.org/
[7] http://eor.berkeley.edu/

§3.3 I present our technique for primary beam calibration. I show the results of applying the method to simulated and actual observations in §3.4 and §3.5, respectively, and I conclude in §3.6.

3.2 Motivation

For non-tracking arrays with static pointings such as PAPER, every celestial source traces out a repeated "source track" across the sky, and across the beam, each sidereal day. The basic relationship between the perceived source flux density (which we shall call a measurement, M) measured at time t and the intrinsic source flux density (f) is:

$$M(t) = b(\hat{s}(t))f, \tag{3.1}$$

where b is the response of the primary beam toward the time-dependent source location $\hat{s}$.

If the inherent flux density of each source were well-known, it would be straightforward to divide each $M(t)$ by f to obtain b along the source track $\hat{s}(t)$. To form a complete beam model, one would then need enough well-characterized sources to cover the entire sky. In the 100-200 MHz band, however, catalog accuracy for most sources is lagging behind the need for precise beam calibration (Vollmer et al. 2005; Jacobs et al. 2011), which in turn is necessary for generating improved catalogs. Without accurate source flux densities, both b and f are unknowns, and Equation 3.1 is underconstrained. The problem becomes tractable if several sources pass through the same location on the sky, and therefore are attenuated by the same primary beam response. However, the density of bright sources at low frequencies is insufficient to relate sources at different declinations. Additional information is necessary to break the degeneracy between primary beam response and the inherent flux densities of sources.

3.3 Methods

One way to overcome the beam-response/flux-density degeneracy described above is to assume 180° rotational symmetry in the beam. Under this assumption, each source creates two source tracks across the sky: one corresponding to the actual position of the source, and the other mirrored across the beam center. Under this symmetry assumption, tracks overlap at "crossing points," as schematically illustrated in Figure 3.1 for two sources: Cygnus A and Virgo A. At a crossing point, there are only 3 unknowns, since each source is illuminated by the same primary beam response, b. Since two circles on the sky cross twice (if they cross at all), there will be two independent relations that together provide enough information to constrain both the primary beam response at those points and the flux densities of sources passing through them (relative to an absolute flux scale). By using multiple sources, it is possible to build a network of crossing points that covers a large fraction of the primary beam. Furthermore, such a network allows data to be calibrated to one well-measured fiducial

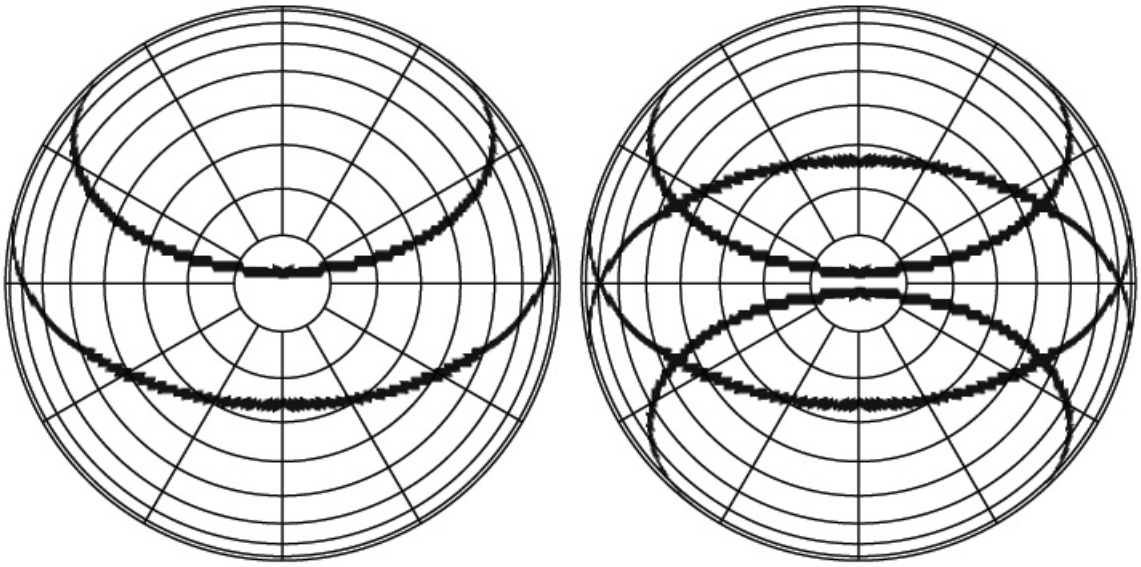

Figure 3.1: The building up of source tracks and crossing points. *Left:* The paths Cygnus A (top) and Virgo A (bottom) take across the beam for the PAPER array in Green Bank. The projection is orthographic with zenith at the center; dotted lines are 10° and 30° steps in elevation and azimuth, respectively. *Right:* The tracks of Cygnus A and Virgo A, and their 180° rotated images overlaid. There are now 6 crossing points. These are the locations where the beam and source parameters can be solved for independently.

calibrator. For observations with the Green Bank deployment of PAPER, source fluxes are related to Cygnus A from Baars et al. (1977).

In the rest of this section, I discuss the algorithmic details of how this approach to primary beam calibration is implemented. These steps are:

1. extracting measurements of perceived source flux densities versus time from observations (§3.3.1),
2. entering measurements into a gridded sky and finding crossing points (§3.3.2),
3. solving a least-squares matrix formulation of the problem (§3.3.3), and, finally,
4. deconvolving the irregularly sampled sky to create a smooth beam (§3.3.4).

I also discuss prior information that can be included to refine the beam model in §3.3.5.

3.3.1 Obtaining Perceived Source Flux Densities

The principal data needed for this approach are measurements of perceived flux density versus time for multiple calibrator sources as they drift through the primary beam. In practice, any method of extracting perceived individual source flux densities (such as image plane analysis) could be used; the beam calibration procedure is agnostic as to the source of these measurements. In this chapter, I use delay/delay-rate (DDR) filters to extract estimates of individual source flux densities as a function of time and frequency. The frequency information can be used to perform a frequency-dependent beam calibration if the SNR in the observations is high enough. These filters work in delay/delay-rate space (the Fourier dual of frequency/time space) to isolate flux density per baseline from a specific geometric area on the sky; for a full description of the technique, the reader is referred to Parsons & Backer (2009).

The first step of our pipeline to produce perceived flux density estimates is to filter the Sun from each baseline individually in DDR space. This is done to ensure that little to no signal from the Sun, which is a partially-resolved and time-variable source, remains in the data. (In principle, data from different observing runs separated by several months could provide complete sky coverage while avoiding daytime data altogether. We chose to use data from only one 24 hour period to minimize the effect of any long timescale instabilities in the system). A Markov-Chain Monte Carlo (MCMC) method for extracting accurate time- and frequency-dependent source models via self-calibration using DDR-based estimators is then used to model the 4 brightest sources remaining: Cygnus A, Cassiopeia A, Virgo A, and the Crab Nebula (Taurus A). The MCMC aspect of this algorithm iterates on a simultaneous model of all sources being fit for to minimize sidelobe cross-terms between sources. After removing the models of Cygnus A, Cassiopeia A, Virgo A and Taurus A from the data, a second pass of the MCMC DDR self-calibration algorithm extracts models of the remaining 22 sources listed in Table 3.1.

3.3.2 Gridding the Measurements

One can increase the signal-to-noise at a crossing point by combining all measurements within a region over which the primary beam response can be assumed to be constant. To define these regions, we grid the sky. In this work, the beam model is constructed on an HEALPix map (Górski et al. 2005) with pixels $\approx 0.9°$ on a side. The choice of grid pixel size is somewhat arbitrary. Using a larger pixel size broadens the beam coverage of each source track, creating more crossing points and helping to constrain the overall beam shape. However, when the pixels are too large, each pixel includes data from sources with larger separations on the sky. Since the principal tenet of this approach is that each source within a crossing point sees the same primary beam response, excessively large pixels can violate this assumption, resulting in an inaccurate beam model. For PAPER data, a HEALPix grid with $0.9°$ pixels is found to be a good balance between these competing factors, as will be explained in §3.4.1. For other experiments with narrower, more rapidly evolving primary beams, smaller pixels may be necessary.

To introduce measurements of perceived flux density into the grid, I first recast Equation 3.1 into a discrete form:

$$M_i = b_i f_k, \tag{3.2}$$

where M_i and b_i are the respective perceived source flux densities and primary beam responses in the pixel i, and k is a source index labelling the inherent source flux density, f. To generate a single measurement of a source for each pixel i, I use a weighted average of all measurements of that source falling with a single pixel. The weights are purely geometric and come from interpolating the measurement between the four nearest pixels.

3.3.3 Forming a Least-Squares Problem

Once all the data are gridded, I solve Equation 3.2 for all crossing pixels simultaneously. To do this, I set up a linearized least-squares problem using logarithms:

$$\log(M_i) = \log(b_i) + \log(f_k). \tag{3.3}$$

Because thermal noise in measurements becomes non-Gaussian in Equation 3.3, the solution to the logarithmic formulation of the least-squares problem is biased. In §3.4.1, I investigate the effect of this bias using simulations and find that the accuracy of our results is not limited by this bias, but instead by sidelobes of other sources. Therefore, while Equation 3.1 can in principle be solved without resorting to logarithms using an iterative least-squares approach, I find that this is not necessary given our other systematics.

Once the logarithms are taken, one can construct a solvable matrix equation, which can be expressed generally as:

$$\mathbf{WM} = \mathbf{WAX}, \tag{3.4}$$

where $\mathbf{W}$ is a weighting matrix with weights W_i on the diagonal and zero on the off-diagonals, $\mathbf{M}$ is a column-matrix of measurements, $\mathbf{A}$ is the matrix-of-condition, describing which

measurements are being used to constrain which parameters, and $\mathbf{X}$ is a column matrix containing all the parameters one wishes to measure: the beam responses b and the source flux densities f. To illustrate the form of this equation, I present the matrix representation of the Cygnus A/Virgo A system shown in Figure 3.1:

$$\begin{pmatrix} W_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & W_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & W_3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & W_4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & W_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & W_6 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & W_7 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_8 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_{10} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_{11} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_{12} \end{pmatrix} \begin{pmatrix} \log M_1 \\ \log M_2 \\ \log M_3 \\ \log M_4 \\ \log M_5 \\ \log M_6 \\ \log M_7 \\ \log M_8 \\ \log M_9 \\ \log M_{10} \\ \log M_{11} \\ \log M_{12} \end{pmatrix}$$

$$= \begin{pmatrix} W_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & W_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & W_3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & W_4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & W_5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & W_6 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & W_7 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_8 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_{10} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_{11} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & W_{12} \end{pmatrix}$$

$$\times \left(\begin{array}{cccccc|cc} 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ \hline 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ \hline 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{array}\right) \begin{pmatrix} \log b_1 \\ \log b_2 \\ \log b_3 \\ \log b_4 \\ \log b_5 \\ \log b_6 \\ \log f_{cyg} \\ \log f_{vir} \end{pmatrix} \tag{3.5}$$

On the left-hand side of the equation are the two matrices $\mathbf{W}$ and $\mathbf{M}$. The weights, W_i, are

defined in §3.3.3. **M** contains logarithms of the perceived source flux density measurements in each pixel, M_i. Recall that each M_i corresponds to one source.

On the right-hand side, the weighting matrix **W** appears again, followed by the matrix-of-condition **A**, and then **X**, a column matrix containing the parameters we wish to solve for: the primary beam response at the 6 crossing points and the flux densities of Virgo and Cygnus. The matrix-of-condition, **A**, identifies which sources and crossing points are relevant for each equation. I have schematically divided it: to the left of the vertical line are the indices used for selecting a particular crossing point; to the right are those for the sources. The first 4 lines represent the two northern Cygnus/Virgo crossing points, and the next four represent the two southern ones (which are identical copies of the northern ones). Finally, the last 4 lines represent the 2 points where Virgo crosses itself; notice that the Cygnus source column is blank for these 4 rows.

It should be noted that Equations 3.4 and 3.5 contains no absolute flux density reference. The simplest way to set this scale is to treat all the recovered flux densities as relative values compared to the flux calibrator (Cygnus A in the case presented here). One can then place all the flux densities onto this absolute scale. An equally valid approach is to append an extra equation with a very high weight, which sets the flux calibrator to its catalog flux value.

Weighting of Measurements in the Least-Squares Formalism

In a least-squares approach, optimal weights are inversely proportional to the variance of each measurement. To calculate the variance of each measurement, one must propagate the uncertainty in the initial perceived flux density measurements through the averaging and logarithm steps. The noise level in each interferometric visibility is roughly constant, and the DDR filters average over a fixed number of visibilities for each perceived flux density estimate, leading to equal variance at each time sample. To produce optimal weights accounting for the many time samples averaged into each beam pixel and the propagation of noise through the logarithm, each logarithmic measurement in Equation 3.4 should be weighted by:

$$W_i = \sqrt{M_i^2 \sum_j w_j}, \tag{3.6}$$

where j indexes the time step (which is generally fast enough to produce many measurements inside a pixel), w_j is the geometric sub-pixel interpolation weighting of each measurement, and M_i is the weighted average of all measurements of a source's flux density in pixel i. Without the square root, these weights would be proportional to the inverse of the variance in each logarithmic measurement; the sum over the geometric weights is the standard reduction of variance for a weighted average, and the factor of M_i^2 comes from propagating variances through the logarithm. The square root appears because a least-squares solver using matrix inversion will add in an additional factor of the weight, leading to the desired inverse-variance weighting. Other solvers using different methods may require different weights.

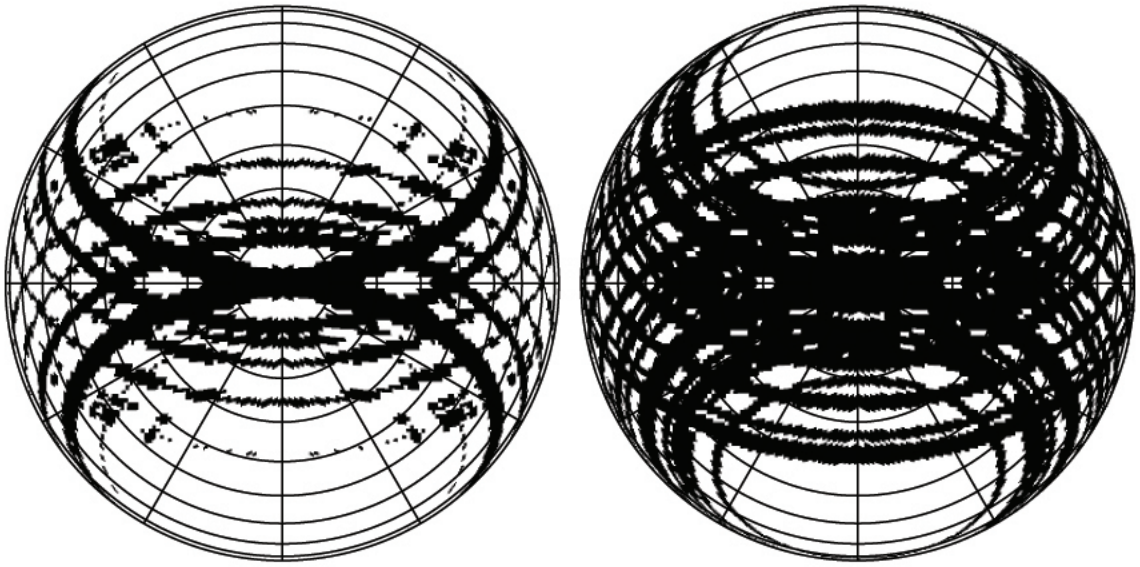

Figure 3.2: The sky coverage of crossing points and source tracks. *Left:* The location of crossing points for the 25 sources listed in Table 3.1 used to calibrate the PAPER primary beam. The projection is orthographic with zenith at the center; dotted lines are 10° and 30° steps in elevation and azimuth, respectively. *Right:* The sky coverage of all 25 source tracks. Although the least-squares inversion solves only for the beam response at crossing points, one can include all source track data using the recovered flux density of each source to create a primary beam estimate along the entire track.

Solving the Least-Squares Equation for a First-Order Beam Model

The solution of Equation 3.4, the matrix $\mathbf{A}$, contains two distinct sets of parameters: the beam responses at crossing points and the flux densities of each source (modulo an absolute scale). There is additional information that may be included to improve the beam model beyond that generated by solving for the responses at crossing points. Given the flux-density solutions for each source, each source track now provides constraints on beam pixels that are not crossing points. By dividing each perceived flux density source track by the estimated inherent flux density from the least-squares inversion, one can produce a track of primary beam responses, with greater coverage than that provided by crossing points alone. I illustrate the difference in coverage in Figure 3.2. The left hand panel shows the locations of crossing points for the 25 sources used to calibrate the PAPER beam, listed in Table 3.1. The right hand panel shows the increased beam coverage that comes from including non-crossing point source-track data. To create an initial beam model, I average within each pixel the estimated

beam responses from each source, weighting by the estimated flux of that source. Given equal variance in each initial perceived flux density measurement, this choice of weights is weighting by signal to noise.

3.3.4 Using Deconvolution to Fill in Gaps in the Beam Model

To produce a model of the beam response in any direction, one must fill in the gaps left by limited sky sampling. I use a CLEAN-like (Högbom 1974) deconvolution algorithm in Fourier space to fill in the holes in the beam. I iteratively fit and remove a fraction of the brightest Fourier component of our beam until the residuals between the model and the data are below a specified threshold. In addition to measured beam responses derived from source tracks, I add the constraint that the beam must go to zero beyond the horizon. In the deconvolution, each pixel is weighted by the estimate of the beam response in that pixel, reflecting that SNR will be highest where the beam response is largest. This beam-response weighting was unnecessary in previous steps, since the least-squares approach solved for each pixel independently. This weighting scheme again represents signal-to-noise weighting, given the equal variance in the initial perceived flux density estimates. The result of the deconvolution is an interpolated primary beam model with complete coverage across the sky.

3.3.5 Introduction of Prior Knowledge

Up to this point, I have made only two fairly weak assumptions about the beam: that it possesses 180° rotational symmetry, and that the response is zero below the horizon. However, if one has additional prior knowledge about the beam, one can better constrain the final model. In particular, one can use beam smoothness constraints to identify unphysically small scale features introduced by sidelobes in the source extraction or by incomplete sampling in the deconvolution. I choose to incorporate additional smoothness information by filtering our model in Fourier space to favor large-scale modes.

PAPER dipoles were designed with emphasis on spatial and spectral smoothness in primary beam shape. To smooth the PAPER beam model, I choose a cutoff in Fourier space that corresponds to the scale at which $> 99.9\%$ of the power is accounted for in a computed electromagnetic model of the beam. While such a filter is not necessarily generalizable to other antenna elements, I find it necessary to suppress the substantial sidelobes associated with observations from a 12-antenna PAPER array that are discussed below.

3.4 Application to Simulated Data

To test the robustness of this approach, I apply it to several simulated data sets. The results of these simulations, with and without Gaussian noise, are described below in §3.4.1. I also simulate raw visibility data to test the effectiveness of source extraction; these simulations are discussed in §3.4.2. The major difference between these two methods of simulation is

that the simulation using raw visibility data allows for imperfect source isolation, leading to contamination of the source tracks by sidelobes of other sources. These sidelobes have a significant effect on the final beam model that is derived.

3.4.1 Simulations of Perceived Flux Density Tracks

I simulate perceived flux density tracks using several model beams of different shapes, including ones with substantial ellipticity and an $\sim 15^\circ$ rotation around zenith. I also input several source catalogs, including a case with 10,000 Jy sources spaced every degree in declination, and a case using the catalog values and sources listed in Table 3.1 that approximately match the sources extracted from observations. In all combinations of beam models and source catalogs, I recover the input beam and source values with $< 5\%$ error. The average error in source flux density is 2.5%. I see no evidence for residual bias, as the distribution of error is consistent with zero mean to within one standard deviation.

I also test the effect of adding various levels of Gaussian noise to a simulation involving a fiducial beam model and the 25 sources listed in Table 3.1. Only when the noise level exceeds an RMS of 10 Jy in each perceived flux density measurement does the mean error in the solutions exceed 10%. As the noise increases beyond this level, there is a general trend to bias recovered flux densities upward; as mentioned earlier, this bias is introduced by the logarithms in Equation 3.3. However, the expected corresponding noise level of DDR-extracted perceived flux density measurements from the 12-element PAPER array is < 1 Jy per sample. At a simulated rms noise level of 1 Jy, the solutions are recovered with a mean error of $< 3\%$. This result validates my previous statement that the bias introduced by the logarithms in Equation 3.3 is not a dominant source of error.

It is also worth noting that these simulations were used to identify 0.9° as the best HEALPix pixel side for the grid; with too large a pixel size (1.8°) the model becomes significantly compromised. This results from combining measurements of sources that are subject to significantly different primary beam responses. I choose not to use a smaller pixel size, since it will reduce the fractional sky coverage of our crossing points and increase the computational demand of the algorithm.

3.4.2 Simulations of Visibilities

I also apply this technique to simulated visibilities in order to test the complete analysis pipeline, including the DDR-based estimation of perceived source flux densities. The visibility simulations are implemented in the AIPY[8] software toolkit. The simulated observations correspond to actual observations made with the 12-element PAPER deployment in Green Bank described in §3.5. I match the observations in time, antenna position, and bandwidth. I also include the expected level of thermal noise in the simulated visibilities. I simulate "perceived" visibilities by attenuating the flux density of each source by a model primary

[8] http://pypi.python.org/pypi/aipy/

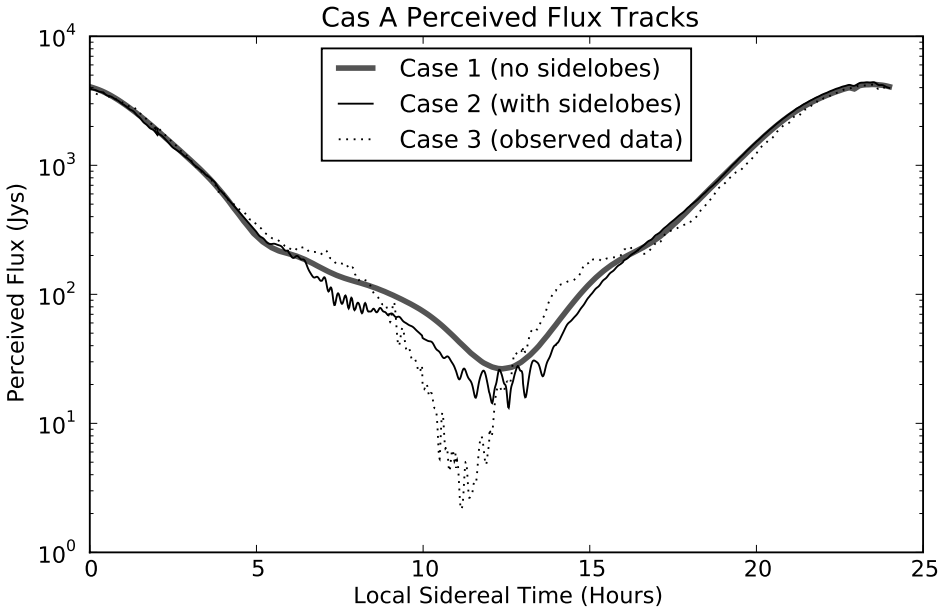

Figure 3.3: Three different perceived flux tracks for Cassiopeia A. Case 1 is simply a cut through a model beam multiplied by a catalog flux density value; this is the simplest case and contains no sidelobe contamination. Case 2 uses DDR filters to extract the source track from simulated visibilities; case 3 does the same, but applied to observed data. The simulated visibilities in case 2 demonstrate the features that arise from sidelobes of other sources during the DDR source extraction. However, the simulation employs a simplistic point-source sky and cannot reproduce all sidelobe features seen in the real data.

beam. The DDR filters return estimates of perceived source flux density for the input primary beam.

