Redundant Baseline Calibration in CHIME

A First Implementation & Application as Beam Probe

by

Deborah C. Good

Bachelor of Science, Colorado School of Mines, 2014

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Astronomy)

The University of British Columbia
(Vancouver)

August 2017

Abstract

The nature of dark energy is one of the most intriguing scientific questions of the twenty-first century. There are many ways to probe dark energy, but one method involves detecting baryon acoustic oscillations (BAO) throughout the universe's history. BAO have a characteristic size scale and therefore act as a "standard ruler," an advantageous property for a method of tracking the universe's expansion history.

While baryon acoustic oscillations can be probed in many ways, one of the most intriguing and promising methods is through twenty-one centimeter hydrogen intensity mapping. Several experiments devoted to twenty-one centimeter hydrogen mapping will be coming on line in coming years, and these experiments have stringent calibration requirements due to the need to remove bright foreground signals. These calibration requirements necessitate new and improved methods for calibration. One proposed method is redundant baseline calibration, a self-calibration method which takes advantage of the massively redundant designs of many hydrogen intensity mapping experiments.

With the Canadian Hydrogen Intensity Mapping Experiment as a test case, we demonstrate that the redundant baseline method is effective in even its simplest implementation for an idealized version of a real telescope. We then show that redundant baseline calibration fails in real CHIME Pathfinder data in a way that is consistent with deviations from redundancy observed in processed CHIME Pathfinder data. These deviations from redundancy are themselves consistent with the effects of feed-to-feed beam pattern variations, a possibility not considered in the conventional redundant baseline calibration algorithm.

We simulate the CHIME Pathfinder including beam width perturbations and verify that similar failures in the redundant baseline calibration can be generated with beam perturbations. We then use the principles of redundant baseline calibration to solve for our simulated beam perturbations. Finally, we compare redundant baseline calibration results to point source holography results and show that the two are equivalent probes of relative feed-to-feed beam variation.

We conclude that redundant baseline calibration is a promising path forward in calibrating hydrogen intensity mapping experiments, both as a conventional calibration method and as a probe of beam structure.

Lay Summary

One goal for cosmology in the twenty-first century is to understand dark energy and the accelerating expansion of the universe. One way to learn about dark energy is by detecting baryon acoustic oscillations (BAO) with 21 cm hydrogen intensity mapping. The Canadian Hydrogen Intensity Mapping Experiment (CHIME) is an experiment designed for this purpose.

Hydrogen Intensity Mapping is difficult, because the desired signal is much dimmer than other sources of detected light. This means that experiments must be calibrated carefully to allow the removal of unwanted sources. One method for calibration is redundant baseline calibration (RBC), which takes advantage of the design of telescopes like CHIME to precisely calibrate without detailed knowledge of the sky or electronics.

RBC is successful in idealized situations, but real telescopes break assumptions underlying the algorithm making RBC inaccurate. However, these inaccuracies predict important properties of the telescope, so it is still a useful tool.

Preface

This thesis is based on work conducted as part of the CHIME Collaboration and the CHIME Pathfinder experiment specifically. None of the text of this thesis is taken directly from previously published articles. The analysis in this thesis is the work of D. Good with the supervision of Kris Sigurdson and J. Richard Shaw as well as incidental input from other members of the CHIME Collaboration.

Several CHIME software modules were used to create the simulated CHIME Pathfinder data in Chapters 2 and 4. These modules were caput, cora, ch_pipeline, ch_util, draco, driftscan, and were developed by members of the CHIME collaboration, largely Kiyoshi Masui and J. Richard Shaw.

The holography data in Chapter 5 was collected by the CHIME collaboration at large and Philippe Berger specifically. Holographic beam mapping for the CHIME Pathfinder is discussed in more detail in [6].

Table of Contents

List of Figures

Acknowledgments

This thesis would never have been possible without the help of a legion of people, both those who supported me scientifically, and those who supported me personally. While I could by no means list everyone who has had a role in my success to date, I will take a moment to acknowledge some of the most pertinent.

First, no thesis is ever completed without the work of a dedicated supervisor. Thank you, Kris, for introducing me to the wonderful world of redundant baselines and to my entire "theory of experiment" niche and for helping me develop as an independent scientist.

Thank you, Richard, for your continual willingness to help, both in figuring out the science itself and in my never-ending stream of software issues. Also, thank you for teaching me how to drive the CHIME car; sorry I left you at the DRAO gate.

Thank you to the entire CHIME collaboration for both providing direction for my thesis and also pushing me to always be able to articulate what I'm working on and why it's important. In particular, thank you to Gary and Mark for serving as my "step-advisors," and to my fellow UBC CHIME graduate students, who have supported me both scientifically and emotionally. Also, thank you to the quadruped graduate student, Looney, for letting me scratch her ears and rant when I was having a bad day.

Finally, thank you to all of my wonderful family and friends, who supported me in a myriad of ways and never wavered in their belief that I would be an excellent physicist.

Chapter 1

A Brief Introduction to ΛCDM Cosmology & 21 cm Hydrogen Intensity Mapping

Before delving into specific discussions of redundant baseline calibration in the Canadian Hydrogen Intensity Mapping Experiment, it is worth discussing more broadly the scientific motivation for and design of CHIME.

We will begin by discussing the theoretical underpinnings of and experimental evidence for dark energy. With this general background and motivation, we will shift our focus to twenty-one centimeter intensity mapping, the technique CHIME will use, and discuss the theory, potential applications, and a few current and upcoming experiments using twenty-one cm intensity mapping. Finally, we will discuss the specific design of CHIME and its unique calibration requirements.

This will provide a structure for the more specific discussion of redundant baseline calibration in Chapter 2 - 5.

1.1 An Introduction to Dark Energy

1.1.1 Vacuum Energy & the Cosmological Constant

The "cosmological constant problem" looms over the past century's attempts to explain the large-scale universe. Shortly after Albert Einstein formulated general relativity in 1915-1916, he attempted to apply his theory to the universe at large, assuming as was commonly thought at the time, that the universe was static. He wrote, "The most important fact that we draw from experience is that the relative velocities of the stars are very small as compared with the velocity of light" [38]. However, Einstein failed to discover a static solution to his equations and therefore resorted to the introduction of a cosmological constant, altering the Einstein equation to be

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R - \lambda g_{\mu\nu} = -8\pi G T_{\mu\nu}. \tag{1.1}$$

In the early 1920s, Alexander A. Friedmann introduced a metric which removed the need for the cosmological constant, the metric which would become known as the FRW metric and which would become the fundamental metric for cosmology in the twentieth century [38]. George Gamow later wrote that Einstein considered that "the introduction of the cosmological term was the biggest blunder he ever made in his life" [1].