In these simulations, I only include bright point sources and a uniform-disk model of the Sun. As a result, source extraction is expected to be more accurate in simulation than in real data, since the sources I extract account for 100% of the simulated signal. However, these simulations do provide a useful test of DDR-filter-based source extraction and of the level of contamination from sidelobes.

As might be expected, the estimates of perceived source flux density versus time from simulated visibilities contain structure that is not attributable to the beam. These features are almost all due to sidelobes of other sources; they persist in real data and are reproducible over many days. Figure 3.3 shows the difference in the source tracks produced for Cassiopeia A by each simulation method (cases 1 and 2 in the figure), as well as the track extracted from the observed data described §3.5 (case 3). I find that the DDR filter only extracts a fraction of the flux density from the Crab Nebula. This bias is unique to Crab, and is most likely a result of the proximity of the Sun to Crab in these observations, coupled with the limited number of independent baselines in the 12-antenna array. For this reason, I choose to

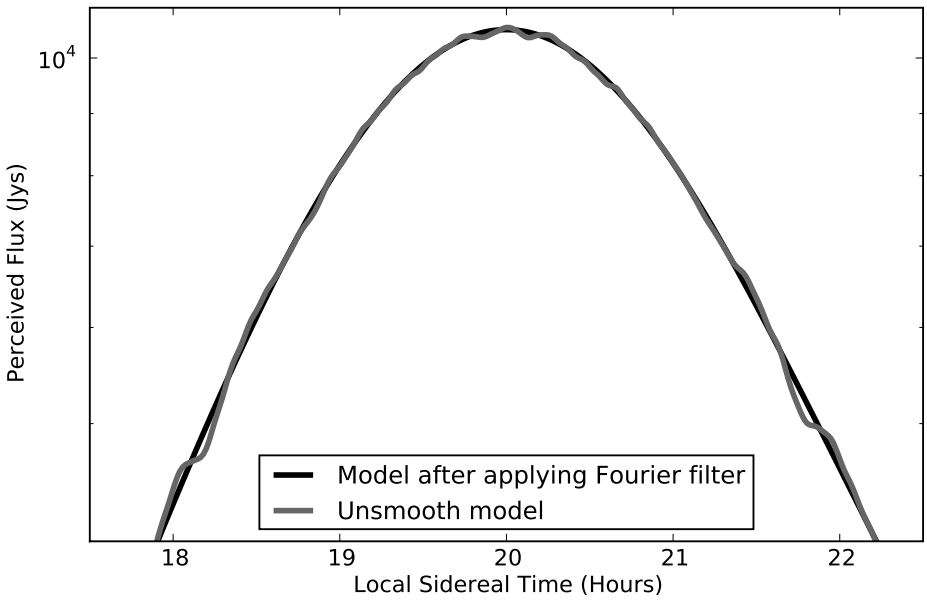

Figure 3.4: The effect of the Fourier domain filter. The deconvolution of §3.3.4 fills in gaps in sky coverage, but leaves small scale structure that is not associated with the primary beam (gray curve). By filtering the solution in the Fourier domain to incorporate prior knowledge of beam smoothness, one can achieve an improved beam model (black curve).

exclude the Crab from our final analysis.

Performing the least-squares inversion described in §3.3.3 on the simulated source tracks, I find that the sidelobes introduce errors into estimates of source flux densities. The average error in flux density is 10%, with the largest errors exceeding 20%. The distribution of errors is consistent with zero mean, indicating no strong biasing of the flux densities.

I also find that the resultant primary beam model matches the input beam to within 15% percent. However, the model contains small scale variations, as seen in Figure 3.4. Using prior knowledge of beam smoothness, one can reject these features as unphysical and apply a Fourier-domain filter, as described in §3.3.5. This filter reduces the effect of sidelobes in the source tracks, as they appear on Fourier scales that are not allowed by smoothness constraints. Even with a relatively weak prior on smoothness (i.e., retaining more Fourier modes than are present in the input model), we substantially reduce the small scale variations in our beam. The effect of the filter on small scale variations is illustrated in Figure 3.4. The filter also reduces large scale errors in the model, bringing the overall agreement to with 10% of the input beam.

In summary, I find that although interfering sidelobes from other sources compromise our source flux density measurements, our approach recovers the input primary beam model at the 15% level. With the introduction of a Fourier space filter motivated by prior knowledge

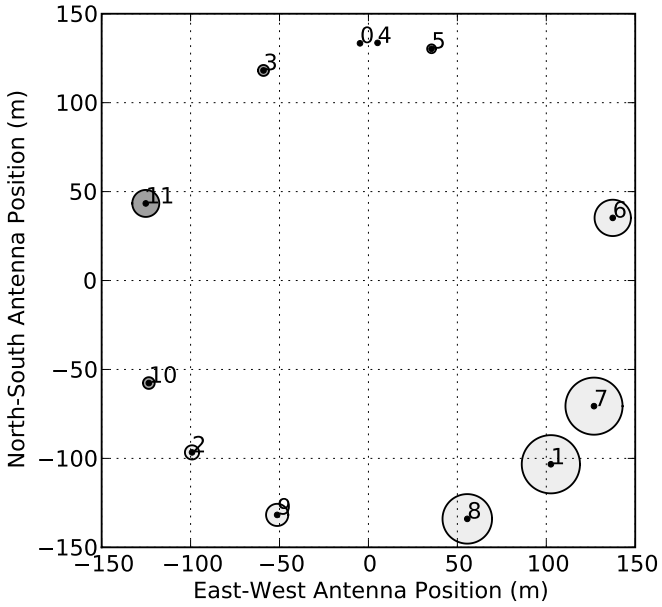

Figure 3.5: The 12 element PAPER array in Green Bank. Shaded circles represent vertical height relative to antenna 0; the radius of the circle is 10 times the height in meters. Light gray is up, dark gray is down.

of the beam smoothness, the output model improves to within 10% of the input. It is also worth noting that the effectiveness of the beam calibration technique presented here will substantially improve with larger arrays, where sidelobe interference will be reduced and source tracks will more closely resemble the ideal case discussed in §3.4.1.

3.5 Observed Data

In this section, I describe the application of this technique to 24 continuous hours of data taken with the PAPER array deployed at the NRAO site near Green Bank from July 2 to July 3, 2009. At this time, the array consisted of 12 crossed-dipole elements. Only the north/south linear polarization from each dipole was correlated. The data used in this analysis were observed between 123 to 170 MHz (avoiding RFI) and were split in 420 channels. The array was configured in a ring of 300 m radius. The longest baselines are 300 m, while the shortest baseline is 10 m; the configuration is shown in Figure 3.5. This configuration gives an effective image plane resolution of 0.4°.

3.5.1 Data Reduction

In addition to the DDR-filter algorithm described in §3.3.1, several other pre-processing steps are necessary. Per-antenna phase and amplitude calibration are performed by fringe fitting to Cygnus A and other bright calibrator sources. I also use a bandpass calibration based on the spectrum of Cygnus A. This calibration is assumed to be stable over the full 24 hours. Two further steps in the reduction pipeline were first described in Parsons et al. (2010): gain linearization to mitigate data quantization effects in the correlator, and RFI excision.

New to this work is the use of the cable and balun temperatures to remove time-dependent gains caused by temperature fluctuations. A thirteenth dipole was operated as the "gain-o-meter" described in Parsons et al. (2010). In brief, the "gain-o-meter" is an antenna where the balun is terminated on a matched load, rather than a dipole. The measured noise power from this load tracks gain fluctuations in the system. We record the temperature of several components of our "gain-o-meter" using the system described in Parashare & Bradley (2009) and Chapter 2 of this dissertation. We correct for the temperature dependent gain changes by applying the linear correction derived from the measured temperature values given in Chapter 2. This step is crucial for the success of the beam calibration; without correcting for them, these gain drifts are indistinguishable from an east/west asymmetry in the primary beam. After these steps, we process the data with the DDR-filtering algorithm to produce estimates of perceived flux density versus time for the 25 sources listed in Table 3.1.

3.5.2 Results

PAPER Primary Beam Model

The first estimate of the PAPER primary beam derived with this approach is shown in Figure 3.6. This figure is produced by dividing each source track by an estimate of its inherent flux density produced by the least squares inversion. This transforms each track into an estimate of the primary beam response, which are the added together into a HEALPix map. There are significant fluctuations from pixel to pixel that are clearly unphysical, and relate to source sidelobes. There are also significant gaps in beam coverage, even with 25 source tracks.

Figure 3.7 shows the results after deconvolving the sampling pattern from Figure 3.6. The beam is now complete across the entire sky, although there is still substantial small-scale structure unrelated to the inherent beam response. I suppress these features for the final model by retaining only the large scale Fourier components, as described in §3.3.4. This final model is shown in Figure 3.8. It is clear that with our deconvolution and choice of Fourier modes, we recover a very smooth, slowly evolving beam pattern, as expected from previous models of the PAPER primary beam.

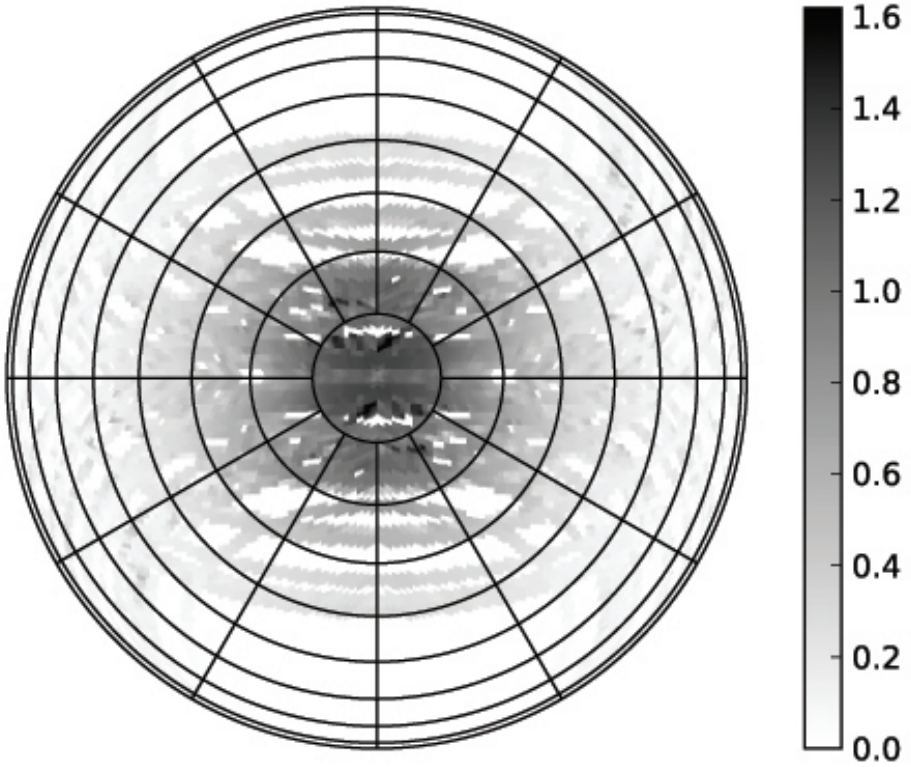

Figure 3.6: The initial solution for the PAPER primary beam response along source tracks. Each perceived flux density track is divided by the estimate of that source's inherent flux density produced by the least-squares inversion. This yields estimates of the primary beam response along each track. With this initial solution, a large fraction of the sky is already covered by the data, although there is a fair bit of variation from point to point. The color scale is linear and normalized to 1.0 at zenith; as in Figure 3.1, dotted lines are 10° and 30° steps in elevation and azimuth, respectively.

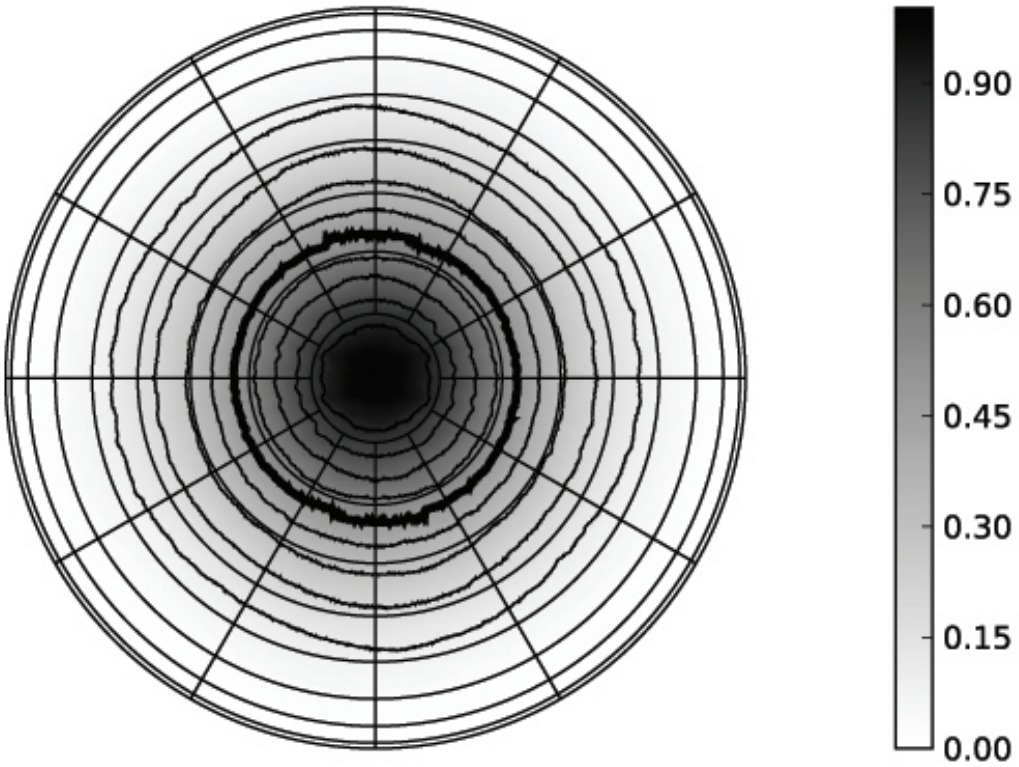

Figure 3.7: The deconvolved PAPER beam model. The grayscale colors are linear and show the primary beam response, normalized to 1.0 at zenith. Solid black lines are contours of constant beam response in 10% increments; the thick black line is the half-power point. As in Figure 3.1, dotted lines are 10° and 30° steps in elevation and azimuth, respectively. This model was produced by deconvolving the sampling pattern from Figure 3.6, interpolating over the gaps in declination where there are no strong calibrator sources.

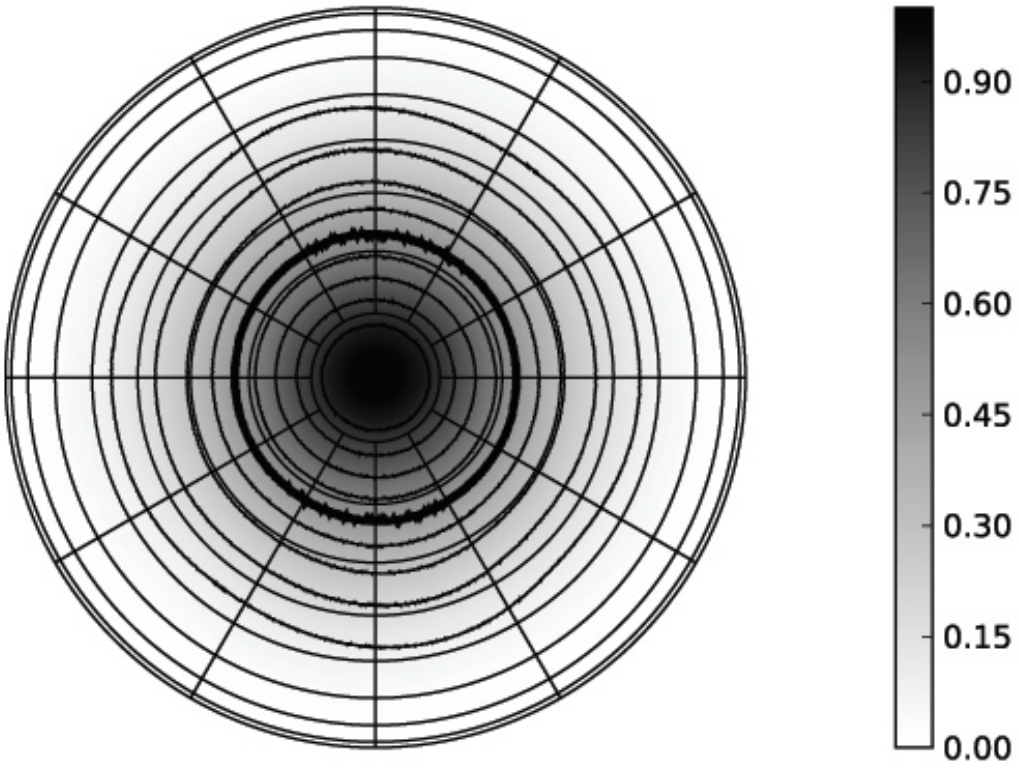

Figure 3.8: The smooth PAPER beam model. As in Figure 3.7, the grayscale colors are linear and show the primary beam response, normalized to 1.0 at zenith. Solid black lines are contours of constant beam response in 10% increments; the thick black line is the half-power point. Dotted lines are 10° and 30° steps in elevation and azimuth, respectively. This smooth model was produced by retaining only the large scale Fourier components from Figure 3.7. The cutoff in Fourier space was derived from a computed electromagnetic PAPER beam simulation. This Fourier mode selection substantially smooths out variations in the initial solution that are unphysical based on the smoothness of beam response in the simulation.

Source Catalog

The measured flux densities of the calibrator sources, produced by the least-squares inversion, are presented in Table 3.1. The overall flux scale is normalized to the value of Cygnus A reported in Baars et al. (1977). Error-bars are estimated by measuring the change in flux necessary to increase $\Delta\chi^2$ between the model and the measured perceived source flux density tracks by 1.0. As noted in §3.4.2, these flux densities can be compromised by sidelobes of other sources. However, the results are generally accurate to within ±10%.

3.5.3 Tests of Validity

I perform several tests to investigate the validity of the results. The comparison of the recovered fluxes to their catalog values shows generally good agreement. As argued in Vollmer et al. (2005) and Jacobs et al. (2011), there is considerable difficulty in accurately comparing catalogs, particularly at these frequencies. Differences in interferometer resolution and synthesized beam patterns can lead to non-trivial disagreement in measured source flux densities. Given this fact, and the complicated sidelobes associated with a 12-element array present in the data, I do not find the lack of better agreement with catalog fluxes troubling.

To test the stability of our pipeline, I test a separate observation spanning the following day. The beam model produced by this data matches the first model to within 2.5%, and source flux densities remain consistent to 5%. This test confirms the suspicion that our errors are not dominated by random noise.

3.6 Conclusions

I have presented a new technique for calibrating the primary beam of a wide-field, drift-scanning antenna element. The key inputs to the method are measurements of perceived flux density of individual sources versus time as they drift through the beam. In the analysis presented in this chapter, I use delay/delay-rate filters (Parsons & Backer 2009) as estimators in a self-calibration loop to extract such measurements from raw interferometric data. However, the remainder of the method is agnostic to how these tracks are derived. The only assumption necessary to make this an over-constrained and solveable problem is 180° rotational symmetry in the beam. With this assumption, one creates “crossing points" where there is enough information for a least-squares inversion to solve for both the inherent flux density of each source and the response of primary beam at the crossing points.

I test this approach using simulated tracks of perceived source flux density across the sky and simulated visibilities. Using the source tracks, the least-squares inversion reliably recovers the primary beam values and source flux densities in the presence of noise 10 times that present in a 12-element PAPER array. The simulated visibilities demonstrate that sidelobes of other bright sources are the source of the dominant errors in source extraction when only 12 antennas are used. The presence of these features in the perceived source flux densities limits the accuracy of estimates of inherent source flux densities. However, in simulation, I

Table 3.1: Measured flux densities for all the sources used in this work. Catalog values come from the 3C catalog (Edge et al. 1959) at 159 MHz unless otherwise noted. Errors are estimated by measuring the change in flux required to increase $\Delta\chi^2$ between the model and measured perceived flux density source tracks by 1.0.

RA	Dec	Measured Flux Density (Jy)	Catalog Flux Density (Jy)	Name
1:08:54.37	+13:19:28.8	$78.1^{+14.7}_{-24.4}$	58, 49[a]	3C 33
1:36:19.69	+20:58:54.8	$47.3^{+12.3}_{-37.3}$	27	3C 47
1:37:22.97	+33:09:10.4	$60.0^{+10.2}_{-13.7}$	50	3C 48
1:57:25.31	+28:53:10.6	$36.2^{+9.8}_{-23.6}$	7.5, 16[b], 23[c]	3C 55
3:19:41.25	+41:30:38.7	$81.7^{+12.0}_{-16.5}$	50	3C 84
4:18:02.81	+38:00:58.6	$75.9^{12.2}_{-14.0}$	60	3C 111
4:37:01.87	+29:44:30.8	$205.3^{+26.9}_{-32.1}$	204	3C 123
5:04:48.28	+38:06:39.8	$100.3^{+17.0}_{-20.2}$	85	3C 134
5:42:50.23	+49:53:49.1	$56.5^{+9.7}_{-11.5}$	63	3C 147
8:13:17.32	+48:14:20.5	$66.3^{+9.5}_{-9.9}$	66	3C 196
9:21:18.65	+45:41:07.2	$48.5^{+7.5}_{-9.0}$	42	3C 219
10:01:31.41	+28:48:04.0	$40.3^{+8.7}_{-15.6}$	30	3C 234
11:14:38.91	+40:37:12.7	$34.5^{+7.0}_{-19.9}$	21.5	3C 254
12:30:49.40	+12:23:28.0	$1108.9^{+29.6}_{-28.1}$	1100	Vir A
14:11:21.08	+52:07:34.8	$64.4^{+6.4}_{-6.7}$	74	3C 295
15:04:55.31	+26:01:38.9	$73.5^{+7.4}_{-9.0}$	72	3C 310
16:28:35.62	+39:32:51.3	$60.5^{+7.2}_{-12.4}$	49	3C 338
16:51:05.63	+5:00:17.4	$461.7^{+78.6}_{-112.7}$	325[a], 378[b], 373[c]	Her A
17:20:37.50	-0:58:11.6	$228.8^{+61.7}_{-147.5}$	180, 236[a], 276[b], 215[c]	3C 353
18:56:36.10	+1:20:34.8	$251.3^{+60.0}_{-135.8}$	680	3C 392[d]
19:59:28.30	+40:44:02.0	$10622.9^{+126.0}_{-128.9}$	10623	Cyg A[e]
20:19:55.31	+29:44:30.8	$62.2^{+16.4}_{-28.0}$	36	3C 410
21:55:53.91	+37:55:17.9	$51.9^{+11.5}_{-16.2}$	43	3C 438
22:45:49.22	+39:38:39.8	$71.3^{+12.4}_{-19.3}$	50	3C 452
23:23:27.94	+58:48:42.4	$9198.1^{+203.4}_{-186.9}$	6230[f]	Cas A

[a] 3CRR (Laing et al. 1983) at 178 MHz.
[b] Culgoora (Slee 1995) at 160 MHz.
[c] PSA32 (Jacobs et al. 2011). Measurements made by PAPER.
[d] Extended (35x27 arcmin) supernova remnant (W44).
[e] Flux density calibrator, using values from Baars et al. (1977).
[f] Extrapolated using the formula in Baars et al. (1977). This is probably a substantial underestimate of the flux density, as suggested by Helmboldt & Kassim (2009).

am able to recover a primary beam model accurate to within 15% percent and source flux densities with an average error of 10%. Using prior information regarding beam smoothness, I improve the model to better than 10% accuracy.