However, though the cosmological advances of the 1920s indicated that the cosmological constant was not necessary for the reasons Einstein originally proposed, the idea never entirely left physics. Indeed, by the late 1980s, it was again a topic of serious discussion. Studying the standard model of particle physics led to increased concern about vacuum energy, the state of least energy density in standard model particle physics [1]. Such a vacuum energy is composed of a bare cosmological constant, the value the cosmological constant would take without any matter in the universe, and the energy density arising from quantum fluctuations [1].

Particle physicists define vacuum as a ground state, and therefore conclude the vacuum must be Lorentz invariant. This in turn means the stress-energy-momentum tensor must be proportional to a Minkowski metric. Knowing that the stress-energy-momentum tensor of a perfect fluid has the diagonal (ρ, P, P, P), we can conclude that the vacuum is a perfect fluid with the equation of state $P_{vac} = -\rho_{vac}$. Assuming adiabatic compression and expansion, ρ_{vac} remains constant and is related to a cosmological constant Λ by $\Lambda = 8\pi G\rho_{vac}$ [8].

This vacuum energy is irrelevant in classical, non-gravitational physics. However, it becomes relevant in quantum theory. We can generalize our description of the vacuum energy to a quantum field theory formulation by modeling a relativistic field as a collection of harmonic oscillators. This allows us to write the vacuum energy as

$$E_0 = \sum_j \frac{1}{2}\hbar\omega_j, \tag{1.2}$$

summing over possible modes of the field. Suppose the system is in a box with side L and that L goes to infinity. We can impose periodic boundary conditions and reformulate Equation 1.2 as

$$E_0 = \frac{1}{2}\hbar\, L^3 \int \frac{d^3\mathbf{k}}{(2\pi)^3}\omega_{\mathbf{k}}. \tag{1.3}$$

To obtain ρ_{vac}, we allow $L \to \infty$ and divide by L^3. We set $\omega_k^2 = k^2 + m/\hbar^2$ and set $k_{max} \gg m/\hbar$. Then we can write an expression for ρ_{vac}

$$\rho_{vac} \equiv \lim_{L\to\infty} \frac{E_0}{L^3} = \hbar\, \frac{k_{\max}^4}{16\pi^2} \tag{1.4}$$

As k approaches ∞, ρ_{vac} diverges. This is not uncommon for low-energy theories at high k, so we treat $k_{\max}$ to be the energy scale at which we remain confident in our theory. This is commonly chosen based on the Planck energy $E^* \approx 10^{19} GeV$, so $k_{\max} = E^*/\hbar$. This indicates that

$$\rho_{vac} \approx 10^{74}\ GeV^4\ ,\hbar^{-3}. \tag{1.5}$$

However, this value is approximately 120 orders of magnitude higher than the value expected from observation [8].

For cosmologists of the late twentieth century, this represented a problem as both theory and evidence from the cosmic microwave background indicated that the total energy density of the universe was close to critical. Thus, theorists concluded that the missing energy density, about 2/3 of critical, must be some kind of smooth unknown energy called "dark energy" [12].

One way to describe the amount energy density in the universe allocated to dark energy is by a cosmological constant Ω_Λ, which should be related to the matter density of the universe by $\Omega_\Lambda + \Omega_m = 1$, where Ω_m is the mass density of the universe.

One issue with thsi Ω_Λ description is that it fails to account for explanations of dark energy apart from a cosmological constant. Therefore, we can also describe dark energy starting from energy conservation in an expanding universe,

$$\frac{\partial \rho}{\partial t} + \frac{\dot{a}}{a}[3\rho + 3P] = 0, \tag{1.6}$$

and defining the quantity

$$w = \frac{P}{\rho}. \tag{1.7}$$

The value of w allows us to describe a variety of universes, not just one dominated by a cosmological constant. For $w = -1$, we have a cosmological constant. For $w = 0$, we have a flat, matter-dominated universe and for $w = 1/3$, we have a radiation dominated universe [12]. We can then write an equation for the evolution of dark energy [12]

$$\rho_{de} \propto \exp\left(-3\int^{a} \frac{da'}{a'}[1 + w(a')]\right). \tag{1.8}$$

1.1.2 Experimental Probes of Dark Energy

Type Ia Supernova

In the last years of the twentieth century, new observational evidence began a new era in the cosmological constant discussion. Observations of Type Ia Supernovae in the late 1990s and early 2000s indicated that the universe was expanding at an accelerating rate, providing evidence that the vacuum energy density Ω_Λ was indeed greater than zero.

Type Ia Supernova are considered standard candles, because their luminosity is well-known. In cosmology, standard candles are highly useful, as they allows cosmologists to calculated the distance to the standard candle object. This process requires a quantity known as the luminosity distance. An object with a given luminosity L a distance d from the observer has a flux

$$F = \frac{L}{f\pi d^2}. \tag{1.9}$$

In comoving coordinates, the flux is

$$F = \frac{L(\chi)}{4\pi\chi^2(a)}, \tag{1.10}$$

where $L(\chi)$ is the luminosity in a comoving spherical shell of radius $\chi(a)$ and a is the scale factor [12]. If all emitted photons have the same energy, $L(\chi)$ is the times the number of photons passing through the comoving spherical shell in a unit of time. Due to the expansion of the universe, the energy of the photons will be less today than when they were emitted and the energy per unit time in our comoving shell will be a^2 smaller than the source luminosity. This means we can write

$$F = \frac{La^2}{4\pi\chi^2(a)}, \tag{1.11}$$

which makes apparent a way to return to the simple form of equation 1.9: define

$$d_L = \frac{\chi}{a} \tag{1.12}$$

the luminosity distance [12]. If we know a luminosity distance and an absolute magnitude of a source, we can find its apparent magnitude using the expression

$$m - M = 5\log\left(\frac{d_L}{10\mathrm{pc}}\right) + K, \tag{1.13}$$

where K is is correction for wavelength shift due to expansion, m is apparent magnitude, and M is absolute magnitude [12].

Type Ia supernovae have almost identical absolute magnitudes, so any set of Type Ia supernova magnitudes can be easily compared. Scott Dodelson's *Modern*

Cosmology textbook offers a brief illustrative example. Given two supernovae, one at $z = 0.83$ with $m = 24.32$ and one at $z = 0.026$ with $m = 16.08$ and the fact that the absolute magnitudes are the same (so the difference in apparent magnitude is due only to luminosity distance), we can write

$$24.32 - 16.08 = 5\log\left(d_L(z = 0.83)\right) - 5\log\left(d_L(z = 0.026)\right). \tag{1.14}$$

The smaller luminosity distance is sufficiently low redshift to be written as $d_L(z = 0.026) = z/H_0 = 0.026/H_0$, making the larger distance our only unknown. Our observation indicates that $H_0\, d_L(z = 0.83) = 1.16$, but for a matter dominated universe ($\Omega_m = 1,\ \Omega\Lambda = 0$), $H_0\, d_L(z = 0.83) = 0.95$. Even our toy problem indicates that we must add a cosmological constant of some sort [12].