While these caveats about the effectiveness of the least-squares technique in the presence of sidelobes may seem worrisome, it bears repeating that these are issues with data quality and not with the technique itself. For example, data from a 32-element PAPER array has shown that DDR filters can extract the sources used in this analysis with little systematic biases. (I do not use this data here do the lack of a temperature record for gain stabilization and the presence of a particularly active Sun when the array was operating.) Therefore, it seems that this technique has significant potential for precise beam calibration on larger arrays.

Another future goal for this technique is to calibrate the frequency-dependence of the beam. Here, I have used the entire bandwidth to improve signal-to-noise in source extractions. For a larger array with higher SNR and lower sidelobes, perceived source flux density measurements can be cut into sub-bands to look at the beam as a function of frequency.

Finally, it is possible to forgo the assumption of 180° rotational symmetry altogether and allow for possible north/south variation in the beam. An experiment in which the dipoles are *physically* rotated on a daily basis can be used to create the same kind of crossing points, since one has changed the section of the beam each source crosses through. Work is progressing on such an experiment using the PAPER array.

Acknowledgments

A version of this chapter was previously published as Pober et al. 2012 AJ 143 53, and is reproduced with the permission of all coauthors and the copyright holder.

Chapter 4

Opening the 21cm EoR Window: Measurements of Foreground Isolation with PAPER

In this chapter, I present new observations with the Precision Array for Probing the Epoch of Reionization (PAPER) with the aim of measuring the properties of foreground emission for 21cm Epoch of Reionization experiments at 150 MHz. I focus on the footprint of the foregrounds in cosmological Fourier space to understand which modes of the 21cm power spectrum will most likely be compromised by foreground emission. These observations confirm predictions that foregrounds can be isolated to a "wedge"-like region of 2D $(k_\perp, k_\parallel)$-space, creating a window for cosmological studies at higher $k_\parallel$ values. I also find that the emission extends past the nominal edge of this wedge due to inherent spectral structure in the foregrounds themselves, with this feature most prominent on the shortest baselines. Finally, I filter the data to retain only this "unsmooth" emission and image specific $k_\parallel$ modes of it. The resultant images show an excess of power at the lowest modes, but no emission can be clearly localized to any one region of the sky. This image is highly suggestive that the most problematic foregrounds for 21cm EoR studies will not be easily identifiable bright sources, but rather an aggregate of fainter emission.

4.1 Introduction

The highly redshifted 21cm line of neutral hydrogen is widely regarded as one of the most promising probes of the high redshift universe, with potential to map out volumes extending from redshift ~ 1 through the Epoch of Reionization (EoR) to the dark ages at redshift 20 and beyond (for reviews of the field, see Furlanetto et al. 2006, Morales & Wyithe 2010, and Pritchard & Loeb 2012). Numerous facilities and experiments targeting the signal from the EoR are already online or under construction, including the LOw Frequency ARray (LOFAR;

Yatawatta et al. 2013)[1], the Murchison Widefield Array (MWA; Tingay et al. 2012)[2], and the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010)[3]. All 21cm cosmology experiments will need to separate bright galactic and extragalactic foregrounds from the neutral hydrogen signal, which can be fainter by as much as 5 orders of magnitude or more (see, e.g., Furlanetto et al. 2006 and Santos et al. 2005).

Almost all of these approaches rely on the spectral smoothness of foreground emission relative to the 21cm signal, which will contain significant structure versus frequency. The purpose of this chapter is to use the delay transform technique (presented in Parsons et al. 2012b — referred to as P12b for the remainder of this chapter — and briefly described in Chapter 1) on observations from PAPER to test the behavior of actual foreground emission. The ultimate goal is to understand the footprint of foregrounds in k-space to determine which modes of the 21cm power spectrum will be most accessible to observation. The structure of this chapter is as follows: in §4.2, I describe the data used in these observations. In §4.3, I review the delay spectrum technique presented in P12b, and then describe the steps used in applying this approach to actual observations. I present the results in §4.4 and conclude in §5.5.

4.2 The Data

I use 4 hours of data collected in 10 second integrations between JD 2455747.48 and 2455747.64 ($4-5$ July 2011), using a 64 element PAPER array located on the SKA site in the Karoo region of South Africa. This data set comes from the same observing campaign described by Stefan et al. (2012), although this specific 4-hour window falls outside of the observations analyzed therein. The dipole antennas are arranged in a "minimum redundancy" configuration optimized for imaging analysis (Parsons et al. 2012a), which is shown in Figure 4.1. This configuration has a maximum baseline length of ~ 300m, corresponding to an image plane resolution of 0.4° at 150 MHz. The PAPER correlator has a 100 MHz instantaneous bandwidth from $100-200$ MHz divided into 2048 frequency channels. In this data set, only one linear polarization on each dipole is correlated, and all data from antennas 40 and 55 are discared due to cross-polarization.

An image of the field transiting during this time period is shown in Figure 4.2. This image spans $140-165$ MHz; 100 sub-bands of 0.25 MHz were individually imaged and summed to make the map shown. No CLEANing or spectral slope correction was performed. The observation is centered on a low-foreground "cold patch," a potential field for an EoR science observation. The Galactic plane, which is just setting at the end of the 4-hour observation, creates prominent sidelobes over the entire map.

Complex antenna based gains are derived using fringe-fitting to Centaurus A, Pictor A, and Fornax A; an overall gain scale is derived from the Helmboldt et al. (2008) source

[1] http://www.lofar.org/
[2] http://www.mwatelescope.org/
[3] http://eor.berkeley.edu/

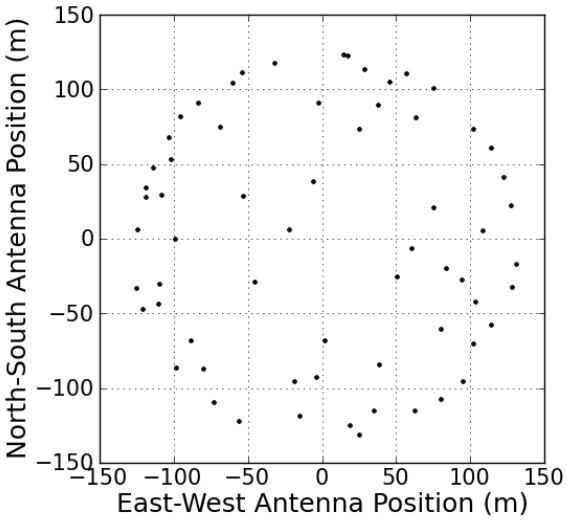

Figure 4.1: The configuration of the 64 PAPER dipoles used in this analysis. The zero-point is the center of the array. The y-axis is North/South, the x-axis is East/West, and distances are in meters.

J2214-170. I perform a small gain linearization correction to mitigate data quantization effects in the correlator (described in Parsons et al. 2010) and a correction for temperature dependent gain drifts (described in Chapter 2). I also perform radio-frequency interference (RFI) excision, manually flagging frequency channels of known transmitters, and flagging any points 6σ above the mean after differencing adjacent channels along the frequency axis.

4.3 Analysis Techniques

At the core of this analysis is the delay spectrum technique presented in P12b. Since cosmological redshifting maps the observed 21cm line frequency into a distance measurement, the Fourier transform of the frequency axis are the $k_{\parallel}$ line-of-sight modes of the 21cm power spectrum. However, this relation is clearly not true for foregrounds, where the frequency axis simply corresponds to the spectra of the sources. The delay transform first presented in Parsons & Backer (2009) provides a framework for mapping foreground emission into a cosmological k-space. The frequency Fourier transform of a single baseline's visibility spectrum (the "delay transform") maps celestial emission to "delay space" where sources appear as Dirac delta functions, located at the delay in arrival time between the two elements of that baseline. These delays must be limited to values below the physical light travel time between the two antennas (the "horizon limit"), and so all emission from the sky maps to a region in the center of delay space determined only by the baseline length. However, any spectral structure in the visibilities acts as a convolving kernel in delay space. Foreground emission is spectrally smooth, translating into a narrow convolving kernel; 21cm emission has

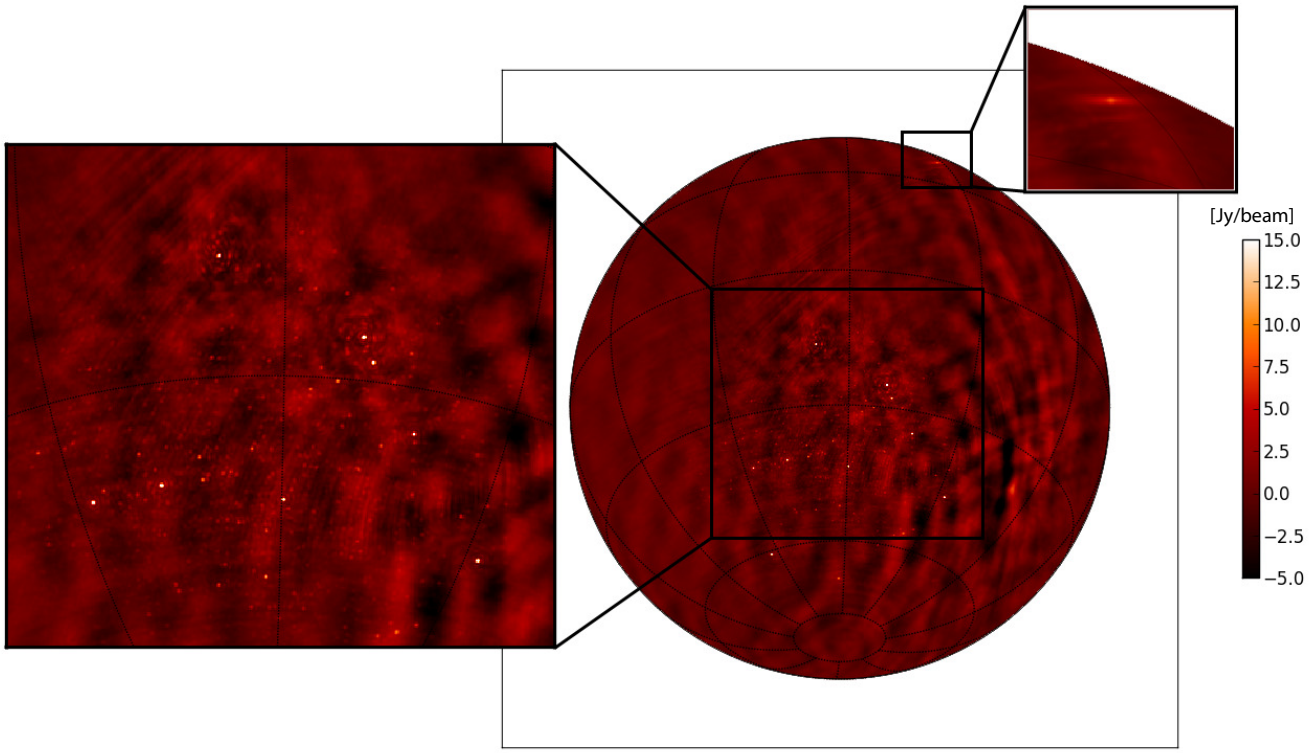

Figure 4.2: A dirty image of the data used in this analysis, centered on RA 21h52m and declination -30°43', the transiting zenith halfway through the observation. Prominent sidelobes from the Galaxy are seen in the right-hand side of the map. Close-ups show Cygnus A at 19h59m and +40°44' and the point source population of a potential EoR cold-patch. The color-scale is linear in Jy, with only the brightest point sources saturating the range.

large amounts of spectral structure, and therefore its kernel scatters power from within the horizon limit to larger delays. P12b also showed that delay has a near one-to-one mapping to $k_\parallel$, meaning that those delay modes free from contaminating foreground emission are effective probes of the 21cm power spectrum. The baseline-length dependence creates better isolation on the shortest baselines, giving rise to the "wedge" structure seen in Datta et al. (2010), Vedantham et al. (2012), Morales et al. (2012), Trott et al. (2012), and P12b; I refer the reader to these works for more detailed and alternative derivations of the "wedge."

4.3.1 Delay Space CLEAN

Some of the practical aspects of implementing the delay-spectrum approach in actual data were described in §3 of P12b. Of particular importance is the implementation of the frequency Fourier transform using a window function and the 1D-CLEAN algorithm first presented in Parsons & Backer (2009) to reduce the effects of RFI flags and band-edge effects. Even if foregrounds are spectrally smooth and easily localized in delay/$k_\parallel$-space, such sharp edges in frequency space will introduce significant covariance in delay space, resulting in the scattering of foreground emission into otherwise uncontaminated regions of k-space. The 1D-CLEAN algorithm treats RFI flags as a sampling function in frequency space, and "fills in" these gaps by iteratively fitting the brightest Fourier components in the delay domain. The end result is a model of the data which is free of RFI flagging gaps. I then form power spectra both of this model and the residuals between it and the raw data, before re-combining them. By separating the two components before the Fourier transform, one reduces the amount of power that can scatter off RFI gaps, minimizing bright sidelobes which would otherwise contaminate the EoR window in the power spectral domain.

I force the foreground model to be smooth-spectrum in frequency by only allowing delay-space components which fall inside a baseline-dependent area (a "CLEAN box"), 50 ns beyond the physical maximum horizon delay on that baseline. This extra 50 ns allows the algorithm to model foreground emission pushed beyond the horizon limit. For this analysis, 50 ns appears to encompass enough foreground emission that the sidelobes of any remaining flux scattering off RFI gaps in the residuals are below the noise. Since our results do detect additional emission beyond our chosen cut-off, its exact value may need to be revisited in future analyses with more sensitivity.

4.3.2 Power Spectra

Once the data have been CLEANed, I form power spectra on a visibility-by-visibility basis. The power spectrum estimates follow from equation (12) of Parsons et al. (2012a):

$$\widehat{P}(k) \approx \tilde{V}_{21}^2 \left(\frac{\lambda^2}{2k_{\mathrm{B}}}\right)^2 \frac{X^2 Y}{\Omega B}, \tag{4.1}$$

where λ is the observing frequency, k_B is Boltzmann's constant, Ω is the solid angle of the primary beam,[4] B is the observing bandwidth, X and Y are cosmological scalars which convert observed angles and frequencies into $h\text{Mpc}^{-1}$, and $\tilde{V}$ is a delay transformed visibility. I avoid introducing a noise-bias by forming our estimator $\tilde{V}_{21}^2$ from adjacent time samples on a baseline. The 10-second interval between integrations is short enough that both measurements can be considered redundant samples of the same k-modes.

The isotropy of the universe allows one to then combine all power spectrum estimates $\tilde{V}_{21}^2$ in annuli of equal $k_\perp$ to form a 2D power spectrum in the $(k_\perp, k_\parallel)$-plane. Note that the method used here does not take advantage of any coherent integration within a uv-pixel. Since foreground emission dominates the observed signal, the loss of sensitivity is not important, and we ignore this effect for computational efficiency. For EoR science runs, however, PAPER explicitly focuses on maximizing the sensitivity boost from coherent integration, using "maximum redundancy" configurations designed to sample select uv-pixels for long periods of time.

4.4 Results

To make power spectra of the 4-hour dataset described above, I first run the CLEAN algorithm over the full 100 MHz band. Using the whole band gives the best resolution in delay space, and since foregrounds are nearly coherent over the whole band, the additional information gives CLEAN the most signal-to-noise to build its model. For the cosmological delay transform to make power spectra, I use only a 25 MHz band from 140 to 165 MHz. This smaller band still exceeds the ~ 8 MHz band over which the $z \sim 8$ universe can be treated as coeval (Furlanetto et al. 2006). Since the main purpose of this analysis is to understand the k-space behavior of foregrounds, I ignore this effect, as the additional bandwidth gives us better k-space resolution. For an EoR analysis, the isolation of foregrounds will be slightly worse than shown here, owing to the necessarily poorer resolution.

Forming individual power spectra from each baseline of the array and binning in annuli of constant $|k_\perp|$ yields the 2-dimensional $P(k_\perp, k_\parallel)$ shown in Figure 4.3. The $k_\perp$-axis is binned with a resolution of $1.38 \times 10^{-4} h\text{Mpc}^{-1}$; gaps are visible (especially at small $k_\perp$-values) where there are no baselines of that length. I do no binning in $k_\parallel$; this resolution is set by the 25 MHz bandwidth used in the analysis.[5] The most prominent feature is the "wedge"-like shape of the foreground emission as predicted. As argued in P12b, this "wedge" footprint in k-space is not a result of imperfect calibration in foreground removal (since we attempt no foreground removal in this work), but a property of the emission itself as measured by an

[4] Derivations of equation 4.1 in Morales (2005), McQuinn et al. (2006), and Parsons et al. (2012a), relate ΩB to an effective cosmological volume, using a top-hat primary beam or effective area as a pedagogical tool to simplify the result. More generally, however, the effective Ω in equation 4.1 is actually $\int A^2(\theta,\phi)\, \mathrm{d}\theta\, \mathrm{d}\phi$, where $A(\theta,\phi)$ is the power response of the primary beam. A derivation of this effect is presented in Parsons et al. (2013). For PAPER, this beam is a factor of ~ 2 smaller than $\int A(\theta,\phi)\, \mathrm{d}\theta\, \mathrm{d}\phi$.

[5] Because of the Blackman-Harris window function used in CLEAN, only every other sample plotted in $k_\parallel$ is statistically independent.

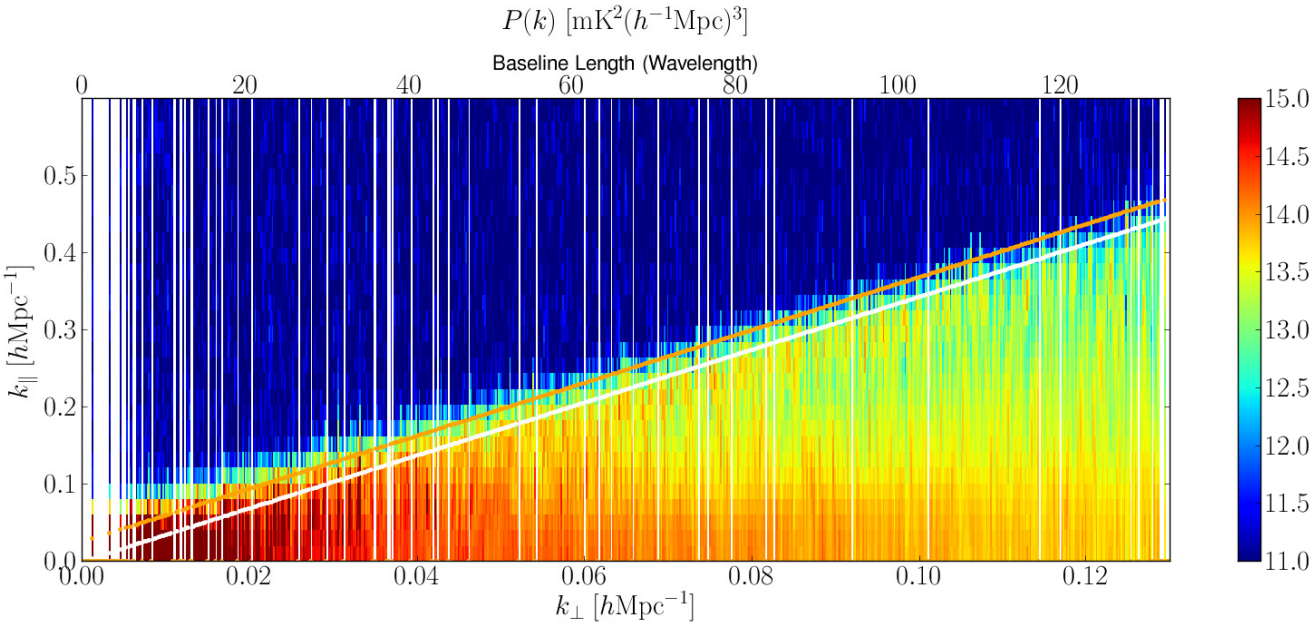

Figure 4.3: A two-dimensional power spectrum of the 4 hours of data analyzed. The wedge-like nature of the foreground emission is clear. The white line marks the horizon limit and the orange line is 50 ns beyond. The colorscale is logarithmic and the units are mK^2 $(h^{-1}\text{Mpc})^3$. The binning is described in the text.

interferometer.

The white diagonal line in Figure 4.3 corresponds to the horizon-limit in $k_\parallel$ for a baseline of corresponding length $k_\perp$; the orange line is 50 ns beyond the horizon, inside of which I allowed Fourier components in the deconvolution described in §4.3.1. As predicted, emission extends beyond the horizon limit due the spectral structure of the foreground emission.

I draw attention to the fact that the supra-horizon emission does not have a constant width in $k_\parallel$ as a function of $k_\perp$. Rather, more emission extends beyond the horizon on the smallest $k_\perp$-values (i.e. the shortest baselines). One expects this behavior to result from two different effects. First, the shortest baselines will resolve out less of the diffuse Galactic synchrotron, so that the emission will be brighter. Therefore, we can see its sidelobes extend further in $k_\parallel$ before they fall below the noise level. The second effect is somewhat more subtle and can be best illustrated with an example. Consider two east-west baselines of length 10 and 100λ at 150 MHz, which correspond to light travel times (i.e. horizon limits) of 66.7 and 667 ns, respectively. A point source 20° above the eastern horizon corresponds to geometric delays of 62.6 and 626 ns on these baselines. If the source spectrum has a given amount of "unsmoothness" and creates a delay space kernel of width 10 ns, then this kernel, centered at the geometric delay of the source, will lead to emission beyond the horizon on the 10λ baseline, but not on the 100λ baseline. This example illustrates how the same sources of emission will naturally lead to more corruption of supra-horizon delays on shorter baselines than on longer ones.

Finally, I also draw attention to the "edge brightening" of the foreground emission in the wedge on the longest baselines as one moves near the horizon limit. This feature can be attributed to the Galactic plane, and moves as expected when the data are viewed as a function of time.

To highlight the steepness of the foreground roll-off, I plot 1-dimensional $k_\parallel$ power spectra for bins of several baseline lengths in Figure 4.4. One sees that the foreground emission can fall by as much as three to four orders of magnitude in a factor of 2 change in $k_\parallel$. It is difficult to explicitly compare this result to the predictions of P12b, due to the different resolutions and binning used. The placement of bin edges can significantly complicate comparison when the fall-off is so steep, as a slight shift in the bin can result in a large change in the average value within. For similar reasons, it is difficult to say exactly where the foreground fall-off falls below the noise. Given these caveats, there is nothing in these data to contradict the predictions of P12b. It is clear, however, that sensitivities will have to increase several more orders of magnitude before anything can be said about the behavior of foreground emission at the tens of mK^2 level where the expected EoR signal lies.

To both demonstrate the effectiveness of the delay-space foreground isolation and to further investigate the nature of the supra-horizon emission, I high-pass filter the data in delay-space, selecting only delay modes more than 50 ns beyond the horizon limit, i.e., I select the emission lying beyond the orange line in Figure 4.3. We then image this data from 140 to 165 MHz in one-hundred 0.25 MHz bins to form a data-cube versus frequency. Finally, I Fourier transform our data-cube versus frequency to create maps of individual $k_\parallel$ modes. Three of the resultant maps for $k_\parallel = 0.06, 0.08$, and $0.51\ h\mathrm{Mpc}^{-1}$ are shown in Figure 4.5.

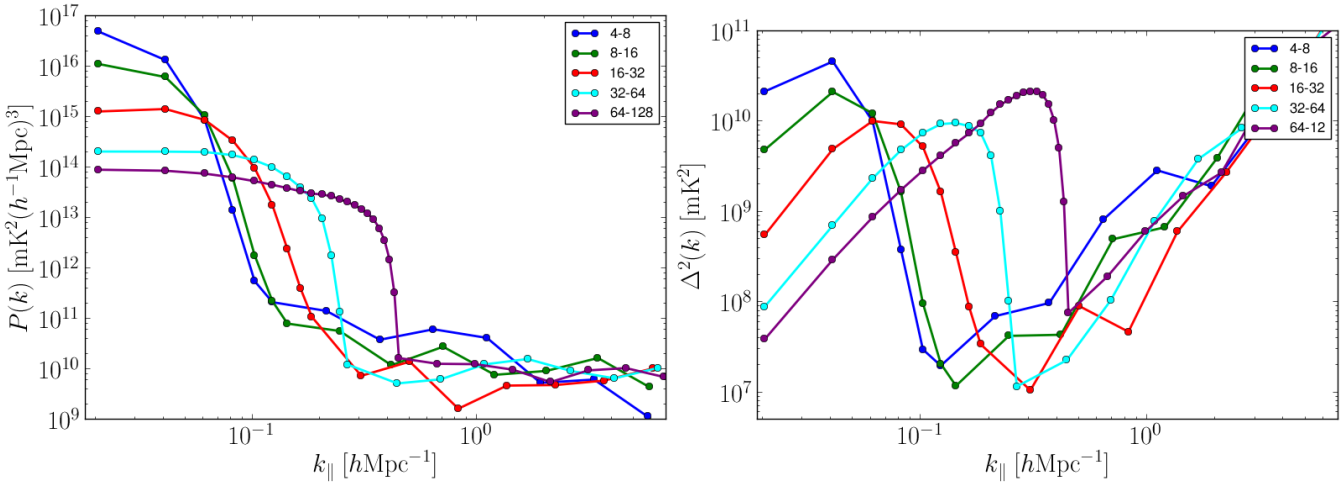

Figure 4.4: *L*eft: A 1-dimensional power spectrum versus $k_\parallel$ for bins of several baseline lengths. To preserve the steep roll-off of the foreground emission, the data are plotted with their natural resolution at lower $k_\parallel$ values, and logarithmically binned at higher values. *R*ight: The same as the left, but the dimensionless power spectrum $\Delta^2(k) = \frac{k^3}{2\pi^2}P(k)$.

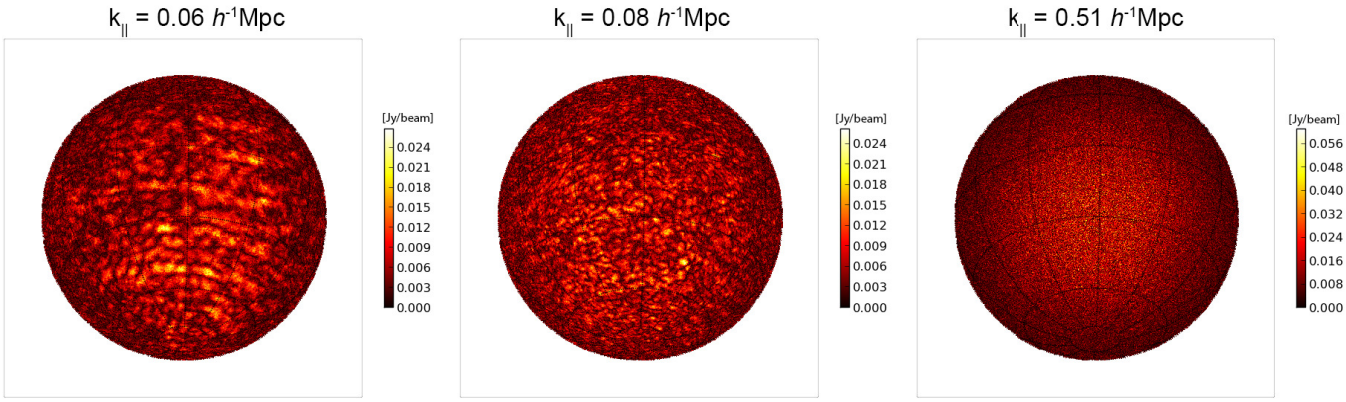

Figure 4.5: Images of the magnitude of select $k_\parallel$ modes in the data, after a high-pass delay space filter has removed all emission interior to 50 ns beyond the horizon limit (i.e. emission below the orange line in Figure 4.3 has been removed). The color-scale is linear in Jy, although the flux scale on the lowest $k_\parallel$-modes is compromised due to the filter. While emission is clearly present in the low $k_\parallel$ modes, it cannot be identified with the bright sources in Figure 4.2, nor are the features obviously correlated with Galactic structure.

The flux scale in these images has been reduced by 2 to 3 orders of magnitude from Figure 4.2, demonstrating the effectiveness of delay-space filtering in isolating foreground emission. Interpreting the flux scale in the lowest $k_\parallel$ modes is complicated, since most of baselines at these modes have been filtered off, whereas, the $k_\parallel = 0.51 h\text{Mpc}^{-1}$ mode is effectively a noise map. It can be seen in Figure 4.3 that this mode lies above the wedge, and therefore nothing has been filtered off from it. The RMS in this map (calculated from the complex data) is 14 mJy. One can estimate the expected noise level using:

$$\Delta\sigma = \frac{2k_B\Omega}{\lambda^2}\frac{T_{\text{sys}}}{\sqrt{N(N-1)Bt}}, \tag{4.2}$$

where k_B is Boltzmann's constant, Ω is the solid angle of the primary beam, λ is the observing wavelength, T_{sys} is the system temperature, N is the number of antennas in the array, B is the bandwidth, and t is the observing time (Thompson et al. 2007). For this observation $\Omega = 0.75$ sr, $t = 4$ hours, $\lambda = 1.96$ m, $B = 25$ MHz, and $T_{\text{sys}} = 1000$ K. This somewhat high value for T_{sys} is reasonable given the Galactic emission in these observations; a more thorough discussion of T_{sys} for PAPER is presented in Parsons et al. (2013). Using these values gives an expected RMS of 14 mJy, in accord with the measurement.

The $k_\parallel = 0.06$ and 0.08 $h\text{Mpc}^{-1}$ maps are clearly not noise dominated. Accounting for the noise that was removed by filter, these data have rough "effective" RMS values of 35 mJy, well in excess of the noise estimate. Given the presence of emission in these modes above the orange line in Figure 4.3, it is not surprising that one sees excess power. What is surprising is that none of this emission can be easily associated with the brightest sources or structures visible in Figure 4.2. Rather, it appears that the bulk of the emission contaminating the EoR window comes from an aggregate of fainter emission. However, I would caution the reader against interpreting these maps as images of true coherent structures on the sky. The fact that the spatial pattern of emission changes significantly from $k_\parallel = 0.06$ to 0.08 $h\text{Mpc}^{-1}$ suggests the emission is a diffuse background and the images are limited by sidelobes.

I note one additional interesting feature in these maps. The noise-dominated $k_\parallel = 0.51$ $h\text{Mpc}^{-1}$ gives a good impression of the PAPER primary beam shape. In the 0.06 and 0.08 $h\text{Mpc}^{-1}$ maps, however, bright emission extends well beyond the half-power point of the beam (roughly 45° FWHM). Under the lens of the delay transform formalism, one might argue that the emission closest to the horizon can most easily create supra-horizon emission. These images show that while emission from well outside the beam FWHM contributes significantly to the supra-horizon emission, the PAPER primary beam roll-off is enough to keep this pattern from extending all the way out the edge of the image.

4.5 Conclusions

I have presented new observations from PAPER measuring the properties of foreground emission in cosmological Fourier space. These observations have confirmed general predictions presented in, e.g., Datta et al. (2010), Morales et al. (2012) and P12b: that foreground

emission occupies a "wedge" in the 2D $(k_\perp, k_\parallel)$ plane, leaving a window at higher $k_\parallel$ values for 21cm EoR studies. I have also confirmed that shorter baselines yield a larger window onto the cosmological signal. However, this the window does not grow perfectly linearly with decreasing baseline length. Therefore, while shorter baselines do make the best probes of the EoR signal, there will be diminishing returns at the shortest baselines. I have also presented an images of several $k_\parallel$ modes of "unsmooth" emission extending past the nominal edge of the wedge. These images are unable to localize any of the emission to known sources on the sky, suggesting that the most problematic foregrounds for EoR observations are a diffuse background.

Acknowledgments

A version of this chapter was previously published as Pober et al. 2013 ApJL XXX XXX, and is reproduced with the permission of all coauthors and the copyright holder.

Chapter 5

The Baryon Acoustic Oscillation Broadband and Broad-beam Array: Design Overview and Sensitivity Forecasts

This chapter describes a new instrument optimized for a detection of the neutral hydrogen 21cm power spectrum between redshifts of $0.5 - 1.5$: the Baryon Acoustic Oscillation Broadband and Broad-beam (BAOBAB) Array. BAOBAB will build on the efforts of a first generation of 21cm experiments which are targeting a detection of the signal from the Epoch of Reionization at $z \sim 10$. At $z \sim 1$, the emission from neutral hydrogen in self-shielded overdense halos also presents an accessible signal, since the dominant, synchrotron foreground emission is considerably fainter than at redshift 10. The principle science driver for these observations are Baryon Acoustic Oscillations in the matter power spectrum which have the potential to act as a standard ruler and constrain the nature of dark energy. BAOBAB will fully correlate dual-polarization antenna tiles over the 600–900MHz band with a frequency resolution of 300 kHz and a system temperature of 50K. The number of antennas will grow in staged deployments, and reconfigurations of the array will allow for both traditional imaging and high power spectrum sensitivity operations. We present calculations of the power spectrum sensitivity for various array sizes, with a 35-element array measuring the cosmic neutral hydrogen fraction as a function of redshift, and a 132-element system detecting the BAO features in the power spectrum, yielding a 1.8% error on the $z \sim 1$ distance scale, and, in turn, significant improvements to constraints on the dark energy equation of state over an unprecedented range of redshifts from $\sim\ 0.5 - 1.5$.

5.1 Introduction

Measurements of the Baryon Acoustic Oscillation (BAO) features in the matter power spectrum are one of the most promising probes for constraining dark energy. So far, these features have been detected with spectroscopic and photometric galaxy surveys, but 21cm measurements in the $z \sim 1-3$ range over a promising complementary or alternative approach. See Chapter 1 for a fuller discusison of BAO and the state of current measurements.

The flexibility in angular and spectral responses of radio interferometers, which measure the power spectrum both parallel and perpendicular to the line of sight (Morales 2005), gives 21cm BAO the ability to survey larger cosmological volumes and operate over a wider range of redshifts than current spectroscopic galaxy redshift surveys. As a result, the 21cm BAO signal has the potential for probing expansion throughout and beyond the critical epoch when dark energy comes to dominate the energy density of the universe. Furthermore, a 21cm intensity mapping experiment can probe redshifts $z > 0.5$ with roughly uniform sensitivity, without complications arising from sky emission lines in the optical/near-infrared. The 21cm signal can be used to constrain the location of the BAO peaks as a function of redshift, and thereby measure the magnitude and time-evolution of dark energy.

21cm BAO experiments can draw on the considerable investments in low-frequency radio astronomy developed in the past decade for studies of the Epoch of Reionization (EoR). In this chapter, I present the Baryon Acoustic Oscillation Broadband and Broad-beam (BAOBAB) array, a new experiment, building on the legacy of the Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2010)[1] and the Murchison Widefield Array (MWA; Lonsdale et al. 2009)[2], for measuring the 21cm HI power spectrum at a redshift of $z \sim 1$. This chapter is structured as follows: in §5.2, I present a system architecture for the BAOBAB instrument. In §5.3, I forecast the sensitivity and cosmological constraints that will be achieved by BAOBAB. I will consider several possible challenges and extensions for this approach in §5.4, and conclude in §5.5. Throughout this work I assume the WMAP7 best fit ΛCDM cosmological model: $h = 0.7$, $\Omega_M = 0.27$, $\Omega_b = 0.046$, $\Omega_{DE} = 0.73$, and $\sigma_8 = 0.8$ (Larson et al. 2011).

5.2 The BAO Broadband and Broad-beam Array

The past decade has seen significant progress in the design, construction, and calibration of low-frequency interferometric arrays toward the goal of detecting the highly-redshifted 21cm signal from the Epoch of Reionization. The technologies used in BAOBAB inherit from two EoR experiments — PAPER and the MWA — but with several significant modifications to optimize the instrument for BAO science. The entire signal chain will be re-tuned to operate between 600–900 MHz. These frequencies corresponds to redshifted 21cm emission between $z = 0.58 - 1.37$, a band chosen for several reasons. First, these moderate redshifts

[1] http://eor.berkeley.edu/
[2] http://www.mwatelescope.org/

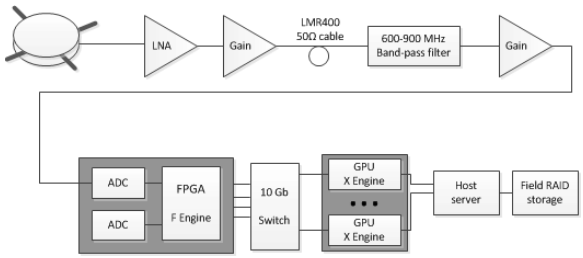

Figure 5.1: System diagram of the BAOBAB interferometer. Dual-polarization antenna signals at -103dBm enter an uncooled low-noise amplifier (LNA) with +12.5dB gain and a noise figure of 0.4dB (30K). Second-stage amplifiers add +36dB gain before transmission through 30 meters of LMR400 50Ω cable to a central enclosure. The signal is bandpass filtered (600–900MHz) and amplified +40dB to the optimal -22dBm input level for the ADCs. Each antenna signal is digitized and channelized in ROACH F-engines, reordered in transmission through a 10Gb Ethernet switch, and sent to GPU GTX580 X-engines for cross-correlation. Raw visibility data are passed to a host computer and written to a RAID storage unit in the MIRIAD UV file format for post-processing.

complement the undergoing lower redshift galaxy surveys like BOSS by probing what is currently a relatively unexplored volume of the universe. Secondly, at these frequencies, commercially available amplifiers and cables provide suitable low-noise performance that would not be obtainable at higher frequencies (lower redshifts). Furthermore, this band also avoids the bright sky noise and ionospheric effects that complicate lower frequency (higher redshift) observations.

BAOBAB will be a non-phase-tracking, broadband array of beam-formed tiles. By lowering the field-of-view of each element, the use of beamformed tiles like the MWA will significantly increase BAOBAB's power spectrum sensitivity over that of an array of the equivalent number of single dipoles, without increasing the correlator demands. Each tile will consist of 4 scaled versions of the PAPER sleeved-dipole design and groundscreens, electronically beamformed to point to zenith. Two linear polarization signals from each tile enter a digital signal processor that computes both auto- and cross-correlation products and outputs the results locally to disk. A block diagram is given in Fig. 5.1; the key properties of the BAOBAB system are listed in Table 5.1 and are described in more detail in the remainder of this section.

5.2.1 Siting

Since radio-frequency interference (RFI) is prevalent at these frequencies, BAOBAB will need to be located at a radio quiet site. The bottom panel of Figure 5.2 shows preliminary measurements made by a prototype 2-antenna BAOBAB interferometer deployed at the Leuschner Observatory near Berkeley, CA. At this site, only 40 MHz of a 400–800 MHz operating band show solar fringes uncorrupted by RFI, demonstrating the need for the primary BAOBAB deployment to be located at a quieter site, such as the NRAO site near

Table 5.1: Proposed BAOBAB Array

Operating Bandwidth	600–900 MHz
Number of Tiles	32–128
Collecting Area per Element	2.6 m^2
Gain per Element	18 dBi
Field-of-View	0.045 str
Receiver Noise Temperature	40 K
System Temperature	50 K
Maximum Imaging Baseline	60 m
Redundant Baseline Scale	1.6 m
$k_{\rm min}$, $k_{\rm max}$	0.01, 2.5 hMpc^{-1}
Array Configuration	Reconfigurable (see Fig. 5.6)
Frequency Resolution	300 kHz
Snapshot Integration Time	10 s

Green Bank, WV. Next-generation activities may take place at the Square Kilometer Array South Africa (SKA-SA) reserve in the Karoo desert. This site is currently occupied by the PAPER and MeerKAT arrays, and has been shown to be a pristine RFI environment (Jacobs et al. 2011).

5.2.2 Analog System

With the drastic reduction in sky noise relative to EoR frequencies, BAOBAB's system temperature will be dominated by the analog electronics. These components must therefore be optimized to reduce receiver noise while maintaining the smooth spatial and spectral responses that are a hallmark of the PAPER design and a key component of the delay spectrum foreground isolation approach presented in Parsons et al. (2012b) (hereafter P12b) and discussed in Chapters 1, 4, and §5.3.3. The analog system will include the collecting element (consisting of 4 antennas and reflectors), low-noise amplifier, coaxial cable, and receiver.

The BAOBAB element will begin with a 1/5-scale PAPER antenna (Parsons et al. 2010), as shown in Figure 5.3. This design is a dual-polarized version of the sleeved dipole design that uses a twin-resonance structure consisting of a pair of crossed dipoles located between a pair of thin aluminum disks. The element's reliability has been demonstrated in PAPER arrays over the past several years. A trough reflector under each dipole will be used to increase the directivity toward zenith. The electromagnetic behavior of the element was modeled extensively for PAPER using CST Microwave Studio, and shown to perform as desired through calibration with celestial sources in Chapter 3. The geometrically re-tuned prototype shown in the top panel of Figure 5.3 will be optimized to operate efficiently over the 600–900 MHz band.

Rather than deploy single elements like PAPER, BAOBAB will use a 2×2 tile of dipoles

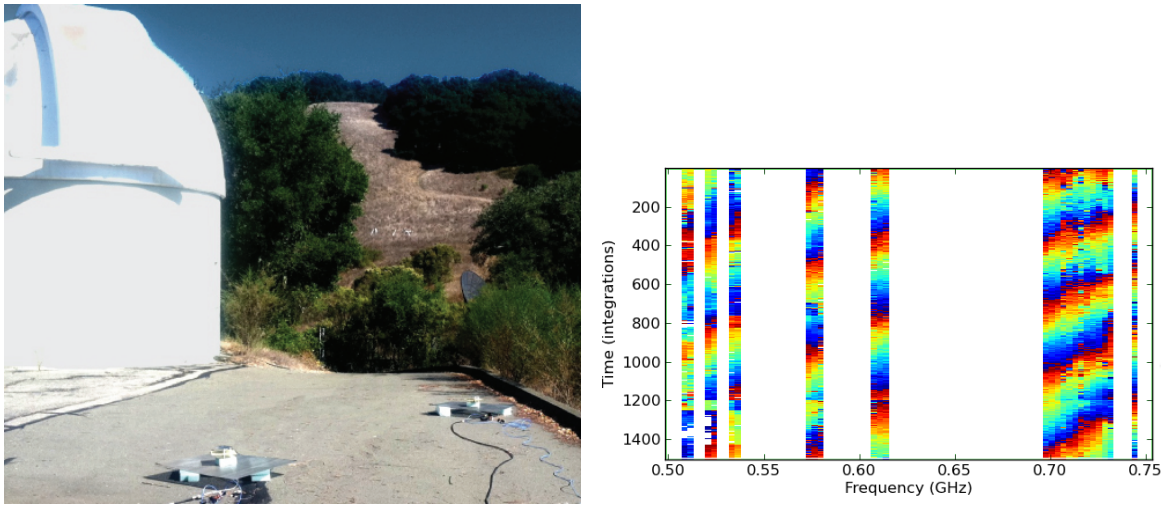

Figure 5.2: *Left*: Leuschner Observatory, with a prototype 2-element BAOBAB interferometer deployed. This system was developed and deployed by students as part of a *Fundamentals of Radio Astronomy* class at UC Berkeley. *Right*: solar fringes measured with the BAOBAB-2 prototype at Leuschner.

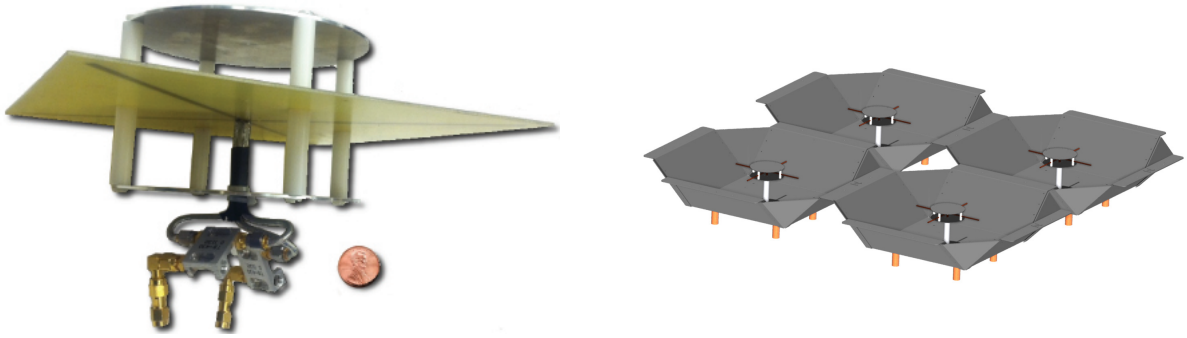

Figure 5.3: *Left*: A prototype BAOBAB dipole antenna, designed as a 1/5 scale model of a PAPER dipole. *Right*: BAOBAB tile design with 4 dipoles and individual ground-screens.

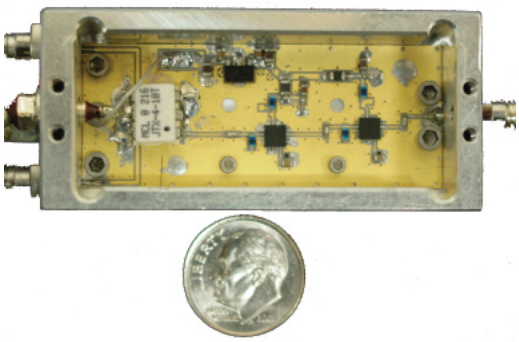

Figure 5.4: The prototype balun and Hittite HMC617LP3 LNA for BAOBAB. The amplifier adds +30dB of gain with a quoted noise figure of 0.5 dB.

and ground-screens, as shown in Figure 5.3. A fixed zenith beamformer will be used to combine the signals, increasing the gain by 6 dB and reducing the field-of-view by a factor of four. Both analog and digital beamformers are being investigated. A key issue is the mutual coupling, which should be reduced by the additional groundscreens between dipoles. The net effect is that for a fixed correlator size, the power-spectrum sensitivity is increased by a factor of four (see §5.3.2).