This same conclusion was reached, much more rigorously by several supernova surveys, beginning with Riess et. al. and Perlmutter et. al. in 1998 and 1999 [26, 29]. These surveys sampled small numbers of high redshift supernovae (around 10 each) and fit for cosmological parameters Ω_M and Ω_Λ, probing whether the universe was matter or radiation dominated. Each group found that there was a significant Ω_Λ component, although the Riess et. al. value was significantly higher.

Since the late 1990s, many more supernova surveys have been conducted, which have corroborated these initial findings and added more supernovae to the total sample. Some notable results include those of the Supernova Cosmology Project, the Sloan Digital Sky Survey, and the Hubble Space Telescope Cluster Supernova Survey, and the Supernova Legacy Survey [13, 15, 34, 39].

Relatively recent supernova surveys have used improved datasets to find improved values for w and Ω_Λ. In 2011, Conley et. al. found $w = -0.91^{+.16}_{-.20}(\text{stat})^{+.07}_{-.14}(\text{sys})$ using only supernova data, assuming a flat universe with constant w, marginalizing over Ω_m [11].

Supernova surveys are a foundational method of measuring Ω_Λ, but there is significant evidence that they are limited by systematic uncertainties. These arise from a variety of quarters including dust extinction, supernova colours, and photometric calibration [37].

Cosmic Microwave Background

While not very useful as a direct probe of dark energy, the cosmic microwave background (CMB) is important in that it constrains other cosmological parameters precisely. In particular, the CMB closely constrains the values of $\Omega_m h^2$ and $\Omega_b h^2$, which are other major components of the energy content in the universe, particularly if the universe is assumed to be flat (i.e. $\Omega_k = 0$). Additionally, the CMB can be used to probe cosmic acceleration models more directly via the Integrated Sachs-Wolfe effect [37].

The CMB's use as a provider of careful measurements is also important in the context of the baryon acoustic oscillations, which are another experimental probe discussed below. CMB anisotropy data allows careful measurement of the physical scale of oscillations deriving from baryon density, which allows the oscillations to be used as a standard ruler. Without the careful constraints from the CMB such analysis might be impossible [31].

Baryon Acoustic Oscillations

Immediately after inflation, the baryons and photons in the universe were locked together, and reactions between the photon pressure and matter density fluctuations produced sound waves. At recombination, the baryons and photons decoupled. However, the previously generated sound waves remained frozen in the structure of the baryonic matter. These fluctuations are known as baryon acoustic oscillations or BAO. The size of the BAO was established by the size of the sound horizon at recombination, the distance that sound could travel by the time of last scattering. This created a characteristic, fixed scale which can be used as a "standard ruler" on the universe [3, 24, 37].

Use of BAO as a standard ruler is significant to the study of dark energy, as tracing baryon acoustic oscillations provides a definite size scale, allowing determination of the angular diameter distance and the Hubble parameter at a variety of redshifts [3]. The comoving distance to an object in the line-of-sight and transverse directions can be written as

$$r_{\parallel} = \frac{c\Delta z}{H(z)} \tag{1.15}$$

and

$$r_\perp = (1+z)\; D_A(z)\; \Delta\theta, \tag{1.16}$$

where $\Delta\theta$ and Δz are the observed dimensions of the object [31]. If we know $r_\perp$ and $r_\parallel$ as is the case when measuring the BAO scale, we are able to determine $D_A(z)$ and $H(z)$.

This method may be preferable to other methods such as supernovae and clustering because it is minimally affected by non-linear gravitational clustering, galaxy biasing, and redshift distortions [31]. Therefore, it could be a vey clean probe of dark energy.

The baryon acoustic oscillation can be detected in multiple ways. All methods depend on a clean determination of the BAO scale by CMB experiments such as Planck, currently estimated to be 147.50 ± 0.24 [2]. To date, there have been several surveys to detect baryon acoustic oscillations from galaxies, many with an eye towards dark energy constraints. These include WiggleZ, BOSS, SDSS main galaxy survey, and the 6dF survey [4, 7, 17, 30].

Weak Lensing

Gravitational lensing is a phenomenon in which mass between an astronomical object and its observer distorts the image of the background source, deflecting it from its proper position by some angle. This angle is generally too small to observe directly, but weak lensing analysis can observe the gradient of the angle, which makes circular galaxies appear elliptical [3]. Weak lensing is a small distortion (about 1%) in the size and shape of images of distant galaxies, generated by distortion from light bending past galaxies or clusters of galaxies in the foreground of the lensed galaxy, It can be used to measure either the galaxy's shearing (distortion in shape) or magnification (distortion in size), though shearing measurements are far more common [37].

Individual galaxies are not circular, so we cannot determine the deformation and thus the lensing signal from any individual galaxy. Instead, we compile large samples of galaxies and detect a lensing signal as a pattern of aligned shapes in a region of galaxies. This means that weak lensing measurements require a very large galaxy survey [3].

Weak lensing can be used as a probe of dark energy because the deflection angle from the expected position of the lensed object depends on the mass of the foreground lensing object and the distances between the observer, the lens, and the lensed source. With careful analysis, then, weak lensing can be used to constrain both the angular-diameter distance as a function of redshift and the growth rate of structure [3].

However, weak lensing measurements are challenging. They require a large sample of lensed objects. They also have very complicated error analyses, requiring the consideration of statistical errors arising from cosmological model, concerns about observational bias, and systematic errors arising from astrophysical processes in the sources and intervening space [37].

Galaxy Clusters

Measuring galaxy clusters is one of the oldest techniques for investigating dark energy, as measurements from galaxy clusters indicated that $\Omega_m < 1$ as early as the 1980s. Today, galaxy clusters continue to be important ways of understanding dark energy. They are the largest gravitationally collapsed objects in the universe, marking locations with large density fluctuations in the early universe. We can analytically predict their mass function, the number of galaxy clusters per unit comoving volume per cluster mass. We can then measure their actual abundance in a region of the sky and compare the two [3].

Cluster results scale sensitively with both the comoving volume and the matter density of the universe, meaning that cluster results are sensitive to changes in either the matter density or the size of the universe, both of which are of relevant observables for dark energy [3]. In particular, the halo mass function which serves as a predictor for cluster abundances is sensitive to the combination of cosmological parameters $\sigma_8\Omega_m$, and is therefore very sensitive to the initial value of w. Small changes in w can alter the evolution of the universe at late times, in turn altering the development of clusters [37].

Galaxy clusters can and have been detected in a variety of ways, including x-ray emissions the Sunyaev-Zeldovich effect, and even gravitational lensing. However, the primary challenge in using galaxy clusters to study dark energy is not detecting the clusters but understanding the relationship between the observables of the galaxy cluster and a cosmological model. Galaxy cluster measurements are not directly measuring the mass of the cluster, but rather proxies like x-ray flux or galaxy counts [3]. Calibrating the observable-mass relationship is the largest challenge facing cluster analysis. It can be accomplished by simulations of galaxy clusters, by extrapolating the few direct mass measurements to larger samples, or by relying on statistical methods involving either additional observables or weak lensing [37].

1.2 21 cm Intensity Mapping

One of the most exciting new categories of cosmology experiments is 21 cm hydrogen intensity mapping experiments. While challenging, 21 cm experiments options for examining a variety of cosmological questions, including both dark energy via the universe's accelerating expansion and the epoch of reionization.

1.2.1 Understanding 21 cm Neutral Hydrogen

The 21 cm line of hydrogen arises from the hyperfine transition of hydrogen, during which the electron spin flips. It was a triumph of theoretical astrophysics when discovered in the 1940s, one of very few spectral lines discovered following a precise theoretical prediction. At redshift 0, the line is at $\nu = 1420.4057$ MHz.

It is potentially useful to cosmology for a few reasons. First, it is a spectral line, so by mapping it at different frequencies (different redshifts), it is possible to trace its full three dimensional evolution. For example, at the era when dark energy becomes dominant (around $z = 1-3$), the 21 cm line is found at $\nu \approx 400-800$ MHz and at the Epoch of Reionization, it is found around $\nu \approx 30-200$ MHz [14]. Second, at higher redshift, it probes the IGM, which is the dominant location of baryonic matter [14]. At lower redshifts, it probes the clustering of collapsed halos and therefore the underlying matter density distribution [28]. At either epoch, it serves as a tracer of baryonic matter in the universe.

Mapping 21 cm hydrogen emission is a potentially useful cosmological probe for several epochs. It may provide a way to examine the dark ages, between recombination and reionization, which are difficult to observe, as they precede the formation of astrophysical objects. It may also be a useful way of examining the epoch of reionization, with the advantage that appropriate frequency coverage allows maps at a progression of redshifts. Finally, at more recent times, it can be used to detect the BAO and therefore constrain dark energy, but without requiring surveys to resolve individual galaxies [14, 28].

1.2.2 21 cm Intensity Mapping as a Probe of Dark Energy

One method for constraining dark energy with BAO is by using 21 cm intensity mapping. This method removes one of the significant challenges in measuring the BAO at appropriate redshifts: it does not require detection of individual galaxies, just large scale variations in HI mass [9]. Therefore, it requires less resolution than galaxy surveys with the same goal. The smallest spatial scale such an experiment would need to consider is the third BAO peak, past which nonlinear evolution attenuates BAO structure. Its wavelength is $35h^{-1}$Mpc, meaning that a Nyquist sampled map need only have $18h^{-1}$ MPc sized pixels. For relevant redshifts such as $z = 1.5$, this means an angular wavelength of 20 arcminutes, corresponding to a 200 wavelengths or 100 m, is necessary to resolve BAO structure [9].

The mean brightness temperature of the 21 cm line at redshift $z = 1 - 3$ can be estimated as

$$T_b = 0.3 \left(\frac{\Omega_{\mathrm{HI}}}{10^{-3}}\right) \left(\frac{\Omega_m + a^3\Omega_\Lambda}{0.29}\right)^{-1/2} \left(\frac{1+z}{2.5}\right)^{1/2} \mathrm{mK}, \tag{1.17}$$

where Ω_m and Ω_Λ are cosmological parameters and $\Omega_{\mathrm{HI}} \approx 1 \times 10^{-3}$ at $z = 1$ [9, 28]. Calculating the mean brightness temperature for 21 cm hydrogen at $18h^{-1}$ Mpc scales indicates that signal should be expected to be about 150μK, dim relative to foregrounds. This means that detecting BAO will require careful removal of foregrounds and Fourier analysis of large sections of sky. This foreground problem is common to all 21 cm experiments and has been thought about carefully by many in the field, e.g. [18].

Bright foregrounds, nonlinearity, and limits on observable volume all constrain the redshift space in which 21 cm intensity mapping to detect BAO is useful. Though finite, this area is large enough to be cosmologically interesting. Intensity mapping is a viable approach between redshifts of about $z = 0.5 - 2.5$ and for scales on the order of $k = 0.01 - 0.1$. At low redshift, the boundary is drawn by limited observable volume, at high k the limit arises from nonlinearity obscuring BAO wiggles, and the lower k limit arises from bright foregrounds making foreground removal infeasible [9].

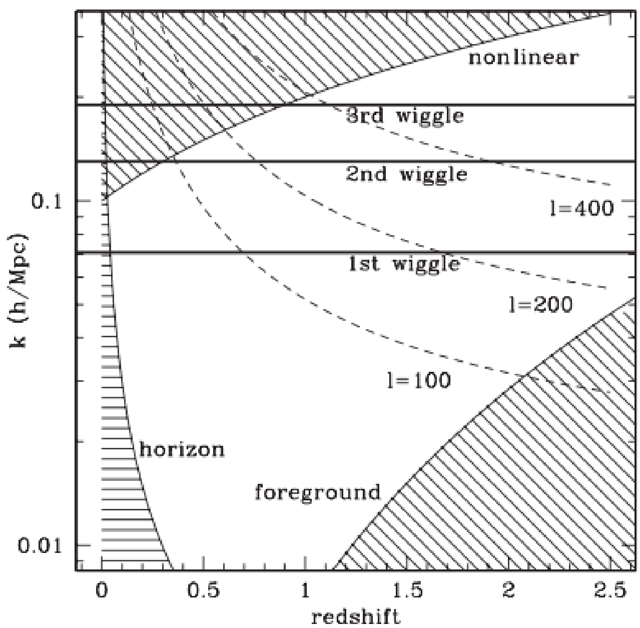

Figure 1.1: Achievable parameter space for BAO detection with 21 cm intensity mapping. The left exclusion arises from limited observable volume, the top exclusion from nonlinearity obscuring the BAO wiggles, and the bottom exclusion from bright foregrounds exceeding the removable threshold. Image from [9].