The amplifier designed for PAPER has a measured noise temperature of 110 K with 30 dB of gain across the 120-170 MHz band (Parsons et al. 2010). For application to BAO at $z \sim 1$, we will modify this amplifier design to operate from 600–900 MHz. Besides re-tuning the filter and amplifier circuits, however, one of the major activities in this modification will be to reduce the noise temperature of the front-end amplifier in order to obtain a target system temperature of 50 K. This change reflects one of the key differences between the BAO and EoR foregrounds. System noise in the EoR band is dominated by ∼300 K sky noise from galactic synchrotron emission. In the BAO band, the sky temperature is reduced to ∼10 K, making the front-end amplifier the leading source of noise. Uncooled commercial UHF-band amplifier transistors based on GasFET or HEMT technology can reliably achieve noise figures of 0.4 dB, corresponding to a receiver temperature of 30K. A prototype BAOBAB balun/amplifier using a Hittite HMC617LP3 LNA with a quoted noise figure of 0.5 dB is shown in Figure 5.4; tests are underway to determine the noise temperature of the complete system.

5.2.3 Digital System

The BAOBAB correlator will follow the scalable correlator design used by PAPER and other members of the international Collaboration for Astronomy Signal Processing and Electronics Research (CASPER)[3], a real-time digital correlator based on Field-Programmable

[3] https://casper.berkeley.edu

Figure 5.5: PAPER's 128-input correlator (shown) follows the packetized frequency–cross-multiply (FX) architecture developed by the Center for Astronomy Signal Processing and Electronics Research (CASPER). Shown are 16 ROACH F-engines (left) and 2 dual-GPU box X-engines (right). The first-generation BAOBAB correlator will modify the 64-input, 100-MHz PAPER correlator to become a 32-input, 300-MHz correlator. It will employ eight ROACH boards for spectral processing and four dual-GPU boxes for cross-multiplication. A 10-GbE switch is used to route data between boards.

Gate Array (FPGA) processors and graphics processing units (GPUs) (Parsons et al. 2008; Clark et al. 2011). The correlator architecture we employ uses modular signal processing electronics and packetized communication protocols to build correlators that are flexible in the number of antennas correlated and the bandwidth correlated by each antenna. A photograph of a 128-input FPGA/GPU correlator is shown in Fig. 5.5.

The generic FX correlator architecture implemented consists of modules responsible for digitizing and channelizing each set of antenna inputs (F-Engines), followed by a set of signal processing modules responsible for cross-multiplying all antennas and polarizations for a single frequency (X-Engines) and accumulating the results. Unique to this architecture, signal processing engines transmit packetized data through commercially available 10 Gb Ethernet switches that are responsible for routing data between boards. This architecture, along with analog-to-digital converters, modular FPGA-based signal processing engines, and a software environment for programming, debugging, and running them, were developed in collaboration with CASPER at the University of California, Berkeley (Parsons et al. 2008). The flexibility and modularity of this correlator design shortens development time, and in this case, allows an existing 64-input, 100-MHz PAPER correlator with 8 ROACH-boards and 4 dual-GPU boxes to be straight-forwardly modified to become a 32-input, 300-MHz BAOBAB correlator using the same boards and signal processing libraries. A forthcoming publication on this correlator will be presented by Ali et al.

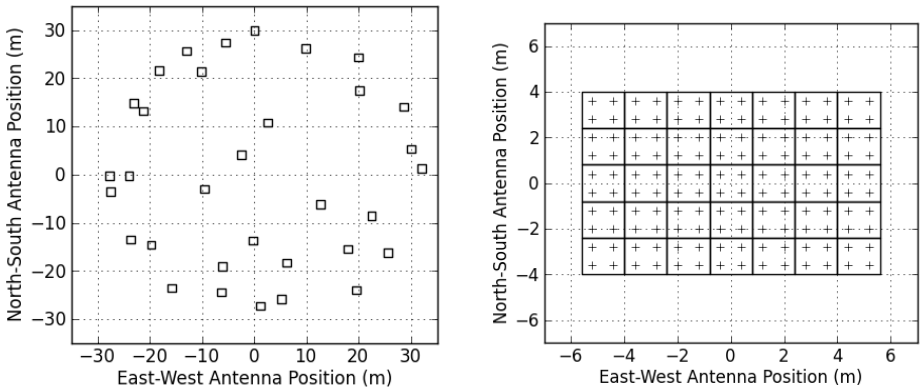

Figure 5.6: BAOBAB array configurations, plotted in meters east-west (horizontal axis) and north-south (vertical axis). BAOBAB will use above-ground cabling to allow antennas to be moved into a minimum-redundancy configuration (left) for imaging foregrounds, or a maximum-redundancy configuration (right) for enhanced power-spectrum sensitivity. Each square represents one tile, and each "+" one dipole.

5.2.4 Configuration

BAOBAB will employ small antennas and above-ground cabling with relatively inexpensive LMR400 50-Ohm coaxial cables; these cables will not be buried, allowing BAOBAB to easily change between different array configurations by moving antenna elements. Following the principles outlined in (Parsons et al. 2012a, hereafter P12a), BAOBAB will employ a minimum-redundancy imaging configurations for characterizing foregrounds with minimal sidelobes and maximum-redundancy configurations to repeatedly sample the same locations in the *uv*-plane, substantially improving sensitivity to the three-dimensional power spectrum of 21cm emission at $z \sim 1$. Although future experiments may target a range of angular scales to map 21cm emission in the plane of the sky, by focusing on a limited number of Fourier modes, these maximum-redundancy configurations can improve sensitivity to the power spectrum by an order of magnitude or more in mK^2, relative to an equivalent observation with a minimum-redundancy configuration.

However, as will be discussed further in §5.3.2, the mapping of baseline length to a transverse k-mode on the sky is significantly larger for BAOBAB than for PAPER. In order to probe the relatively large-scale BAO features, then, BAOBAB will use the most compact configurations possible for its power spectrum measurements. Such an array configuration for a 35-element system, as well as that of a 32-element imaging configuration, are shown in Figure 5.6. The tiles are spaced 1.6m apart, effectively touching end-to-end. Investigations of cross-talk and mutual coupling will take place during an early prototype of the system; it may be the case that a phase-switch or additional shielding between tiles will be necessary to accommodate the short baselines required by BAO science.

With a modular CASPER correlator increasingly able to process larger numbers of

antenna inputs, BAOBAB naturally lends itself to a staged approach. Early ≤ 16 tile-element prototypes will characterize system performance, while a subsequent ~ 32-element array will study foreground emission and constrain the neutral hydrogen fraction as a function of redshift with a measurement of the 21cm power spectrum (§5.3.4). A ~ 128-tile version of BAOBAB will measure BAO features and provide substantial improvements over our current constraints on the equation of state and time evolution of dark energy (§5.3.5).

5.3 Predicted Cosmological Constraints from BAOBAB

In this section I present predictions for forthcoming cosmological constraints for several iterations of the BAOBAB instrument. I begin by reviewing the predicted signal strength for the cosmological 21cm power spectrum in §5.3.1. In §5.3.2, I adapt the power spectrum sensitivity calculations of P12a for an array operating at $z \sim 1$, including the effects of sample variance and shot noise. In 5.3.3, I briefly review the delay spectrum foreground removal procedure of P12b. While a detailed study of foregrounds is beyond the scope of this paper, it is worthwhile to discuss the implications of the technique on which Fourier modes of the 21cm power spectrum will be accessible. I conclude the section by presenting forecasts for a 35- and 132-element BAOBAB system in §5.3.4 and §5.3.5 respectively, including Fisher matrix predictions for constraints on the dark energy equation of state in the latter section. In the discussion of §5.4, I explore possible directions for improvement with larger BAOBAB arrays.

5.3.1 The 21cm Power Spectrum

As with galaxy redshift surveys, a 3D map of the neutral hydrogen in the universe will serve as a tracer of the underlying dark matter power spectrum. The brightness of the observable radio 21cm signal will depend on the cosmological neutral hydrogen fraction, as well as the bias of hydrogen containing halos with respect to the dark matter (Barkana & Loeb 2007; Madau et al. 1997; Ansari et al. 2012):

$$P_{T_{21}}(k,z) = \left[\tilde{T}_{21}(z)\right]^2 b^2 P(k,z), \tag{5.1}$$

$$\tilde{T}_{21}(z) \simeq 0.084\text{mK}\frac{(1+z)^2 h}{\sqrt{\Omega_m(1+z)^3+\Omega_\Lambda}}\frac{\Omega_B}{0.044}\frac{f_{\rm HI}(z)}{0.01}, \tag{5.2}$$

where $\tilde{T}_{21}(z)$ is the mean 21cm brightness temperature at redshift z; $P(k,z)$ is the linear matter power spectrum; b is the bias factor of HI containing halos with respect to the dark matter; $f_{\rm HI}(z)$ is the mass fraction of neutral hydrogen with respect to the overall cosmological baryon content (i.e., $\Omega_{\rm HI} = f\Omega_b$); Ω_Λ is the cosmological constant, and Ω_m and Ω_B are the matter and baryon density in units of the critical density, respectively.

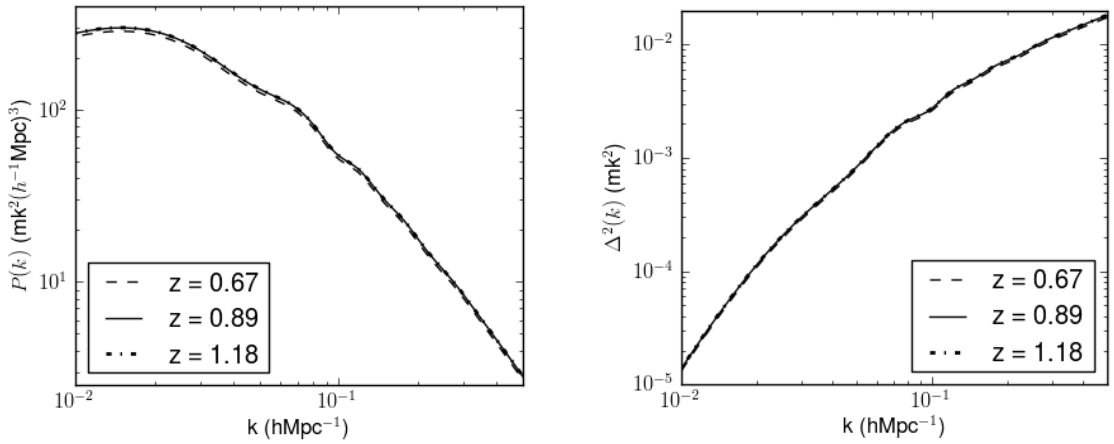

Figure 5.7: *Left*: The predicted real space (i.e., $\mu = 0$) power spectrum of 21cm emission in the fiducial model at three redshifts: 0.67, 0.89 and 1.18 (corresponding to frequencies of 850, 750 and 650 MHz). The $z = 0.89$ and $z = 1.18$ power spectra overlap, and, effectively, the signal is at the same strength for all three redshifts. Matter power spectrum predictions come from CAMB. *Right*: The equivalent dimensionless power spectra, $\Delta^2(k) = \frac{k^3}{2\pi^2}P(k)$.

I plot the predicted 21cm brightness temperature power spectrum, $P(k)$, for the fiducial model at redshifts of 0.67, 0.89 and 1.18 (corresponding to frequencies of 850, 750, and 650 MHz, respectively) in Figure 5.7. Predictions for the matter power spectrum come from CAMB (Lewis et al. 2000)[4]. I also plot the dimensionless power spectrum, $\Delta^2(k) = \frac{k^3}{2\pi^2}P(k)$. For the remainder of this chapter I will primarily express the results in terms of $\Delta^2_{T_{21}}(k)$, due to its more intuitive units (mK^2) and simple physical interpretation as the variance per logarithmic k bin.

In practice, the 21cm power spectrum is measured in redshift space, which can be, at linear order, related to the real-space $P_{T_{21}}(k, z)$ as (Kaiser 1987):

$$\tilde{P}_{T_{21}}(k, z) = (1 + \beta\mu^2)^2 P_{T_{21}}(k, z) \tag{5.3}$$

where $\mu = \hat{k} \cdot \hat{z}$ is the wavevector $\hat{k}$ projected along the line-of-sight $\hat{z}$, and $\beta \equiv f(\Omega)/b$ from linear theory where $f(\Omega) \approx \Omega_m(z)^\gamma$ is the dimensionless linear growth rate and $\gamma = 0.557$ for ΛCDM cosmologies. To incorporate the effects of redshift-space distortions in 1D plots of the power spectrum, I reduce our thermal noise error bars in k-space by the factor of $(1 + \beta\mu^2)^2$; however, the cosmological analyses keep full 2D information. I will not attempt to constrain cosmological parameters like β by measuring the power spectrum as a function of μ. I assume fiducial values of $f_{\mathrm{HI}} = 0.015$ and $b = 0.75$, chosen to agree with $f_{\mathrm{HI}}b = 0.012 \pm 0.003$ as measured by Chang et al. (2010). To obtain individual constraints on these parameters, one will need to measure redshift-space distortions themselves.

[4] http://camb.info/

5.3.2 Sensitivity of an Array to the 21cm Signal

There are three independent sources of statistical uncertainty in a 21cm power spectrum measurement: thermal noise in the interferometric visibilities, sample variance, and shot noise, of which the last is in some sense "signal," but still inhibits measurements of cosmological parameters. For the first-generation of 21cm experiments, thermal noise will be the dominant source of uncertainty in measurements of the power spectrum. I therefore calculate the effects of thermal noise independently in §5.3.2. I add in the effects of sample variance in §5.3.2 and argue in §5.3.2 that shot noise can be neglected for these experiments.

Thermal Noise

Thermal noise, in addition to being the dominant source of uncertainty in first generation 21cm BAO experiments, is also likely to be least familiar to those used to optical redshift surveys. Given the limited collecting area of early experiments, reducing thermal noise contributions is of paramount importance, even at the expense of the number of Fourier modes measured. Much of the work in this section explicitly follows the derivation of an interferometric array's power spectrum sensitivity presented in P12a (only the prefactors have changed to account for different fiducial values at $z = 1$). The approach of P12a is to treat each baseline as an independent probe of the 21cm power spectrum at one $k_\perp$ set by the baseline length, and many $k_\parallel$-values along the line-of-sight. I therefore first derive the sensitivity of a single-baseline; the sensitivity of the array is the aggregate sum of all these independent measurements.

I begin with a version of equation 16 from P12a, which gives the power spectrum of the thermal noise obtained from one integration of a dual-polarization baseline:

$$\Delta^2_{\rm N}(k) \approx X^2 Y \frac{k^3}{2\pi^2} \frac{\Omega}{2t} T^2_{\rm sys}, \tag{5.4}$$

where X^2Y is a scalar translating observed units to cosmological distances in h^{-1}Mpc (X converts from angles on the sky to transverse distances, Y from bandwidth to line-of-sight distance), Ω is the solid angle of the power-squared primary beam of one element in steradians (see footnote 4 in Chapter 4 and the discussion in Chapter 1), $T_{\rm sys}$ is the system temperature, and t is the amount of integration time of the sample. It is again worth emphasizing that this equation is not valid for all k-modes (which would imply a white noise power spectrum throughout Fourier-space), but rather only those modes sampled by the one baseline in question.

I also note that this relationship is very similar to equation 31 in Ansari et al. (2012). Quantitatively, however, they differ by a factor of 8. There are two separate effects contributing to this difference. Firstly, equation 5.4 is explicitly for a dual-polarization receiver, giving us twice as many independent samples of the same k-modes and therefore half the noise. Secondly, for the receiver design of the BAOBAB system, the RMS fluctuations in a measurement are given by $T_{\rm rms} = \frac{T_{\rm sys}}{\sqrt{2Bt}}$ (Kraus 1966; Thompson et al. 2007). Ansari et al. (2012) include

an additional factor of 4 in their equation (21), perhaps due to the design of their system including phase switching or some other effect.

The science goals of a BAO experiment are to actually measure X and Y; that is, since the exact values of X and Y depend on the underlying cosmology, one can combine the known physical scale of BAO with the angular and frequency scales in the observed signal to extract the detailed expansion history of the universe. For the purpose of a sensitivity derivation, however, the behavior of X and Y can be considered well enough known to compute fiducial values for equation 5.4.

X is related to the angular size distance, $D_{\rm A}$, as

$$X = D_{\rm A}(1+z) \equiv \int_0^z \frac{c\ dz}{H(z)}, \tag{5.5}$$

with $H(z)$ in the matter/dark-energy dominated epoch being approximately given by $H(z) = H_0\sqrt{\Omega_{\rm M}(1+z)^3 + \Omega_\Lambda}$. Numerical integration for a flat universe with $\Omega_{\rm M} = 0.27$, $\Omega_\Lambda = 0.73$, and $H_0 = 70$ yields $D_{\rm A}(z=1) \approx 1680$ proper Mpc (Wright 2006). Ignoring the evolution in the angular diameter distance around $z \sim 1$, one can write:

$$X \approx 1700(1+z)\frac{\rm Mpc}{\rm radian}. \tag{5.6}$$

Note that although I use this (admittedly somewhat poor) approximation for simplicity in deriving the relations for thermal noise power spectra in this subsection, all the subsequent results include the full z-dependence of the angular diameter distance.

A few more words are warranted concerning the magnitude of X at $z = 1$. Given the scaling of equation 5.5, a 16λ-baseline (6.8m a $z = 1$) corresponds to a transverse wavemode of $k_\perp = 0.042\ h{\rm Mpc}^{-1}$, a non-negligible value compared to the first BAO peak at $\sim 0.08\ h{\rm Mpc}^{-1}$. Therefore, baselines longer than $\sim 32\lambda$ will lose access to the first peak and be less effective probes of cosmology, regardless of foreground effects to be discussed later. This scaling motivates the extremely compact configurations proposed for BAOBAB in §5.2, despite the possible systematics associated with such short baselines.[5]

To compute the scaling between frequency, ν, and comoving line-of-sight distance, $r_{\rm los}$, one can use

$$dr_{\rm los} = \frac{cdz}{H(z)}. \tag{5.7}$$

Since $\nu(1+z) = \nu_{21}$, one has that $dz/(1+z) = -d\nu/\nu$, so

$$Y \equiv \frac{dr_{\rm los}}{d\nu} = \frac{c(1+z)^2}{\nu_{21}H(z)}. \tag{5.8}$$

[5]Epoch of Reionization experiments at $z = 9$, however, do not find themselves limited by the transverse modes probed. $X(z=9) \approx 9360\frac{\rm Mpc}{\rm radian}$ (P12a), so that a 16λ-baseline corresponds to $k_\perp = 0.015\ h{\rm Mpc}^{-1}$. With this scaling, the effect of the $k_\perp$ component on the measured power spectrum will always be sub-dominant to the foreground effects described in P12b.

Evaluating the above numerically, one gets

$$Y = 3.0 \frac{(1+z)^2}{\sqrt{\Omega_{\rm M}(1+z)^3 + \Omega_\Lambda}} \frac{\rm Mpc}{\rm MHz}. \tag{5.9}$$

Finally, one can compute the product X^2Y used in equation 5.4:

$$X^2Y \approx 2.93 \frac{(1+z)^4}{\sqrt{\Omega_{\rm M}(1+z)^3 + \Omega_\Lambda}} \frac{(h^{-1}{\rm Mpc})^3}{\rm sr \cdot Hz}. \tag{5.10}$$

Nominally, $X^2Y = 28(h^{-1}{\rm Mpc})^3\ {\rm sr}^{-1}\ {\rm Hz}^{-1}$ at $z = 1$.[6]

The other values in equation 5.4 are system-dependent parameters. The BAOBAB tiles have a considerably sized primary beam on the sky, $\Omega \approx 150{\rm sq.deg} \approx 0.045{\rm sr}$, yielding a power-squared beam of approximately twice this size.[7] However, this beam is significantly narrower than that of a single dipole, so that the use of beamforming results in a considerably lower noise level, since our noise power spectrum scales as Ω into equation 5.4.

The issue of system temperature is another instance where a BAO experiment at $z = 1$ is fundamentally different from an EoR experiment at $z = 9$. In the EoR case, Galactic synchrotron emission is the primary source of noise at 150 MHz, with a value of $\sim 300{\rm K}$ toward the galactic poles. However, synchrotron emission scales approximately as

$$T_{\rm sync} \approx 300{\rm K} \left(\frac{\nu}{150{\rm MHz}}\right)^{-2.5}. \tag{5.11}$$

At frequencies around 750 MHz, $T_{\rm sync} \approx 5{\rm K}$. This value is substantially below typical receiver temperatures of 50K, so that receiver temperatures will dominate $T_{\rm sys}$. As described in §5.2, BAOBAB will have a system temperature of approximately 50K across its entire band.

One can express the noise power spectrum of one dual-polarization baseline integrating on one Fourier mode by substituting these fiducial values in equation 5.4. However, before doing so, it is worthwhile to look ahead and emphasize that the remaining results in this section are intended to give both a quantitative sense of the level of thermal noise in $z \sim 1$ 21cm observations and scaling relations for the effect of various instrumental and observational parameters. In the sensitivity calculations of §5.3.4 and §5.3.5, however, I fully simulate the *uv*-coverage of our arrays, including the effects of earth-rotation synthesis. I then use

[6] The magnitude of X^2Y at $z = 1$ represents an often under-appreciated gain in sensitivity between a 21cm BAO experiment and a similar reionization experiment. At $z = 9$, $X^2Y \approx 540(h^{-1}{\rm Mpc})^3\ {\rm sr}^{-1}\ {\rm Hz}^{-1}$, meaning that the quoted EoR signal strength of $\sim 10{\rm mK}^2$ is normalized over a much larger volume. The smaller volume scalar at $z = 1$ means that over an order of magnitude less thermal noise is picked up per equivalent integration.

[7] This value for primary beam was estimated based on models of the PAPER single dipole primary beam and a simple array-factor to estimate a tile beam. The actual value will depend on the illumination pattern of the ground screen flaps by the dipole. If prototype systems substantially under-perform in this regard, a second-stage of element design may be necessary to reduce the size of the primary beam in order to achieve the sensitivities presented here.

equation 5.4 to evaluate the thermal noise level in each uv-pixel given an effective integration time in that pixel.

Substituting fiducial values into equation 5.4 yields the following result for the sensitivity of a single baseline to the 21cm power spectrum:

$$\begin{aligned}\Delta^2_{\rm N}(k) \approx 8\times10^{-3}&\frac{(1+z)^4}{\sqrt{\Omega_{\rm M}(1+z)^3+\Omega_\Lambda}}\\ &\times\left[\frac{k}{0.1\ h{\rm Mpc}^{-1}}\right]^3\left[\frac{\Omega}{0.09\ {\rm sr}}\right]^{\frac{3}{2}}\\ &\times\left[\frac{180\ {\rm days}}{t_{\rm days}}\right]\left[\frac{|\vec{u}|}{20}\right]\left[\frac{T_{\rm sys}}{50\ {\rm K}}\right]^2{\rm mK}^2,\end{aligned}\tag{5.12}$$

where $|\vec{u}|$ is the length of the baseline in wavelengths. I have also replaced t, the integration time in a sample, with a combination of factors of $|\vec{u}|$ and Ω. These quantities are related to the amount of time a single-baseline will sample a uv-pixel: $\Omega^{\frac{1}{2}}$ sets the diameter of the pixel and $|\vec{u}|$ sets the length of time the baseline samples that same pixel before earth rotation causes it to move into an independent pixel (longer baselines drift through the uv-plane faster). The full derivation of this relation is given in §2.2 of P12a. The choice of a 20-wavelength baseline is arbitrary. I have also added a $t_{\rm days}$ factor; each day of an observation provides an identical measurement to the previous day, resulting in a linear increase in the power spectrum sensitivity. A fiducial "long observation" is 180 days; I set this value as a hard maximum for the number of days BAOBAB can observe in one calendar year. This choice comes from the fact that observations will be compromised by foreground emission when either the Galactic plane or the sun is in view.