After detecting the BAO in neutral hydrogen, the analysis proceeds as in other dark energy experiments using BAO analysis. Calculations from Chang et.al. forecast that BAO from intensity mapping combined with Planck could constrain key dark energy parameters to the same level as other "Stage 2" methods, as outlined in the Dark Energy Task Force Report [9]. See Figure 1.2 for a graphical representation of the confidence intervals for potential IM + Planck results.

CHIME, the Canadian Hydrogen Intensity Mapping Experiment, is one exam-

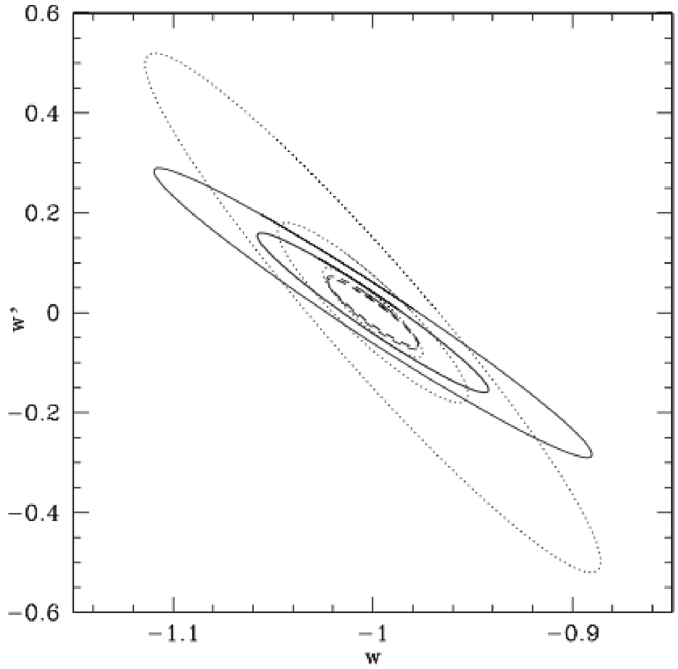

Figure 1.2: Constraints on the dark energy equation of state and its redshift dependence from DETF stage I + Planck (outermost line), DETF stage I+ III + Planck (intermediate dotted line), DETF stage I + III + IV + Planck (inner dotted line), intensity mapping + Planck (inner solid line for best case, outer solid line for worst case, and all options combined (dashed lines for best and worst - almost indistinguishable). Image from [9].

ple of an experiment using intensity mapping to detect the BAO and constrain dark energy. It will be discussed further in Section 1.3.

1.2.3 Designing 21 cm Experiments

While 21 cm intensity mapping can be used for more than one cosmological purpose, depending on a given experiment's frequency range, there are common principles present in many 21 cm experiments. All 21 cm experiments are mapping experiments. All are advantaged by having significant frequency coverage, corresponding to significant redshift coverage. All are concerned about removal of foregrounds.

In principle, 21 cm signal can be observed with a single dipole. Indeed, structure in the 21 cm signal has been detected using only single dish telescopes such as Parkes Observatory and Green Bank Telescope [10, 25]. However, such surveys are not ideal for creating multi-frequency full sky maps of the 21 cm emission. As the objective is ultimately a full sky map, steerability is unimportant, so transit telescopes, in which antenna are in fixed positions and the sky rotates past them, are common. Additionally, heavily redundant array configurations are preferred, generally with close-packed feeds. Most 21 cm experiments do incorporate reflectors, either cylinders or dishes, but in principle a sufficiently large block of antenna could be used as a 21 cm experiment. Many attempt to make use of modern computation efficiency to either rapidly process cross-correlations or to limit the amount of correlations necessary to proceed, as discussed in [35, 36].

Though 21 cm intensity mapping is being used for a variety of cosmological applications, the similarity in experimental design and challenges such as foreground removal allow experiments even with different science goals to share best practices in design, calibration, and data analysis.

1.2.4 A Brief Census of 21 cm Experiments

While 21 cm cosmology is still a new field, a few experiments have already been conducted. These were predominantly smaller, less ambitious experiments designed to inform larger, upcoming experiments, but they were still worth noting as introductions to the field.

In the higher frequency, late-time expansion focused range, there have been both dedicated experiments and surveys involving conventional radio telescopes. The first of these was the Pittsburgh Cylinder Telescope, a very early cylinder telescope prototype, was composed of two 10 m by 25 m cylinders 25 m apart, with a fixed, transit telescope design and sixteen dipoles per cylinder [27]. Early efforts at intensity mapping in this band also included survey at the Green Bank Telescope and the Parkes Telescope, with both surveys detecting the 21 cm signal [10, 25]. At present, the CHIME Pathfinder is an active telescope in this frequency range [5].

There have also been several first generation or pathfinder experiments in the lower frequency, higher redshift range. Perhaps the most prominent is the Precision Array for Probing the Epoch of Reionization, located in South Africa. PAPER pioneered many practical elements of EoR 21 cm analysis and also created a competitive map at their frequency range [23]. Other firsts generation telescopes in this frequency range include MITEoR, 21 CentiMeter Array, LOw Frequency ARray, the Giant Metrewave Radio Telescope EoR Experiment and the Murchison Widefield Array [20, 22, 40–42]

Additionally, upcoming years will see more 21 cm experiments coming online, investigating both the higher redshift EoR era and the lower redshift late-time expansion era. At higher frequencies, CHIME in the northern hemisphere and HIRAX in the southern hemisphere will focus on detecting BAO with 21 cm measurements. At lower frequency, HERA and SKA Low will be able to learn more about the Epoch of Reionization.

1.3 An Introduction to CHIME

The Canadian Hydrogen Intensity Mapping Experiment is a 21 cm intensity mapping experiment designed to detect baryon acoustic oscillations and thereby constrain dark energy as discussed in previous sections. CHIME is a collaboration between University of British Columbia, University of Toronto, McGill University, and the Dominion Radio Astrophysical Observatory (DRAO). It is located at DRAO in Kaleden, BC.

CHIME is a transit telescope, composed of cylindrical reflectors with dual polarization feeds located along a focal line suspended above the cylinder surface. The CHIME Pathfinder, pictured in Figure 1.3, consists of two 20 m by 36 m cylinders with 128 dual-polarization feeds. It has been operational since 2014, so data from the Pathfinder can be used to test analysis and calibration methods. Full CHIME, pictured in Figure 1.4, is currently under construction but expected to begin taking data later in 2017. It consists of four 20 m by 100 m cylinders with 1024 dual-polarization feeds. On both telescopes, these dual-polarization feeds are spaced about 30 cm apart on the focal line, in a manner that is intentionally highly redundant.

Figure 1.3: The CHIME Pathfinder at DRAO. It consists of two 20 m by 36 m cylinders and 256 total inputs, and operates between 400-800 MHz.

Figure 1.4: CHIME under construction at DRAO. It consists of four 20 m by 100 m cylinders and 2028 total inputs. It will operate at the same frequency range as the CHIME Pathfinder.

The Pathfinder and CHIME both observe at a frequency range from 400-800 MHz. This range is covered in 1024 390 kHz channels distributed evenly across the band. As observing neutral hydrogen at varied frequencies is equivalent to observing it at different redshifts, CHIME observes from $z = 0.8 - 2.5$, a range which is of significant importance in observing baryon acoustic oscillations to constrain dark energy.

Both telescopes see a narrow "hotdog shaped" primary beam, approximately 100 degrees north-to-south and approximately two degrees east-to-west. As CHIME is a transit telescope, it does not point at specific objects. Instead, the CHIME beam passes over the entire northern sky in these two degree strips as the earth rotates. CHIME's frequency dependent resolution is approximately 0.25-0.5 [21].

Like other 21 cm experiments, foreground removal is a major concern for CHIME. To successfully measure BAO, CHIME must be not allow the systematic errors from foreground filtering or calibration to dominate statistical errors. The foreground signals in particular are quite bright, up to 700 K and generally between 10-20 K. In contrast, 21 cm signal from BAO is expected to be about 100 μK [21]. The stringent foreground filtering necessary to detect the BAO requires excellent calibration techniques.

CHIME will require precise calibration in several different instrumental components. First, CHIME must know the beam response precisely. Frequency dependent structure in the antenna beam will transform angular structure in the foregrounds into spectral structure to CHIME visibilities. Polarized foregrounds will also undergo Faraday rotation and introduce spectral structure to the visibilities. Simulations indicate that CHIME must understand the beam width to approximately 0.1% to avoid biasing the derived power spectrum by an amount greater than statistical uncertainty [21, 32].

CHIME must also calibrate the relative complex gain as a function of time to achieve acceptable brightness accuracy. End-to-end simulations again shine light on the maximum random variations allowed in the complex gains. These simulations indicate that random gain variations must be less than 1% on 60 second timescales [21, 32]. CHIME is, however, saved from conducting an absolute calibration as our BAO measurements do not require an absolute sky brightness value.

Cross-talk, coupling between channels is expected to occur at various points on the CHIME system, including between feeds, cables, and digitzer and correlator boards. Cross-talk in cross-correlation measurements will add signal and noise with stable coefficients. This cross-talk will not appear as a constant additive offset but as a signal at random phases. CHIME will have to measure cross-talk coefficients and then extract the effects from the data [21].

Finally, CHIME will need to understand the instrument passband to a part within 10^5. CHIME will accomplish this by assuming that bright foreground regions of our maps have smooth frequency spectrums, but passband effects are intrinsically tied to beam calibration and therefore cannot be fully estimated without a solid estimate of beam calibration [21].

Chapter 2

Implementing Redundant Baseline Calibration

As discussed in Chapter 1 as well as [21], [32], and [18], 21 cm intensity mapping experiments such as CHIME have stringent calibration requirements, to allow them to carefully remove foregrounds and ultimately extract cosmological information. There are a myriad of potential calibration methods, but they can roughly be combined into three categories.

The first is point source or other sky based calibrations. Point source calibration involves using the telescope to view a source which is very well-known and comparing the data to the known solution to extrapolate instrumental factors. Point source calibration, however, is constrained by external knowledge of the point source in question.

The second method is noise source calibration. In this method, an external source of noise is broadcast into the interferometer and switched off and on at short intervals. Data from the intervals where the noise source is off is then compared to data where the noise source is on, eliminating sky data and leaving only instrumental information. This can be the most analytically straightforward method (although is not always, depending on the noise source configuration) but requires external hardware and can therefore be more complicated than it initially appears.

The third category of calibration method is self-calibration, wherein the intrinsic properties of the telescope and the data ordinarily collected are used to calibrate the telescope. Methods in this category, including redundant baseline calibration, are desirable because they require neither detailed knowledge of the sky nor external noise sources, but instead use the properties of the instrument itself to calibrate. This means it can be effective even in the absence of sufficiently well-known point sources or external noise sources. However, self-calibration methods are generally relative and unable to determine the absolute calibration of an instrument, so they are often used in conjunction with other methods.

Redundant baseline calibration, specifically, exploits the property that there are many baselines in CHIME which generate redundant information. To understand what we mean by redundant baselines and why they are potentially useful for calibration, we must also understand the structure of data in CHIME, discussed in Section 2.1. In Sections 2.2 and 2.3 we discuss the structure of the redundant baseline algorithm and in Section 2.4 we discuss an implementation of the amplitude calibration on data from a simulation of the CHIME Pathfinder.

2.1 Constructing Visibilities

At its most fundamental level, a transit telescope like CHIME can be viewed as a collection of feeds measuring electric field from the sky. To keep track of polarization information, we would like to know what the contribution to the electric field from every direction is and therefore define ε, which is an electric field density in a frequency interval and solid angle. We can then write the electric field as

$$d\mathbf{E} = (\mu_0 c)^{1/2}\, \varepsilon\,(\hat{\mathbf{n}}, \nu)\, d^2\hat{n}' d\nu \tag{2.1}$$

and subsequently write the Poynting Vector as

$$\mathbf{S}_p = \frac{1}{\mu_0 c}\mathbf{E}\times\mathbf{H} = \int d^2\hat{n}\; d^2\hat{n}'\; d\nu\; d\nu'\; \hat{\mathbf{n}}\; \langle \varepsilon(\hat{\mathbf{n}})\cdot\varepsilon(\hat{\mathbf{n}}')\rangle \tag{2.2}$$

.

The astronomical radio signals we are interested in are generally incoherent, enabling us to define ε in terms of the Stokes parameters

$$\begin{aligned}\langle \varepsilon_a(\hat{\mathbf{n}},\nu)\,\varepsilon_b^*(\hat{\mathbf{n}}',\nu)\rangle = \frac{2k_B}{\lambda^2}\delta\left(\hat{\mathbf{n}}-\hat{\mathbf{n}}'\right)\delta\left(\nu-\nu'\right)\times[P^I_{ab}I(\hat{\mathbf{n}})+\\ P_a b^Q Q(\hat{\mathbf{n}})+P_a b^U U(\hat{\mathbf{n}})+P_a b^V V(\hat{\mathbf{n}})],\end{aligned} \tag{2.3}$$

where indices *ab* represent basis vectors transverse to the line of sight. The polarization tensors P^X_{ab} are known variations on the Pauli matrices.