In order to calculate the sensitivity of an entire interferometric array, one must sum the contributions of all the Fourier modes sampled by every baseline, paying careful attention to the number of times a unique Fourier mode is sampled by distinct redundant baselines. If each baseline measured an independent Fourier mode, the overall power spectrum sensitivity of an array would grow proportionally to the square root of the number of baselines (or, alternatively, linearly with number of antennas). To first approximation, every baseline can contribute an independent measurement of each $|k|$-mode in the power spectrum, since the frequency axis covers a very broad range in k.[8] However, coherent integration on a particular Fourier mode beats down thermal noise as the square-root of the number of samples in *temperature*, and hence linearly in the temperature-squared units of a power spectrum.

[8] The situation is more complicated than this for two reasons. First, each baseline has a minimum $|k|$-mode it can measure, corresponding to its length in $k_\perp$. The analytic formulae below ignore this effect, which becomes important for small values of k. However, I emphasize that the cosmological sensitivity results in the subsequent sections do take this effect into account. Second, redshift space distortions break the isotropy between $k_\parallel$ and $k_\perp$ so that modes of equal $|k|$ cannot naïvely be combined. As stated previously, I maintain 2D information in the full analysis and only combine modes in annuli of equal $k_\perp$. Only for the purposes of making 1D power spectrum plots, do I combine all modes of equal $|k|$ by reducing the noise in each mode by $(1+\beta\mu^2)^2$.

Therefore, redundant measurements can improve the power spectrum sensitivity of an array to select Fourier modes faster than two non-redundant baselines measuring independent modes of the same magnitude. In the formalism used here, this additional sensitivity boost from redundant sampling enters through the f/f_0 metric for the amount of redundancy in an array configuration, defined in P12a as:

$$\frac{f}{f_0} \equiv \frac{\sum_i n_i^2}{\sum_i n_i}, \tag{5.13}$$

where i labels individual uv pixels, and n_i the number of one-second integration samples falling within pixel i. The ratio f/f_0 measures the increase in sensitivity for a redundant array over one in which there is no sampling redundancy from either redundant baselines or redundant time samples. (This hypothetical "reference" array is obviously non-physical, as it assumes each independent integration in time will measure a statistically independent sky; in practice, this means that the f/f_0 metric incorporates the length of time an array can observe the same patch of sky.) An f/f_0 factor of 10^4 is representative of the 32-element drift-scanning maximum-redundancy arrays described in P12a.

Using this metric, one can express the resultant sensitivity of an arbitrary array as:

$$\begin{aligned}\Delta^2_{\mathrm{N}}(k) \approx 2 \times 10^{-4} & \frac{(1+z)^4}{\sqrt{\Omega_{\mathrm{M}}(1+z)^3+\Omega_\Lambda}} \\ & \times \left[\frac{k}{0.1\ h\mathrm{Mpc}^{-1}}\right]^3 \left[\frac{\Omega}{0.09\ \mathrm{sr}}\right] \left[\frac{T_{\mathrm{sys}}}{50\ \mathrm{K}}\right]^2 \\ & \times \left[\frac{8\ \mathrm{hrs}}{t_{\mathrm{per\_day}}}\right]^{\frac{1}{2}} \left[\frac{180\ \mathrm{days}}{t_{\mathrm{days}}}\right] \left[\frac{32}{N}\right] \left[\frac{10^4 f_0}{f}\right]^{\frac{1}{2}} \mathrm{mK}^2. \end{aligned} \tag{5.14}$$

This equation is derived in Appendix B.2 of P12a. The f/f_0 term is computed from uv sampling patterns including earth rotation aperture synthesis effects over a period of one hour phased to a single pointing center. One hour roughly corresponds to the time it takes a point on the sky to drift through the width of the BAOBAB primary beam, after which a statistically independent patch of sky comes to dominate the data, so that additional integration time only grows the sensitivity as $(t_{\mathrm{per\_day}})^{\frac{1}{2}}$. In general, f/f_0 accounts for most all effects regarding array configuration, so that the additional factors of $\Omega^{\frac{1}{2}}$ and $|\vec{u}|$ that appeared in equation 5.12 do not appear in 5.14. The factor of $t_{\mathrm{per\_day}}$ sets the total integration time per day, which will likely be limited by the size of a low foreground emission region (i.e., a "cold patch"). I choose 8 hours as the maximum time that can spent observing cold patches per day, a value influenced by existing all-sky maps and observations with PAPER.

Sample Variance

In galaxy redshift surveys, the sample variance can be calculated relatively simply by counting Fourier modes over an effective survey volume. However, in the case of BAOBAB,

not all modes are equal, since it will use redundant samples of certain modes to beat down thermal noise. In creating a 1- or 2-D power spectrum out of the full 3-D Fourier space, one must take a weighted combination of these modes, since the ratio of thermal noise to sample variance can vary between every k-mode measured.

Using inverse-variance weighting to combine each measurement at a particular k-mode, one finds that the optimal estimator of the power spectrum results in an error that can be calculated by combining the errors on each measured mode in inverse-quadrature:

$$\delta\Delta^2(k) = \left(\sum_i \frac{1}{\left(\Delta^2_{\mathrm{N,i}}(k) + \Delta^2_{21}(k)\right)^2} \right)^{-\frac{1}{2}}, \tag{5.15}$$

where $\delta\Delta^2(k)$ is the resultant uncertainty on a given k-mode, $\Delta^2_{\mathrm{N},i}$ is the per-mode thermal noise calculated with equation 5.4 taking the full uv-coverage and earth-rotation synthesis into account, $\Delta^2_{21}(k)$ is the cosmological 21cm power spectrum (which is also the sample variance error), and i is an index labeling the independent k-modes measured by the array over which we are summing (note that I never combine modes into bands with significantly varying $|\vec{k}|$, which is why I can sum inverse absolute variance instead of inverse fractional variance).

Since the sample variance is completely a function of the uv-coverage of an array, it is best calculated numerically, as described in the preceding section. In Figure 5.8 I plot the 1D thermal noise and sample variance uncertainties for two maximum-redundancy configurations of BAOBAB — 35- and 132-tiles — shown in Figures 5.6 and 5.10. To calculate the sample variance, I use the 21cm brightness power spectrum from CAMB, also plotted for comparison. At the scale of the first acoustic peak, sample variance has clearly become the dominant source of error for a long integration with 132-elements; I discuss possible methods for improving this situation in §5.3.5 and §5.4.2. At the smaller scales, however, the analytic functions given in the previous section accurately reproduce the array sensitivity. Note that the thermal noise only curve in Figure 5.8 is not a perfect power law because not all baselines can probe the largest scales. I also plot the effect of our foreground model of §5.3.3; the achievable sensitivity if no modes are excluded by foregrounds is shown by the dot-dashed lines. The other two curves do include the effects of the foreground model.

Shot Noise

Measurements of the 21cm power spectrum will also be affected by the discrete nature of the neutral hydrogen distribution at low redshift. Only overdensities self-shielded to the ~ 1 Ry ionizing background contain neutral hydrogen, so BAOBAB will be subject to the same galactic shot noise as optical redshift surveys. Following Seo et al. (2010), I assume that galaxy positions and luminosities are distributed with probability proportional to $1 + b\ \delta_m$, where b is the bias and δ_m is the mass overdensity. This allows for treating shot noise as a scale-independent contribution to the power spectrum with $P(k) = 1/\bar{n}$. In terms of 21cm

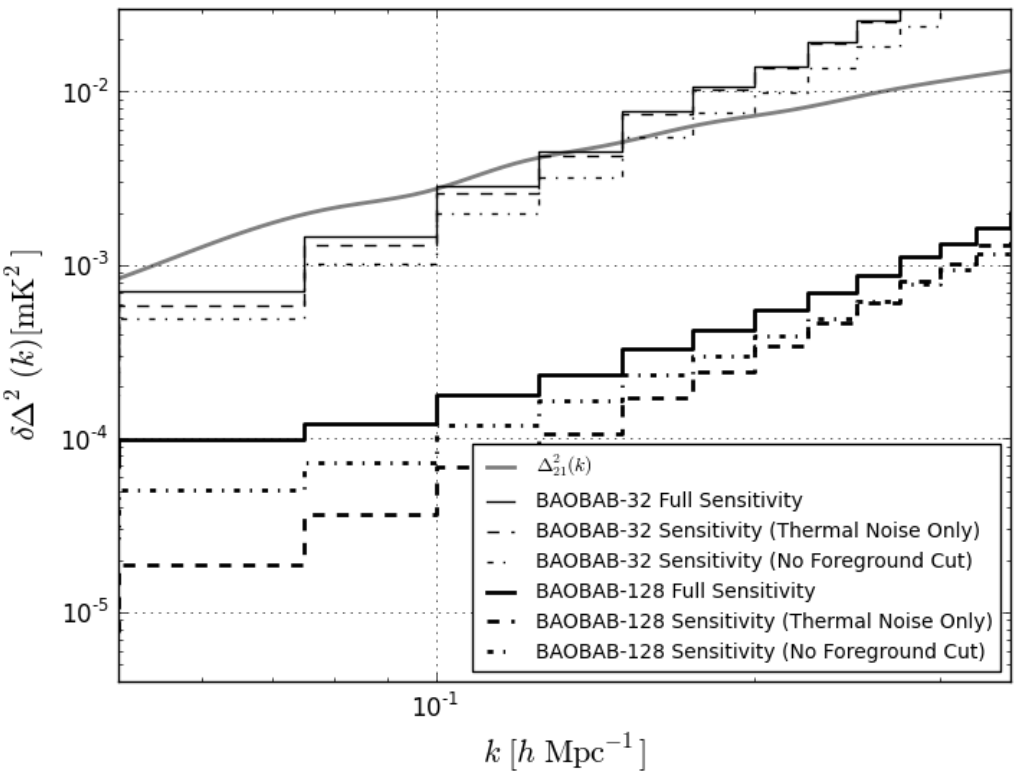

Figure 5.8: Noise levels for two fiducial observations of one declination range with BAOBAB at $z = 0.89$. The dashed thin (thick) line shows the sensitivity given thermal noise only for a 30 days observation with a 35-element system (180 days with a 132-element system); solid lines show the effect of including sample variance. I also plot the effect of our foreground model of §5.3.3; the achievable sensitivity if no modes are excluded by foregrounds is shown by the dot-dashed lines. The other two curves do include the effects of the foreground model. For comparison, I also plot the dimensionless 21cm power spectrum as the thick, gray line. At the larger scales, the 180 day observation with 132-tiles is completely dominated by sample variance. At the smaller scales, the analytic expression for thermal noise given in equation 5.14 accurately reproduces the sensitivity. The plotted thermal noise only curve is not a perfect power-law because longer baselines cannot probe the largest scale k-modes.

brightness:

$$\Delta^2_{\rm shot}(k) \approx \left[\tilde{T}_{21}(z)\right]^2 \frac{1}{\bar{n}} \frac{k^3}{2\pi^2} \tag{5.16}$$

Using the result of Seo et al. (2010) (who use a fit to the neutral hydrogen mass function of Zwaan et al. 2005), $\bar{n} = 0.01\ h^3{\rm Mpc}^{-3}$ and $\Delta^2_{\rm shot} \approx 6.5 \times 10^{-5}{\rm mK}^2$ at $k = 0.1\ h{\rm Mpc}^{-1}$ and $z = 1$. The number density of hydrogen-containing halos is substantially higher than for the bright galaxies used in optical/NIR surveys, making shot noise a substantially smaller contaminant; it will only begin to dominate the signal at $k > 2\ h{\rm Mpc}^{-1}$. Regardless of uncertainties in this calculation, shot noise is clearly a subdominant effect, and I neglect it for the remainder of this chapter.

5.3.3 The Delay Spectrum Technique at $z \sim 1$

Before combining the results of the last two sections, one must consider the effect of foregrounds on power spectrum measurements. The presence of foreground emission orders of magnitude brighter than the cosmological 21cm signal has been one of the major impediments for high-redshift 21cm tomography. P12b presented a per-baseline delay-spectrum technique for isolating foreground emission solely on the basis of its spectral smoothness. In this section I briefly recapitulate the principles of the delay-spectrum technique, and present a simple approximation for the behavior of foregrounds in the 600-900 MHz band.

The delay spectrum technique is a methodology for using each baseline of an interferometer as a independent probe of the 21cm power spectrum. The most powerful aspect of this approach is that the frequency dependence of a baseline's Fourier sampling pattern, typically regarded as a major complication for 21cm experiments, naturally gives rise to an isolation of foreground emission in Fourier space. The ability to remove foregrounds on a per-baseline basis allows multiple baselines to be tuned to target the same Fourier mode for greater sensitivity, as opposed to more traditional techniques that use overlapping uv-coverage at multiple frequencies to avoid the issue of frequency-dependent sampling.

At the heart of the delay transform is a dual interpretation of the Fourier transform of interferometric visibilities along the frequency axis. On the one hand, for 21cm experiments, frequency maps directly into redshift since the observed signal is a spectral line. Therefore, the Fourier transform along the frequency axis gives $k_\parallel$, the Fourier wavemode along the line of sight. However, the frequency dependence of a baseline's length (as measured in wavelengths), gives rise to the delay transform interpretation of the frequency Fourier transform presented in Parsons & Backer (2009). If performed over a wide enough bandwidth, this transform maps sources to Dirac delta functions in "delay space," corresponding to the geometric delay of signal arrival time between the two elements of the baseline. There is thus a maximum delay at which any signal coming from the sky can appear, set by the physical length of the baseline. Furthermore, each source delta-function will be convolved by a kernel representing the Fourier transform of that source's intrinsic spectrum (as well as any spectral features introduced by the instrument). Therefore, as long as the instrumental frequency structure is

kept to a minimum, sources with smooth intrinsic spectra (such as foreground emission) will have their emission confined within the region of delay space set by the maximum delays (the so-called "horizon-limit"). Sources with unsmooth emission, like the 21cm signal, will be convolved by a broad kernel, scattering "sidelobes" well beyond the horizon limit, and creating a window for detecting 21cm emission free of smooth-spectrum foregrounds.

A major component of P12b was to calculate the mapping between cosmological k-space and delay space. To phrase the same question in other terms, we explicitly calculated the effect of "mode-mixing" due to the frequency dependence of a baseline's $k_\perp$ sampling on the recovery of the 21cm power spectrum. For short baselines like those used in BAOBAB, delay-modes proved an effective probe of the 21cm power spectrum, recovering the signal without corruption due to mode-mixing.

A full simulation quantifying the effects of the delay transform on foregrounds is beyond the scope of this present analysis. Rather, I choose to assign a minimum $k_\parallel$ which depends on baseline length, below which we consider modes as being wholly contaminated by foregrounds. These contaminated modes are treated as not "measured" by the array, that is, they are excluded from the sum in equation 5.15. Since the horizon limit described above is a linear function of baseline length, I use a $k_{\parallel,\mathrm{min}}$ which linearly increases on longer baselines. I model the choice for the exact value of $k_{\parallel,\mathrm{min}}$ on the simulations presented for PAPER in P12b, which finds for 16λ-baselines at 150 MHz foregrounds contaminate modes with $k_\parallel \lesssim 0.2\ h\mathrm{Mpc}^{-1}$. At EoR frequencies of 150 MHz, this cutoff maps to delay-modes of 400 ns. Since BAOBAB baselines are physically shorter by a factor of 5, this reduces the maximum delay-space contamination to 80 ns, which in turn maps back to $k = 0.1\ h\mathrm{Mpc}^{-1}$ at 750 MHz, using the Y parameter from equation 5.8. There are two important factors which will further serve to reduce this number for BAOBAB. First, celestial foregrounds should have power law spectra with steeply decrease in intensity versus frequency, and so will be fainter than at EoR frequencies. Although the signal has also fallen a similar amount, this reduced foreground structure will still make the delay transform even more effective at isolating foreground emission. Secondly, the narrower primary beam of BAOBAB will limit the delay modes from which there can be appreciable celestial emission, as sources near the horizon will be significantly attenuated.

To better determine the scale of $k_{\parallel,\mathrm{min}}$ in the foreground model, I perform a cursory calculation in which the delay-transform is applied to a simulated sky model. In these simulations, I assume the sky is entirely composed of point sources, where the source strength distribution follows a power-law with a slope of -2.0, normalized to a 2 Jy source per 10 steradians, a distribution derived empirically from PAPER data with extrapolation to the BAOBAB band. I also model the frequency spectrum of each source as a power-law with a normal distribution of spectral indices centered on -1.0 and a standard deviation of 0.25. I refer to these simulations as "cursory," since they exclude instrumental effects such as RFI flagging and frequency-dependent beam structure. Instead, I use a single, frequency-independent Gaussian to model the primary beam of the BAOBAB tile; the potential effects of a more realistic beam model are discussed in §5.4. I find the delay transform confines foregrounds to k-modes below a value of $k_\parallel = 0.045\ h\mathrm{Mpc}^{-1}$ for baselines of 16 wavelengths. The $k_\parallel$ value

for the maximum delay of a 16λ baseline at 750 MHz (i.e., the horizon limit) is 0.028 hMpc^{-1}, which implies that the intrinsic spectral behavior of foreground emission corresponds to a kernel of width ~ 0.02 hMpc^{-1}. In this analysis, the foreground model is to exclude k-modes smaller than the sum of the maximum realizable delay on a baseline (converted from seconds of light-travel time to hMpc^{-1} using equation 5.8) and this kernel. The maximum realizable delay scales linearly with baseline length, while the additive kernel remains constant. In effect, this model states that intrinsic spectral structure in foregrounds corrupt 0.02 hMpc^{-1} beyond a naïve prediction based only on the physical length of the baseline. I explore the effects of modifying this model in §5.4.1.

5.3.4 Detecting the HI Power Spectrum

The first major science result from BAOBAB will be the detection of the 21cm power spectrum near $z \sim 1$. I present predictions for the power spectrum error bars using the formalism outlined above: I fully simulate the uv-coverage of our arrays, including the effects of earth-rotation synthesis, over a one-hour period. I then use equation 5.4 to evaluate the thermal noise level in each uv-pixel given an effective integration time in that pixel. I incorporate redshift-space distortion effects by reducing the magnitude of these thermal noise errors according to equation 5.3. I include the effects of sample variance in our measurements using equation 5.15, combining measurements with the same $|k_\perp|$, but maintaining 2D information in the $(k_\perp, k_\parallel)$-plane. These error bars are further reduced by both the square root of the number of independent 1-hour pointings available per-day (8 in our fiducial calculation) and by the number of days observed. For the plots below, I further compress these errors into 1D, but use the full 2D information for our calculations of detection significance and cosmological parameter extraction in §5.3.5.

A high significance detection will be achievable with a short $\sim$ 1-month observation with a 35-tile system operating in the maximum redundancy shown in Figure 5.6.[9] The predicted measurement for a 30-day observation (240 hours) is shown in Figure 5.9. These observations assume a 100-MHz bandwidth centered on 750 MHz ($z = 0.89$). The net result is an 5.6σ detection of the 21cm power spectrum when the model for foreground emission from §5.3.3 is used to exclude contaminated modes. Results for bands centered on 650 ($z = 1.18$) and 850 MHz ($z = 0.67$) are similar, yielding 5.8σ and 5.0σ detections, respectively. Although a small effect, I do modify the system temperature in each band to represent the change in sky temperature; a spectrally-flat $T_{\rm sys}$ is therefore only assumed on 100 MHz scales. Rather, the lower significance detection at the lowest redshifts results primarily from the scaling of the angular diameter distance; at redshift $z = 0.67$, a 16λ baseline corresponds to a $k_\perp$ wavemode of ~ 0.06 hMpc^{-1}, limiting the number of baselines that can probe the largest-scale k-modes where thermal noise is lowest. Over the $z = 0.5 - 1.5$ range, both the 21cm signal and, the noise remain roughly constant in magnitude, the latter because it is dominated by a frequency-independent front-end amplifier noise temperature. This trend does not continue

[9] Correlator inputs have traditionally been in powers of 2; hence this array is our $\sim$ 32-tile configuration. The sensitivity calculations do assume that all 35 elements are correlated.

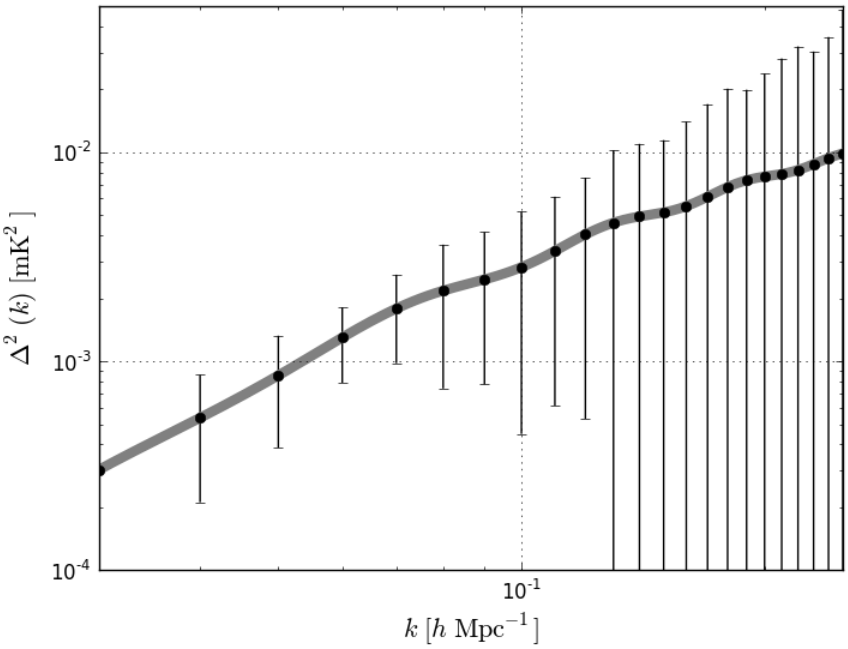

Figure 5.9: Predicted constraints on $z = 0.89$ 21cm power spectrum from a 30 day observation with a 35-tile BAOBAB system. The net result is a $\sim 5\sigma$ detection of the power spectrum. Results are comparable in the other 2 redshift bins.

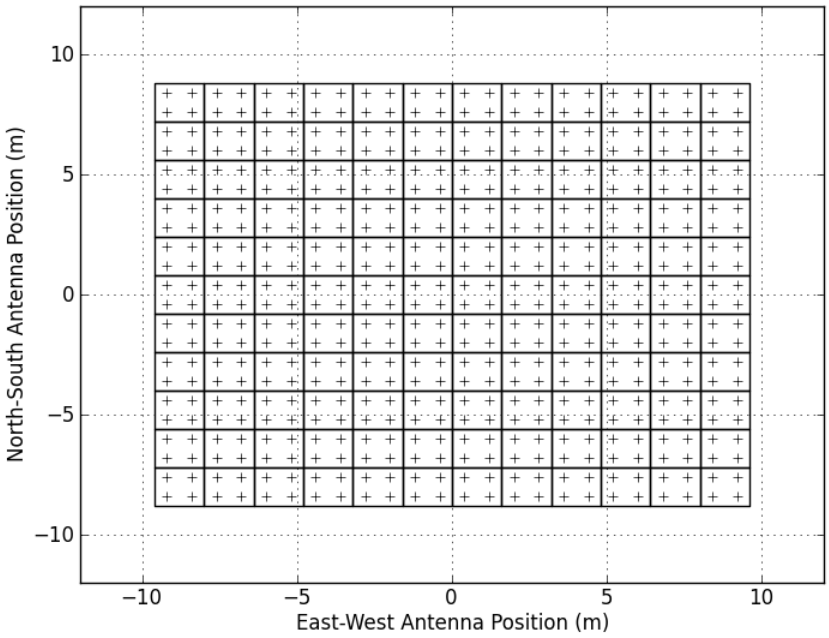

Figure 5.10: The maximum redundancy configuration of a 132-element BAOBAB system. The close-packed tiles are chosen to produce the shortest possible baselines.

indefinitely, however, as sky noise increases with increasing redshift, eventually dominating the system temperature.

Measurements of this significance will allow for an accurate determination of $f_{\rm HI}\ b$, the combination of the cosmic neutral hydrogen fraction and the bias of neutral hydrogen containing regions, as a function of redshift. Breaking the degeneracy between these parameters will require additional information. Measuring redshift-space distortions can, in principle, separate the effects of the two terms. Constraints from a longer integration or a system with ~ 64 elements will further improve constraints on the neutral hydrogen power spectrum, and it will be possible for these systems to measure redshift-space distortion effects. Measuring these effects accurately requires more careful systematic control, which may warrant different configurations and observing strategies, so we postpone an exploration of this science to a future work (see, e.g., Masui et al. 2010).

5.3.5 Detecting Baryon Acoustic Oscillations

As shown in in Figure 5.8, a 132-tile BAOBAB array with the configuration shown in Figure 5.10 has effectively reached the sample variance limit in 180 days (1440 hours) of observing time. Using the methodology described at the beginning of §5.3.4, I calculate that this observation yields a 3.3σ detection of the BAO features at $z = 1.18$, with effective

2.1− and 2.7σ non-detections at the $z = 0.67$ and $z = 0.89$ bands, respectively, where I have isolated the BAO features from the broad-band shape of the power spectrum by removing a model fit using the transfer function from Eisenstein & Hu (1998). The effect of sample variance is most dominant at the lower redshifts, because the angular diameter distance scaling means that fewer samples of the BAO scale can be found in the same area of sky (this observation of 8-independent fields with a 0.045 steradian primary beam corresponds to an effective survey area of $\sim$ 1200 square degrees). While longer observations with the same array configuration can improve these constraints by reducing thermal noise on the smaller scale modes, a better approach will be to observe additional independent fields.

For a fiducial BAOBAB observation, I use an array which observes 24 independent fields (i.e., three independent declination observations for 8 hours per day), yielding an effective survey area of $\sim$ 3600 square degrees. I discuss the motivation for this particular approach to increasing survey volume in §5.4.2. There are several equivalent ways an experiment can probe this additional area. One approach would be to conduct three 1-year surveys, with the dipoles pointed towards a different declination; this could be achieved by physically placing the array on an platform inclined by $\sim 15^\circ$, or potentially by adding a steering component to the tile-beamformers. If the beamformers are designed to allow multiple beams, one could in principle achieve similar sensitivities with only one year of observing, although at the expense of additional degrees of complexity in the system. If funding permits, the simplest approach might be to build three 132-tile BAOBAB arrays, each tilted towards a different declination; this would also yield the subsequently predicted sensitivities in only one season of observing. Note that for BAO science these independent configurations are potentially more desirable than an array with a larger number of fully correlated tiles. Since we are using a close-packed configuration, the addition of more tiles can only yield new modes at corresponding larger $k_\perp$ where the amount of BAO information is significantly diminished. It may be possible that more information could be recovered from these larger k modes using a reconstruction method (e.g. Eisenstein et al. 2007, Padmanabhan et al. 2009, Noh et al. 2009), but I do not explore this option in this chapter. If there is significant BAO information beyond $k \sim 0.2$, then a $\sim$ 256-tile array could possibly yield tighter constraints than two 132-tile arrays observing for the same amount of time. In §5.4.2, I discuss other ways to increase the survey footprint and further reduce the effect of sample variance.

I plot the expected constraints for the three-declination range fiducial BAOBAB observations on the $z = 0.89$ BAO features in Figure 5.11, where the broad band shape of the power spectrum has been removed using the transfer function fit of Eisenstein & Hu (1998), which neglects BAO. The measurements from a 180-day integration at each declination range with this array amount to an 4.7σ detection of these features. Results are similar for the other redshift bins, with expected 3.6σ and 5.7σ detections at redshifts of 0.67 and 1.18, respectively. While it is clear from Figure 5.8 that sample variance dominates the errors on the largest scale modes after 180 days of observing one declination range, I find the additional sensitivity towards the higher BAO peaks with this observing duration yields better constraints on the signal than e.g., observing twice as many declination ranges for 90 days.

With a significant BAO detection, BAOBAB can also begin to place constraints on

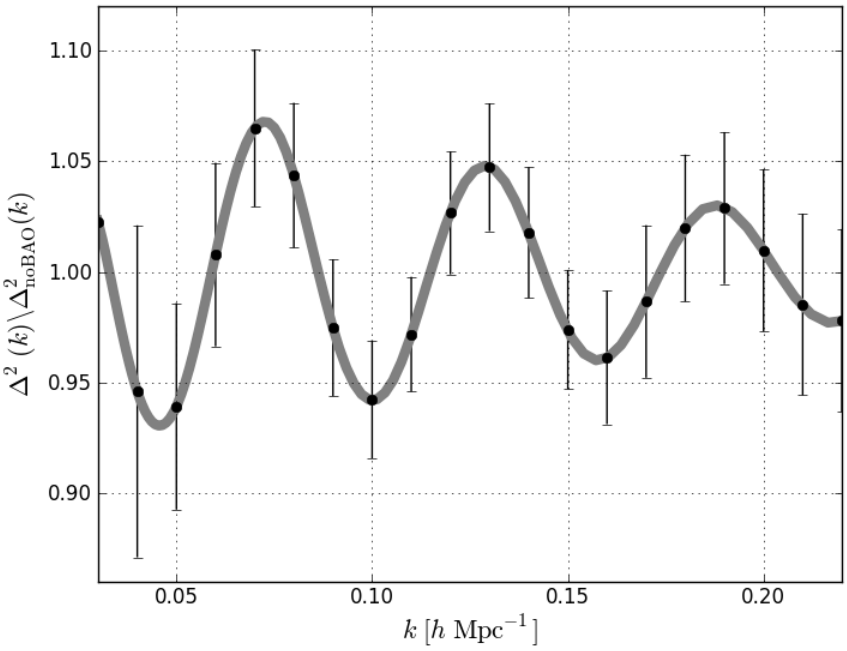

Figure 5.11: Predicted constraints on the $z = 0.89$ BAO features from a 180 day observation of three-declination fields with a 132-element BAOBAB system. The net result is a 4.7σ detection of the BAO features. The sensitivity is comparable in the other 2 redshift bins.

cosmological parameters. To quantify the effect of such measurements, I use the Fisher matrix formalism of the Joint Dark Energy Mission (JDEM) Figure of Merit Science Working Group (FoMSWG; Albrecht et al. 2009), defining our Fisher matrix as:

$$\mathcal{F}_{ij} = \sum_b \frac{1}{\sigma_b^2} \frac{\partial f_b}{\partial p^i} \frac{\partial f_b}{\partial p^j}, \tag{5.17}$$

where f is some observable measured at some b values, σ^2 is the variance in a measured value of f, p^i are cosmological parameters, and we sum over all measured f_b values.

I propagate our power spectrum measurements into constraints on the Hubble parameter $H(z)$ and angular diameter distance D_A using this formalism, where $f = \Delta^2(k)$, the measured power spectrum, $b = k$, the set of k-modes measured, and $p^i = [H(z), D_A]$. The derivatives with respect to D_A and $H(z)$ are straightforward to calculate, as they affect the measurements through the X and Y parameters defined in equations 5.5 and 5.8. In effect, changing $H(z)$ or D_A changes the k-modes sampled by BAOBAB. For these calculations, I exclude k-modes deemed contaminated by the foreground model of §5.3.3. I also model the nonlinear degeneration of higher k-modes using the elliptical Gaussian formula from Seo & Eisenstein (2007). To isolate the constraints provided by the BAO features from the broad-band shape of the power spectrum, I again remove a model fit using the transfer function from Eisenstein & Hu (1998). I split the data into three redshift bins centered at $z = 0.67, 0.89$ and 1.18. Although BAOBAB's frequency coverage will be continuous between $z = 0.58$ and 1.37, I find that there is minimal penalty for using only three bins in a Fisher matrix study. The result of this calculation is that the fiducial 3-declination, 180-day integration yields measurements of $H(z)$ with an error ranging from 9% to 4.5% across the three redshift bins, from low to high redshift, and measurements of D_A with errors effectively constant at 17% over the same range (note that errors $\gtrsim 5\%$ should be understood in the usual formal Fisher matrix sense – these measurements would not correspond to significant detections, taken in isolation). The particularly poor constraints on D_A come from the loss of modes due to foreground emission; I further explore the effects of our foreground model in §5.4.1. Taking correlations between $H(z)$ and D_A into account, these measurements amount to 2.5% to 4.5% errors on a "dilation factor," which scales D_A and $H^{-1}(z)$ in proportion. Expressed as a single constraint on the $z \sim 1$ distance scale, these measurements correspond to a dilation factor error of 1.8%. The exact uncertainties are given in Table 5.2.

It is also straightforward to propagate these errors on $H(z)$ and D_A into errors on underlying cosmological parameters through the Fisher matrix formalism of equation 5.17. In this case, f is now $H(z)$ or D_A, and b corresponds to the redshift bin it was measured in. $H(z)$ is given by:

$$H^2(z) = H_0^2 \Big[\Omega_m (1+z)^3 + \Omega_k (1+z)^2 + \Omega_{DE} \exp\left(3 \int_0^z \frac{\mathrm{dz'}}{1+\mathrm{z'}} [1 + \mathrm{w(z')}] \right) \Big]. \tag{5.18}$$

Table 5.2: Percent errors on the distance scale from BAO measurements, for a three declination BAOBAB survey. The correlation is the correlation coefficient between the $H(z)$ and D_A measurements. R is the "dilation factor," a single estimate of the distance scale which scales D_A and $H^{-1}(z)$ in proportion. BOSS-LRG constraints come from Schlegel et al. (2009) and BOSS-Lyα Forest constraints come from the method of McDonald & Eisenstein (2007) and Dawson et al. (2013).

Survey	Redshift	$H(z)$ Error	D_A Error	Correlation	R Error
BAOBAB	0.67	8.9%	17.1%	0.71	4.4%
BAOBAB	0.89	6.1%	16.4%	0.72	3.3%
BAOBAB	1.18	4.5%	17.5%	0.73	2.6%
BOSS-LRG	0.35	1.8%	1.0%	0.41	0.7%
BOSS-LRG	0.6	1.7%	1.0%	0.41	0.7%
BOSS-Lyα	2.5	3.1%	7.4%	0.58	2.0%

D_A is given by:

$$D_A(z) = \frac{1}{1+z}\int_0^z \frac{cdz}{H(z)} \tag{5.19}$$

The parameters of interest, p^i, are now the underlying cosmological parameters. I use the parameterization of the JDEM FoMSWG, which include w_m, w_b, w_k and w_{DE} ($w_X = \Omega_X h^2$, where m, b, k, and DE correspond to the matter, baryon, curvature and present day dark energy density, respectively) although I use the simpler 2-component form for the dark energy equation of state:

$$w(a) = w_0 + (1-a)w_a. \tag{5.20}$$

Following the convention of Dark Energy Task Force report (Albrecht et al. 2006), I marginalize over all other parameters after combining our Fisher matrices with constraints from other experiments, to create 2×2 matrix representing constraints on w_0 and w_a. As a Figure of Merit I use FoM $= |\mathcal{F}'|^{1/2}$, where $\mathcal{F}'$ is the original Fisher matrix $\mathcal{F}$, marginalized to a 2D (w_0, w_a)-space; this FoM is proportional to the inverse of the error ellipse area in the $w_0 - w_a$ plane.

As a baseline for current dark energy constraints, I use the JDEM FoMSWG predictions for the Planck satellite, combined with constraints from the BOSS-LRG survey listed in Table 5.2 and a 5% error on H_0. This combination of experiments yields an FoM of 8.7. Including the measurements from the fiducial BAOBAB observation increases this FoM value to 16.6. For comparison, BAOBAB combined with Planck and the H_0 constraint only yields a FoM of 4.4; the strength of BAOBAB therefore lies in adding complementary high-redshift information to the BOSS-LRG survey: the high-redshift constraints from BAOBAB can significantly improve our current measurements of the dark energy equation of state. (And, since the BOSS experiment is already underway and yielding high-quality data (Dawson et al. 2013), this is the more interesting comparison to make.) These constraints are plotted as 1 and 2σ error ellipses in Figure 5.12. If we include the Lyα forest survey of BOSS, our baseline constraint FoM becomes 17.8 which is improved to 23.0 with the inclusion of BAOBAB data.

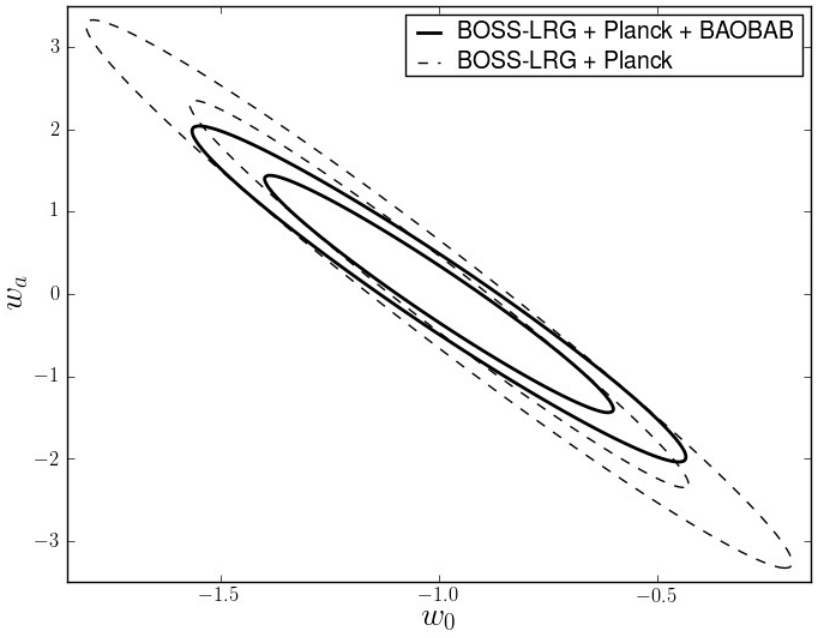

Figure 5.12: 1 and 2σ error ellipses in the $w_0 - w_a$ plane for various surveys. The dotted line shows constraints from Planck, BOSS-LRG and a 5% error on H_0; the solid line shows the effect of including a 1440-hour integration on three independent declination fields with a 132-tile BAOBAB array. These error ellipses correspond to Figures of Merit of 8.7 and 16.6, respectively.

Even when the BOSS-Lyα forest constraints between redshifts 2 and 3 are added, BAOBAB still provides valuable information, serving to increase the FoM by $\sim 25\%$.

5.4 Discussion

I break this discussion into two parts. In §5.4.1, I consider two components of the analysis which may be overly simplistic: the model of the BAOBAB primary beam as a frequency-independent Gaussian, and the foreground emission model. I discuss the effects any shortcomings in these models could have on our conclusions. In §5.4.2, I consider the fact that sample variance is the dominant source of uncertainty in our measurements, and present future approaches that could improve the dark energy constraints possible with the BAOBAB instrument.

5.4.1 Potential Shortcomings in the Analysis

A effect that could modify the predictions presented in this chapter is the use of an overly simplistic model for the BAOBAB beam: a frequency independent Gaussian. Although this model is sufficient to calculate sensitivities, the effects of a more realistic beam model on the delay spectrum foreground removal technique will necessitate further investigations. The principal cause for concern is from the existence of frequency-dependent grating-lobe structure associated with the tile. These sidelobes can introduce apparent frequency structure into otherwise spectrally smooth foreground emission. If this structure represents a significant increase in the size of the delay-space convolving kernel, it will move the foreground contaminated region to higher k values. If the size of the effect is large enough to push foreground emission on the shortest baselines beyond the first BAO peak, the predicted cosmological constraints could be reduced.

Note, however, that the frequency-independent Gaussian beam is not as bad an assumption as it might first appear. The choice to neglect the frequency evolution of the beam is partially motivated by experience with the PAPER dipole beam, which, like BAOBAB, uses a modified dual-polarization "sleeved" dipole design to limit the frequency evolution of the beam to only $\sim 10\%$ over the $120 - 180$ MHz band (Chapter 3). Of more concern are the grating lobes introduced by beamforming in tiles. If one were trying to image the sky, ignoring the grating lobes would be unjustified. With the delay spectrum approach, though, the issue is not the existence of the sidelobes, but their frequency dependence. If the frequency Fourier transform of the beam pattern is particularly broad — corresponding to rapid evolution of the beam pattern with frequency — then foreground emission will have a similarly broad footprint in delay space, compromising the 21cm signal. Of course, the grating lobes themselves will change position as a function of frequency, introducing additional structure not in the PAPER beam. However, as argued in P12b, it is difficult for an element only several wavelengths across to possess such frequency structure. As stated in §5.3.3, I will postpone a detailed investigation of these effects to a future paper with empirical studies of the beam shape, and

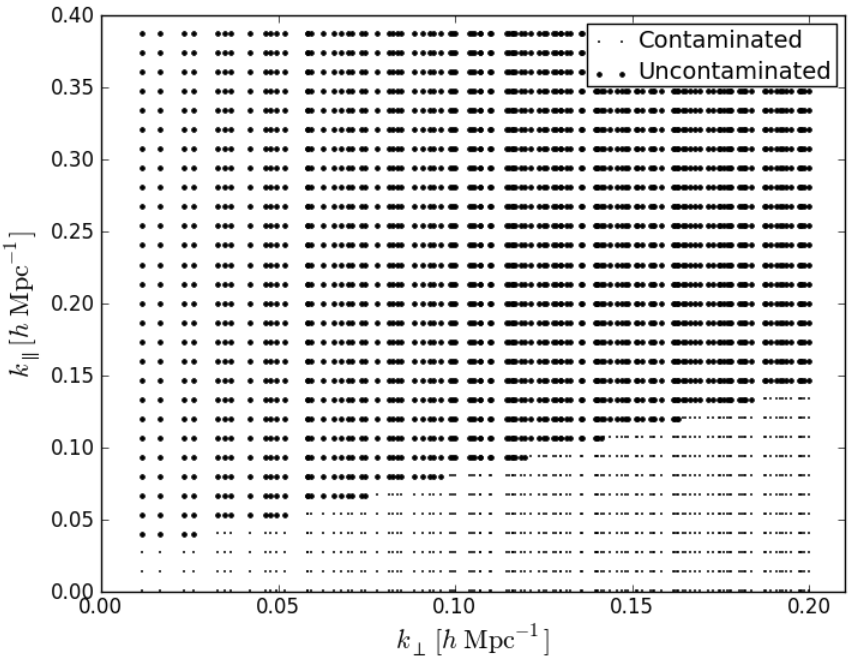

Figure 5.13: The sampling pattern of BAOBAB-132 at $z = 0.89$ in the ($k_\perp$-$k_\parallel$)-plane. Modes are marked as either uncontaminated or contaminated by foregrounds using our model of §5.3.3. Foregrounds limit μ to, e.g., $\gtrsim 0.65$ for $k \sim 0.1$, leading to the relatively poor D_A measurement and high correlation.

focus here on array sensitivities, for which the Gaussian model is sufficient.

As noted in §5.3.5, the constraints on D_A are significantly limited by foreground emission. I illustrate this effect in Figure 5.13, which shows the sampling pattern of BAOBAB-132 at 750 MHz in the ($k_\perp$-$k_\parallel$)-plane, highlighting those modes discarded as contaminated by our foreground model of §5.3.3. Foreground contamination effectively excludes modes where $k_\perp \lesssim k_\parallel$, i.e., transverse modes. This has the effect of significantly degrading the achievable constraints on D_A.

While the foreground model presented in §5.3.3 is empirically motivated, accurate predictions for foreground emission will need to come from early BAOBAB arrays or other 21cm experiments. I will, however, briefly explore the effect of changing the size of the additive component of our foreground model arising from the spectral smoothness (or lack thereof) the emission. For the fiducial model, this term has a magnitude of $k_\parallel = 0.02\ h\mathrm{Mpc}^{-1}$; as test cases, I analyze the constraints obtainable with BAOBAB if this term is changed by $\Delta k_\parallel = \pm\ 0.01\ h\mathrm{Mpc}^{-1}$. This has the effect of moving the cutoff between contaminated and uncontaminated modes in Figure 5.13 up or down by $0.01\ h\mathrm{Mpc}^{-1}$.

The effect of increasing this term (i.e., if foreground emission is not as spectrally smooth as predicted) is to degrade the significance of the BAO detections in each of the redshift bins

by $\sim 0.4\sigma$, and the $H(z)$, D_A, and R constraints by ~ 0.5, ~ 5, and ~ 0.3 percentage points, respectively. Reducing the foreground emission footprint in k-space (i.e., if foregrounds are spectrally smoother than predicted) has similar effects with the opposite sign: the significance of the BAO detections in each band are increased by $\sim 0.3\sigma$, and the $H(z)$, D_A, and R constraints are improved by ~ 0.5, ~ 5, and ~ 0.3 percentage points, respectively. From this analysis, it is clear that foreground emission can significantly alter the achievable constraints on D_A, but ultimately, the success of BAOBAB will not be determined by the details of foreground spectral properties.

5.4.2 Improving The Constraints

As discussed in §5.3.5, sample variance dominates the uncertainties in power spectrum measurements from BAOBAB. However, measuring a new set of independent modes is not a trivial undertaking. By using a close-packed array, BAOBAB completely samples the uv-plane out to some maximum baseline length. Therefore a simple array rotation or reconfiguration will not yield new samples. Furthermore, in the highest redshift bin of $z = 1.18$, the longest baseline in the 132-tile array probes a transverse k-mode of $k_\perp = 0.15\ h\mathrm{Mpc}^{-1}$ (and at lower redshift, this longest baseline corresponds to an even larger value of $k_\perp$). At this scale and smaller, most BAO information is being lost to non-linear damping effects. Therefore, while a larger array will beat down thermal noise faster, the cosmological returns from increasing the array size beyond ~ 128 tiles are limited, since effectively no new modes with significant BAO information will be probed.