Any feed i on the telescope is collecting a weighted combination of the electric fields it receives, designated F_i and given by

$$F_i(\phi) = \int d^2\hat{n}A^a_i(\hat{\mathbf{n}},\phi)\varepsilon_a(\hat{\mathbf{n}})e^{2\pi i\hat{\mathbf{n}}\cdot\mathbf{u_i}(\phi)}, \tag{2.4}$$

where ϕ is a rotation angle, and u_i is a physical location. This beam has a solid angle $\Omega_i = \int d^2\hat{n}|\mathbf{A}(\hat{\mathbf{n}})|^2$

Real interferometric data , however, is not presented as individual feed responses, but as visibilities, cross correlations between multiple feeds: $V_{ij} = \langle F_i F_j^*\rangle$. These visibilities can be written as

$$V_{ij}(\phi) = \int \sum B^S_{ij}(\hat{\mathbf{n}},\phi)S(\hat{\mathbf{n}}), \tag{2.5}$$

where S represents each Stokes Parameter (I, Q, U, V) and B^S_{ij} is a beam transfer function. These beam transfer functions compile all the necessary instrumental information into one quantity, which is given by

$$B^S_{ij}(\hat{\mathbf{n}},\phi) = \frac{2}{\Omega_{ij}}A_i(\hat{\mathbf{n}},\phi)A^*_j(\hat{\mathbf{n}},\phi)P^S_{ab}e^{2\pi i\hat{\mathbf{n}}\cdot\mathbf{u_{ij}}(\phi)} \tag{2.6}$$

where A_i is the antenna response, $\Omega_{ij} = \sqrt{\Omega_i\Omega_j}$, P^S_{ab} is the previously mentioned polarization tensor, and $\mathbf{u_{ij}}$ is a vector of the length of the baseline between feeds i and j [32].

While a number of factors influence the actual value of the visibility, relatively few change within a given telescope. We assume, quite reasonably, that the sky varies little relative to the spatial scale of our telescope in a given time step, so for a given rotation angle ϕ, the visibilities should vary based on beam transfer matrices rather than sky variation. Within the beam transfer function, if we assume that all feeds have the same beam pattern A, we are left with a function that varies only based on baseline distance $\mathbf{u_{ij}}$. Therefore, any two sets of feeds with the same baseline distance $\mathbf{u_{ij}}$ should produce the same visibility, i.e. we expect the information from these baselines to be redundant. One consequence of redundancy is that any information that is not identical amongst redundant visibilties must originate from instrumental factors, most prominently the instrumental gains.

The visibility calculated above is not precisely the visibility measured by the telescope. The measured visibility combines the actual visibility calculated above with both instrumental gains and noise. In other words,

$$V_{ij}^{\text{meas}} = g_i g_j^* V_{ij} + n_{ij}, \tag{2.7}$$

where $g_{i,}$ and g_j are the complex gains of feeds i and j, V_{ij} is the actual visibility, and n_{ij} is a noise term. We assume this noise term to be uncorrelated, and on the order of $T_{sys}/\sqrt{\tau\Delta\nu}$, where T_{sys} is the system temperature, τ is integration time, and $\Delta\nu$ is bandwidth [19].

We linearize this equation by taking the logarithm and separating the real and imaginary parts so that we have two equations for each measured baseline:

$$\ln|V_{ij}^{meas}| = \ln|g_i| + \ln|g_j| + \ln|V_{ij}| + \mathrm{Re}(n_{ij}) \tag{2.8}$$

$$\arg\left(V_{ij}^{meas}\right) = \arg(g_i) + \arg\left(g_j^*\right) + \arg(V_{ij}) + \mathrm{Im}(n_{ij}). \tag{2.9}$$

This enables us to separately solve for the real and imaginary parts of the gain and the actual visibility [19].

2.2 Amplitude Calibration

Once the set of measured visibilities has been linearized as in Equation 2.8, we can write a matrix equation:

$$\begin{pmatrix} \ln|V_{01}^{meas}| \\ \ln|V_{12}^{meas}| \\ \ln|V_{02}^{meas}| \\ \vdots \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 & 0 & \cdots & 1 & 0 & 0 & \cdots \\ 0 & 1 & 1 & 0 & \cdots & 0 & 1 & 0 & \cdots \\ 1 & 0 & 1 & 0 & \cdots & 0 & 0 & 1 & \cdots \\ \vdots & & & & \vdots & & & \vdots & \end{pmatrix} \begin{pmatrix} \ln|g_0| \\ \ln|g_1| \\ \ln|g_2| \\ \vdots \\ \ln|V_0| \\ \ln|V_1| \end{pmatrix} + \begin{pmatrix} \mathrm{Re}(\eta_{01}) \\ \mathrm{Re}(\eta_{12}) \\ \mathrm{Re}(\eta_{02}) \\ \vdots \end{pmatrix} \tag{2.10}$$

More succinctly,

$$\mathbf{d} = \mathbf{M}\mathbf{x} + \eta, \tag{2.11}$$

where $\mathbf{d}$ is a vector containing the measured visibilities, $\mathbf{M}$ is a matrix containing the various combinations of gain and true visibility, and $\mathbf{x}$ is a vector consisting of the instrumental gains for all feeds and the actual visibilities as defined in equation 2.5

We can solve this matrix equation using least squares methods, so that

$$\mathbf{x} = \left(\mathbf{M}^T\mathbf{N}^{-1}\mathbf{M}\right)^{-1}\mathbf{M}^T\mathbf{N}^{-1}\mathbf{d} \tag{2.12}$$

where $\mathbf{N} = \langle \eta\eta^T \rangle$ is the noise covariance matrix.

This noise covariance matrix $\mathbf{N}$ can be set in a number of ways. In the simplest approximation, we use the identity matrix as the noise covariance matrix. We are therefore assuming that noise is uncorrelated (by selecting a diagonal matrix as $\mathbf{N}$). This can be a valid way to think about the noise covariance, but is also an over-simplification of the scenario and therefore can be a detriment.

As noted above, this procedure yields a matrix, $\mathbf{x}$ composed of the gain amplitudes for each feed and the amplitude of the actual visibility for each feed. However, the amplitude calibration is affected by at minimum one degeneracy. This is apparent from calculating the null space of the $\mathbf{M}$ matrix, which is non-zero and arises because redundant baseline calibration is in itself a relative not absolute calibration technique. We therefore apply a gauge fixing condition of some form. The simplest is to impose a value for one gain in the set and set all gains relative to that value. The primary disadvantage to this method is simple. We do not necessarily know a priori the value of any single gain. The concern is less that we might choose a wrong numerical value for this gain (as we will only consider other gains relative to it in our analysis) but that we may accidentally choose an ill-behaved gain, one with very unusual structure. In the case that the reference gain experiences dramatic fluctuations, this will induce dramatic fluctuations in our other gains as well.