As an upper-limit to the constraints obtainable with a single-declination, ~ 1200 square degree, 132-tile BAOBAB observation (as opposed to the fiducial observation targeting three declination fields) I compute the results of a completely sample variance limited survey, i.e., one where thermal noise uncertainties have been set to 0 (although modes are still excluded using our foreground emission model). This sample variance limited 132-tile BAOBAB observation yields distance scale uncertainty of $\sim 1.5\%$, averaged over the whole band (compared with 1.8% for the fiducial 1-year observation of 3 independent declination fields). For comparison, a 10-year (14,400 hour), one declination observation with the same array yields a distance scale uncertainty of 1.9%. Therefore, while better constraints can come from a longer observation, obtaining measurements of new modes to beat down sample variance is clearly the optimal way to proceed. There are two ways forward to achieve this goal: map a different volume of the universe (as with our fiducial experiment) or recover foreground corrupted modes. We consider each of these approaches in turn.

Since BAOBAB is a zenith-pointing, drift-scanning telescope, to map a new area of the sky, it will either need to be relocated to a different latitude or tilted to point towards a different patch of sky. Either option is potentially feasible, as even a 132-tile BAOBAB experiment spans less than 20 meters. With a primary beam full-width half-max of $\sim 15°$, there are ~ 10 independent pointings in declination that BAOBAB can target. Since the fiducial observation already targets three declination fields, mapping every declination could in principle yield up to $\sim 70\%$ reductions in the error bars over the results presented.

The other way to potentially measure new modes and beat down sample variance is to recover samples I have considered corrupted by foregrounds. There are two ways foregrounds compromise BAOBAB observations. The first is the limited observing time per-day, set by Galactic emission, which I have treated as irreparably corrupting all samples, even those in principle recoverable with the delay transform. If it is possible to observe all 24 hours of right ascension, as opposed to the 8 considered here, the constraints from a single observing season will increase by a factor of $\sqrt{3}$. While it is unlikely that all 24 hours of right ascension will be workable, the fiducial value of 8 hours per day, motivated by observations with present EoR experiments, may well be conservative, since Galactic synchrotron emission has significantly fallen in brightness compared to EoR frequencies.

Even when observing a "cold patch," foregrounds corrupt large scale k-modes with a footprint moving to smaller scales as baseline length increases (§5.3.3). If these modes could be retrieved, they could significantly increase the volume of Fourier space that BAOBAB can probe. As an upper-limit, I calculate the obtainable power spectrum constraints ignoring all foreground contamination. The result is that a 1-year (1440-hour) observation in each of three independent declination fields yields a distance scale uncertainty of 1.4% combined over the entire redshift range, an increase of $\sim 25\%$ over the same observation including foreground emission. In particular, note that a foreground-free observation yields errors of $\sim 5\%$ on D_A at redshifts of 0.67, 0.89 and 1.18, respectively — a factor of $\gtrsim 3$ improvement over the predictions for an observation including the effects of foreground emission. While an analysis of foreground removal techniques is beyond the scope of this work, this result suggests that foreground removal may be the way to improve constraints on D_A.

As an order of magnitude estimate, one can consider whether a foreground removal or subtraction scheme might be more effective in the BAO band than at EoR frequencies. At $k \sim 0.1\ h\mathrm{Mpc}^{-1}$, the 750-MHz BAO 21cm power spectrum reaches $\sim\ 3 \times 10^{-3}\ \mathrm{mK}^2$, compared with a peak brightness at EoR frequencies of 150 MHz reaching $\sim\ 10\ \mathrm{mK}^2$. The steep spectrum Galactic synchrotron emission has a spectral index of -2.5, and so will fall by a factor of $(5^{-2.5})^2 = 3.2 \times 10^{-4}$ in units of temperature squared. (Extragalactic point sources are less steep spectrum, and so will not fall off in brightness as steeply. Therefore, this estimate can be considered a lower limit on the foreground-to-signal ratio). Roughly speaking, then, the foreground-to-signal ratio is unchanged compared with EoR experiments, suggesting that a foreground isolation scheme like the delay-spectrum technique is still likely the most viable approach for first-generation experiments limited in collecting area.

5.5 Conclusions

In this chapter, I have presented a concept for a new experiment using the redshifted 21cm line of neutral hydrogen to probe cosmology at $z \sim 1$. The BAO Broadband and Broad-beam Array (BAOBAB) will incorporate both the hardware and analysis infrastructure developed for 21cm experiments at higher redshifts.

The hardware design will borrow heavily from the Precision Array for Probing the Epoch

of Reionization (PAPER) and the Murchsion Widefield Array (MWA), using a scaled version of the PAPER dipoles as a feed element, tiling dipoles as done by the MWA, and modifying the CASPER FPGA/GPU PAPER correlator to perform full dual-polarization cross-correlations of all elements. Significant improvements to the system temperature will be brought about through state-of-the-art uncooled, low-noise amplifiers. Relative to PAPER, the collecting area will be substantially increased through the use of tiles of 4 dipoles combined through a beamformer, as demonstrated by MWA efforts. Although I have largely avoided specific cost-estimates, it is fair to say that this infrastructure is obtainable at a fraction of the cost of the $\sim$ 100-million dollar ground-based optical redshift surveys.

On the analysis side, BAOBAB will use the maximum redundancy configurations and delay spectrum foreground removal techniques presented in Parsons et al. (2012a) and Parsons et al. (2012b) to enhance sensitivity to Fourier modes along the line-of-sight. Motivated by the science of Baryon Acoustic Oscillations, BAOBAB will utilize extremely close-packed arrays to maximize the number of short baselines. The sensitivity calculations presented here show that BAOBAB will achieve several milestone measurements over our anticipated staged deployment process. A $\sim$ 32-element BAOBAB system will yield high significance detections of the HI power spectrum, and constrain the evolution of the cosmic neutral hydrogen fraction from $z = 0.5$ to 1.5 as well as the bias of DLAs. Over the same wide redshift range, a $\sim$ 128-element system will allow for a first detection of the BAO features in the power spectrum, and yield errors on the distance scale R at the several percent level. When combined with our current constraints on dark energy, including those forthcoming from the BOSS and Planck experiments, BAOBAB's measurements result in substantial increase in the Dark Energy Task Force Figure of Merit, representing constraints on the nature and time evolution of dark energy over a wide range of cosmic history.

Acknowledgments

A version of this chapter was previously published as Pober et al. 2013 AJ 145 65, and is reproduced with the permission of all coauthors and the copyright holder.

Chapter 6

Conclusions

6.1 Summary

Before summarizing the main results of this dissertation, it is worthwhile to review the status of PAPER when I started work on the project. Parsons et al. (2010) gives a clear description of the state of the PAPER project at this time. The analog and digital systems were clearly defined, although the number of elements and correlator inputs were both slated to grow in the next few years. A cursory all-sky map existed, although no real cataloging efforts had been made to date. The clear next steps involved calibrating the system to a significantly higher level than had been previously achieved.

The work in the first two chapters of this dissertation solved two of the most significant calibration challenges facing the PAPER array at this time. The temperature dependence of the system gain had been measured in the lab, but in order to move forward with robust cataloging efforts and long integrations for building power spectrum sensitivity, the correction scheme needed to be verified as improving the data itself. The main result of Chapter 2 is a conclusive verification the we understand the temperature dependence of the PAPER system gain and that it can be corrected for with a simple model and a measurement of the ambient temperature.

Far more daunting a calibration challenge was the issue of the PAPER primary beam. As outlined in Chapter 3, wide-field, drift-scanning arrays like PAPER posed new calibration challenges that could not be solved with traditional approaches. The most problematic issue was that existing source catalogs in the PAPER band did not contain accurate measurements of enough sources to truly calibrate the beam over the sky. And, since a good beam model was necessary for any cataloging effort with PAPER, the problem appeared to be ill-conditioned. One needed a good beam model to make accurate catalogs, but one needed accurate catalogs to derive a good beam model. The work presented in Chapter 3 solved this problem for PAPER and similar arrays by designing (and robustly testing) a new beam calibration method that could solve for both source flux densities and a beam model simultaneously. Although the model of the PAPER primary beam presented in this Chapter is accurate to only $\sim 10\%$, work is currently being undertaken using this method and higher quality data from a 64

element PAPER array in South Africa to further improve the model.

Simultaneous to the calibration efforts described in Chapters 2 and 3, the PAPER group made significant strides in understanding the 21cm power spectrum sensitivity of the array, and developing the delay spectrum formalism that gives rise to the predictions for the EoR window. These results were formalized in Parsons et al. (2012a,b), and were shortly thereafter implemented as observing strategies for the PAPER array in South Africa. The main result of Chapter 4 of this dissertation was to produce the first observations of the "wedge" and the EoR window in actual data from any array. The predictions of Parsons et al. (2012b) were spectacularly met, as the EoR window appeared with little to no foreground contamination, at least at the sensitivity level of the analysis. Higher sensitivity observations will be necessary to measure how fast the wedge rolls-off into the EoR window. Parsons et al. (2012b) predicts at the level of ~ 10 mK^2 on 16λ baselines, foregrounds corrupt k-modes below $k_\parallel = 0.02\ h\mathrm{Mpc}^{-1}$ — completely consistent with the measurements presented here.

The final result of this dissertation was to collect the lessons learned from 21cm experiments — both PAPER and the MWA — regarding both analog and digital design, as well as in analysis techniques, and apply them to the design of a new experiment for $z \sim 1$ 21cm intensity mapping. The forecasts for BAOBAB presented in Chapter 5 are extremely encouraging, with high significance detections of the power spectrum and, eventually, significant dark energy constraints achievable with relatively modest sized arrays.

6.2 Future Directions

At nearly the exact same time as this dissertation, a new upper limit on the 21cm power spectrum from PAPER was released (Parsons et al. 2013). This new measurement was highly encouraging for the delay spectrum technique, improving on the previous limit by two of magnitude (in temperature-squared). It also gave tantalizing hints of the newest problems to be solved for continued improvements in the measurements. The lowest bins in this power spectrum are not consistent with zero — indicating potential foreground contamination of the EoR window, or a new systematic, e.g., polarization leakage of Faraday rotated emission (Moore et al. 2013). The PAPER approach has always been to solve problems as they arise, and these new issues are no different.

Upcoming investigations will look at the power spectra of all the Stokes parameters to test the polarization leakage theory. A new beam-weighting technique has the promise to steer the beam away from foreground contaminated regions, to test the extent to which foreground leakage compromises the data. Finally, new, higher sensitivity measurements from a 64 element PAPER observing season will test our ability to keep integrating down.

Lastly, it should be stressed that the era of 21cm cosmology yielding results of cosmological interest is not far away. Parsons et al. (2013) can, for a wide range of models, conclude that heating of the IGM took place by $z = 7.7$. To some level, this conclusion surprises no one, but it represents a transition to a time when upper limits will constrain more and more details about the reionization epoch. Most reasonable models lie only one to two orders of

magnitude away in temperature-squared from the current result, making a detection in the next few years a very real possibility — if not with PAPER or its competitors, than with the planned Hydrogen Epoch of Reionization Array (HERA).[1] And, once this door has opened, the potential for high redshift 21cm observations outlined in Chapter 1 can truly begin to be realized.

[1] http://reionization.org

Bibliography

Adams, J. J., Blanc, G. A., Hill, G. J., et al. 2011, ApJS, 192, 5

Albrecht, A., Bernstein, G., Cahn, R., et al. 2006, ArXiv Astrophysics e-prints, arXiv:astro-ph/0609591

Albrecht, A., Amendola, L., Bernstein, G., et al. 2009, ArXiv e-prints, arXiv:0901.0721 [astro-ph.IM]

Amendola, L., Appleby, S., Bacon, D., et al. 2012, ArXiv e-prints, arXiv:1206.1225 [astro-ph.CO]

Anderson, L., Aubourg, E., Bailey, S., et al. 2013, ArXiv e-prints, arXiv:1303.4666 [astro-ph.CO]

Ansari, R., Campagne, J.-E., Colom, P., et al. 2012, Comptes Rendus Physique, 13, 46

Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, A&A, 61, 99

Barkana, R., & Loeb, A. 2007, Reports on Progress in Physics, 70, 627

Battye, R. A., Browne, I. W. A., Dickinson, C., et al. 2012, ArXiv e-prints, arXiv:1209.0343 [astro-ph.CO]

Beringer, J., Arguin, J.-F., Barnett, R. M., et al. 2012, Phys. Rev. D, 86, 010001

Bolton, J. S., Becker, G. D., Raskutti, S., et al. 2012, MNRAS, 419, 2880

Bolton, J. S., Haehnelt, M. G., Warren, S. J., et al. 2011, MNRAS, 416, L70

Bowman, J. D., Morales, M. F., & Hewitt, J. N. 2009, ApJ, 695, 183

Bowman, J. D., & Rogers, A. E. E. 2010, Nature, 468, 796

Bowman, J. D., Rogers, A. E. E., & Hewitt, J. N. 2008, ApJ, 676, 1

Bradley, L. D., Trenti, M., Oesch, P. A., et al. 2012, ApJ, 760, 108

Burns, J. O., Lazio, J., Bale, S., et al. 2012, Advances in Space Research, 49, 433

Carilli, C. L., Wang, R., Fan, X., et al. 2010, ApJ, 714, 834

Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, Nature, 466, 463

Chang, T.-C., Pen, U.-L., Peterson, J. B., & McDonald, P. 2008, Physical Review Letters, 100, 091303

Clark, M. A., La Plante, P. C., & Greenhill, L. J. 2011, ArXiv e-prints, arXiv:1107.4264 [astro-ph.IM]

Cooray, A., Gong, Y., Smidt, J., & Santos, M. G. 2012, ApJ, 756, 92

Datta, A., Bowman, J. D., & Carilli, C. L. 2010, ApJ, 724, 526

Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, AJ, 145, 10

Dillon, J. S., Liu, A., & Tegmark, M. 2012, ArXiv e-prints, arXiv:1211.2232 [astro-ph.CO]

Dillon, J. S., Liu, A., Williams, C. L., et al. 2013, ArXiv e-prints, arXiv:1304.4229 [astro-ph.CO]
Dodelson, S. 2003, Modern Cosmology (San Diego, CA: Academic Press)
Drinkwater, M. J., Jurek, R. J., Blake, C., et al. 2010, MNRAS, 401, 1429
Edge, D. O., Shakeshaft, J. R., McAdam, W. B., Baldwin, J. E., & Archer, S. 1959, MmRAS, 68, 37
Eisenstein, D. J., & Hu, W. 1998, ApJ, 496, 605
Eisenstein, D. J., Hu, W., & Tegmark, M. 1998, ApJ, 504, L57
—. 1999, ApJ, 518, 2
Eisenstein, D. J., Seo, H.-J., Sirko, E., & Spergel, D. N. 2007, ApJ, 664, 675
Eisenstein, D. J., Zehavi, I., Hogg, D. W., et al. 2005, ApJ, 633, 560
Ellingson, S. W., Clarke, T. E., Cohen, A., et al. 2009, IEEE Proceedings, 97, 1421
Ellis, R., Takada, M., Aihara, H., et al. 2012, ArXiv e-prints, arXiv:1206.0737 [astro-ph.CO]
Fan, X., Narayanan, V. K., Lupton, R. H., et al. 2001, AJ, 122, 2833
Fan, X., Strauss, M. A., Richards, G. T., et al. 2006, AJ, 131, 1203
Field, G. B. 1958, Proceedings of the IRE, 46, 240
Finkelstein, S. L., Papovich, C., Ryan, R. E., et al. 2012, ApJ, 758, 93
Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, Phys. Rep., 433, 181
Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759
Gunn, J. E., & Peterson, B. A. 1965, ApJ, 142, 1633
Harker, G. J. A., Pritchard, J. R., Burns, J. O., & Bowman, J. D. 2012, MNRAS, 419, 1070
Helmboldt, J. F., & Kassim, N. E. 2009, AJ, 138, 838
Helmboldt, J. F., Kassim, N. E., Cohen, A. S., Lane, W. M., & Lazio, T. J. 2008, ApJS, 174, 313
Hibon, P., Malhotra, S., Rhoads, J., & Willott, C. 2011, ApJ, 741, 101
Hinshaw, G., Larson, D., Komatsu, E., et al. 2012, ArXiv e-prints, arXiv:1212.5226 [astro-ph.CO]
Högbom, J. A. 1974, A&AS, 15, 417
Iliev, I. T., Mellema, G., Shapiro, P. R., et al. 2012, MNRAS, 423, 2222
Jacobs, D. C., Aguirre, J. E., Parsons, A. R., et al. 2011, ApJ, 734, L34
Kaiser, N. 1987, MNRAS, 227, 1
Kashlinsky, A., Arendt, R. G., Ashby, M. L. N., et al. 2012, ApJ, 753, 63
Kraus, J. D. 1966, Radio Astronomy (New York, NY: McGraw-Hill)
Laing, R. A., Riley, J. M., & Longair, M. S. 1983, MNRAS, 204, 151
Larson, D., Dunkley, J., Hinshaw, G., et al. 2011, ApJS, 192, 16
Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473
Lidz, A., Zahn, O., McQuinn, M., Zaldarriaga, M., & Hernquist, L. 2008, ApJ, 680, 962
Liu, A., Pritchard, J. R., Tegmark, M., & Loeb, A. 2013, Phys. Rev. D, 87, 043002
Liu, A., & Tegmark, M. 2011, Phys. Rev. D, 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, IEEE Proceedings, 97, 1497
Madau, P., Meiksin, A., & Rees, M. J. 1997, ApJ, 475, 429
Masui, K. W., McDonald, P., & Pen, U.-L. 2010, Phys. Rev. D, 81, 103527

Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, ApJ, 763, L20
McDonald, P., & Eisenstein, D. J. 2007, Phys. Rev. D, 76, 063009
McGreer, I. D., Mesinger, A., & Fan, X. 2011, MNRAS, 415, 3237
McQuinn, M., Hernquist, L., Zaldarriaga, M., & Dutta, S. 2007, MNRAS, 381, 75
McQuinn, M., & O'Leary, R. M. 2012, ApJ, 760, 3
McQuinn, M., Zahn, O., Zaldarriaga, M., Hernquist, L., & Furlanetto, S. R. 2006, ApJ, 653, 815
Moore, D. F., Aguirre, J. E., Parsons, A. R., Jacobs, D. C., & Pober, J. C. 2013, ArXiv e-prints, arXiv:1302.0876 [astro-ph.CO]
Morales, M. F. 2005, ApJ, 619, 678
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., & Wyithe, J. S. B. 2010, ARA&A, 48, 127
Noh, Y., White, M., & Padmanabhan, N. 2009, Phys. Rev. D, 80, 123501
O'Leary, R. M., & McQuinn, M. 2012, ApJ, 760, 4
Ono, Y., Ouchi, M., Mobasher, B., et al. 2012, ApJ, 744, 83
Ord, S. M., Mitchell, D. A., Wayth, R. B., et al. 2010, PASP, 122, 1353
Ouchi, M., Shimasaku, K., Furusawa, H., et al. 2010, ApJ, 723, 869
Paciga, G., Chang, T.-C., Gupta, Y., et al. 2011, MNRAS, 413, 1174
Padmanabhan, N., White, M., & Cohn, J. D. 2009, Phys. Rev. D, 79, 063523
Parashare, C. R., & Bradley, R. F. 2009, in 2009 USNC/URSI Annual Meeting
Parsons, A., Pober, J., McQuinn, M., Jacobs, D., & Aguirre, J. 2012a, ApJ, 753, 81
Parsons, A., Backer, D., Siemion, A., et al. 2008, PASP, 120, 1207
Parsons, A. R., & Backer, D. C. 2009, AJ, 138, 219
Parsons, A. R., Pober, J. C., Aguirre, J. E., et al. 2012b, ApJ, 756, 165
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. 2013, ArXiv e-prints, arXiv:1304.4991 [astro-ph.CO]
Pen, U.-L., Chang, T.-C., Hirata, C. M., et al. 2009, MNRAS, 399, 181
Pentericci, L., Fontana, A., Vanzella, E., et al. 2011, ApJ, 743, 132
Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, ApJ, 517, 565
Peterson, J. B., Bandura, K., & Pen, U. L. 2006, ArXiv Astrophysics e-prints, arXiv:astro-ph/0606104
Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013a, ArXiv e-prints, arXiv:1303.5076 [astro-ph.CO]
—. 2013b, ArXiv e-prints, arXiv:1303.5082 [astro-ph.CO]
—. 2013c, ArXiv e-prints, arXiv:1303.5083 [astro-ph.CO]
Pritchard, J. R., & Loeb, A. 2012, Reports on Progress in Physics, 75, 086901
Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, AJ, 116, 1009
Röttgering, H. 2003, New A Rev., 47, 405
Santos, M. G., Cooray, A., & Knox, L. 2005, ApJ, 625, 575
Schenker, M. A., Stark, D. P., Ellis, R. S., et al. 2012, ApJ, 744, 179

Schlegel, D., White, M., & Eisenstein, D. 2009, in Astronomy, Vol. 2010, astro2010: The Astronomy and Astrophysics Decadal Survey, 314
Schlegel, D., Abdalla, F., Abraham, T., et al. 2011, ArXiv e-prints, arXiv:1106.1706 [astro-ph.IM]
Seo, H.-J., Dodelson, S., Marriner, J., et al. 2010, ApJ, 721, 164
Seo, H.-J., & Eisenstein, D. J. 2007, ApJ, 665, 14
Slee, O. B. 1995, Australian Journal of Physics, 48, 143
Stefan, I. I., Carilli, C. L., Green, D. A., et al. 2012, ArXiv e-prints, arXiv:1212.1624 [astro-ph.CO]
Swarup, G., Ananthakrishnan, S., Kapahi, V. K., et al. 1991, CURRENT SCIENCE V.60, NO.2/JAN25, P. 95, 1991, 60, 95
Switzer, E. R., Masui, K. W., Bandura, K., et al. 2013, ArXiv e-prints, arXiv:1304.3712 [astro-ph.CO]
Tegmark, M., Strauss, M. A., Blanton, M. R., et al. 2004, Phys. Rev. D, 69, 103501
Thompson, A. R., Moran, J. M., & Swenson, G. W. 2007, Interferometry and Synthesis in Radio Astronomy, John Wiley & Sons, 2007.
Tilvi, V., Rhoads, J. E., Hibon, P., et al. 2010, ApJ, 721, 1853
Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2012, ArXiv e-prints, arXiv:1206.6945 [astro-ph.IM]
Totani, T., Kawai, N., Kosugi, G., et al. 2006, PASJ, 58, 485
Trac, H., & Cen, R. 2007, ApJ, 671, 1
Trott, C. M., Wayth, R. B., & Tingay, S. J. 2012, ApJ, 757, 101
Tseliakhovich, D., & Hirata, C. 2010, Phys. Rev. D, 82, 083520
Vedantham, H., Udaya Shankar, N., & Subrahmanyan, R. 2012, ApJ, 745, 176
Vollmer, B., Davoust, E., Dubois, P., et al. 2005, A&A, 436, 757
Wouthuysen, S. A. 1952, AJ, 57, 31
Wright, E. L. 2006, PASP, 118, 1711
Wyithe, J. S. B., Loeb, A., & Geil, P. M. 2008, MNRAS, 383, 1195
Yatawatta, S., de Bruyn, A. G., Brentjens, M. A., et al. 2013, ArXiv e-prints, arXiv:1301.1630 [astro-ph.IM]
Zahn, O., Reichardt, C. L., Shaw, L., et al. 2012, ApJ, 756, 65
Zwaan, M. A., Meyer, M. J., Staveley-Smith, L., & Webster, R. L. 2005, MNRAS, 359, L30