We attempt to mitigate this risk by setting the gain relative to an expected average point for the gains, rather than an individual gain. Specifically, our usual condition is

$$\sum_i \ln |g_i| = 0. \tag{2.13}$$

This condition states that the sum of the natural logarithms of the gains must be equal to zero. This condition then places the gain amplitude values around one, which is a reasonable choice for instrumental gains on the CHIME Pathfinder.

2.2.1 Complications to the Amplitude Calibration

Several complications exist in the amplitude calibration, both from the perspective of non-idealities in a real telescope like the CHIME Pathfinder and from the calibration method itself.

Probably the most significant is the addition of non-identical primary beams between feeds. In summary, if the beam amplitudes A_i for each feed are not identical, one of underlying assumptions of redundancy collapses and the redundant baseline calibration method performs in unexpected, though sometimes useful, ways. This issue will be discussed in greater detail in Chapter 4, where we will present a detailed simulation of such a telescope.

At the simplest level, we can improve our noise covariance matrix by using a diagonal matrix with more carefully considered An individually generated approximation of the noise covariance is generally more exact, if available. One simple option is to use the radiometer noise test to generate a value for σ^2. Another option is to derive the noise covariance matrix directly from the variance of the data itself.

One major assumption made in simple implementations of redundant baseline calibration is that noise is uncorrelated. However, this is quite often not the case, meaning that our diagonal matrix estimate for **N** is inaccurate. A common reason for correlated noise in systems like CHIME is cross-talk. Cross-talk arises when two feeds interact with one another and therefore skew the results from their mutual baseline. This is primarily a concern when dealing with short, intra-cylinder baselines. Therefore, the simplest way to mitigate cross-talk is to ignore either all intra-cylinder baselines or intra-cylinder baselines shorter than a pre-determined minimum baseline distance.

This decision is however not without consequences. Working solely across cylinders in particular introduces an additional degeneracy into the problem. This is evidenced by an additional dimension in the null space of the coefficient matrix, **M**, indicating that the problem is no longer uniquely solved. In the absence of information within a cylinder, the two cylinder's absolute levels can change independently. The initial degeneracy fixing condition present in the algorithm pertains to the overall level of the system's gain. Without knowledge within a single cylinder, the algorithm is free to distribute gain between the two cylinders in any way it pleases, e.g. assigning very high gain to one and a very low gain to the other, so long as the overall level remains consistent with the degeneracy fixing condition. This can be simply and effectively remedied by splitting the degeneracy fixing condition into two complementary conditions, one which sets a level for the first cylinder and one that sets a level for the second cylinder.

2.3 Phase Calibration

Our earlier linearization of the visibility equation split the amplitude and phase of the visibility, allowing them to be solved for separately. Though this thesis' focus is on amplitude calibration using redundant baselines, redundant baseline calibration can also be used as a relative phase calibrator, via a method outlined in Michael Sitwell's PhD thesis [33]. At first glance, it would appear that we could apply the same process to solve for phase, but degeneracies in the problem make this straightforward approach intractable.

To implement redundant baseline phase calibration, we first construct a matrix

$\mathbf{G} = |g\rangle\langle g|$, using an initial estimate $\hat{G_i}j = V_{ij}^{meas}/V_{ij}^{true}$ for each component. This estimate for V_{ij}^{true} can be arbitrary, so one simple option is to use the V_{ij}^{true} calculated from the amplitude calibration. With an estimate for $\mathbf{G}$, we attempt to solve for the gain vector that minimizes chi-squared, which we can write as [33]

$$\chi^2 = \sum_{ijkl} \left(\hat{G}_{ij} - G_{ij}\right) C_{ij,kl}^{-1} \left(\hat{G}_{kl} - G_{kl}\right). \tag{2.14}$$

If the covariance matrix $C_{ij,kl}^{-1}$ is uncorrelated between baselines and gives each baseline the same variance σ^2, we can reduce Equation 2.14 to

$$\chi^2 = \frac{1}{\sigma^2} \sum_{ij} |\hat{G}_{ij} - G_{ij}|^2. \tag{2.15}$$

As a consequence of our definition of $\mathbf{G}$, $|g\rangle$ is an eigenvector of $\mathbf{G}$ with the eigenvalue $\langle g|g\rangle$. We can therefore use the eigenvalue decomposition to find values of the gains. As $\hat{\mathbf{G}}$ is Hermitian and positive semi-definite, its eigenvalue decomposition is the same as its singular value decomposition. This is convenient because it allows us to find which eigenvector we should use as our gain estimate. Specifically, this SVD equivalent means that reducing $\hat{\mathbf{G}}$ to a rank-1 matrix with the SVD is the equivalent to minimizing χ^2. We are able to conclude that the eigenvector corresponding to the largest eigenvalue is our desired gain vector [33].

With an estimate of the gains, we can re-estimate the true visibility for a given baseline $b = i - j$, writing

$$V_b = \frac{\sum_i g_i g_{i+b}^* V_{i,i+b}^{\text{meas}}}{\sum_i |g_i|^2 |g_{i+b}|^2}. \tag{2.16}$$

We then iterate this process until our results converge [33].

There are a few caveats to this method. First, the assumption that covariance matrix $C_{ij,kl}$ is proportional to the identity used to reach Equation 2.15 requires us to assume autocorrelations should receive the same weight as all other correlations, which is distinctly inaccurate. However, if we only need to solve for gain phases, we can replace $\hat{G}_{ij} \to \hat{G}_{ij}/|\hat{G}_{ij}|$ which makes $\hat{\mathbf{G}}$ a matrix with complex elements with unit norms and thus the diagonal elements are all scaled to one. Second, there are a number of degeneracies in the phase solution. It is insensitive to both a rotation of the sky (or equivalently a tilting of the array) as well as the absolute phase level of the system. These degeneracies can be fixed either by calculating the deviation from expected phase at a point source or by simply fixing the first two gains in the gain vector to known values.

Though this method has been previously shown to be successful in small cases (see [33]), in the remainder of this thesis we will focus solely on amplitude calibration.

2.4 Redundant Baseline Calibration on Simulated Data

Although we have at our disposal actual CHIME pathfinder data, we choose to first present redundant baseline calibration results from our carefully constructed simulations, as this allows us to make decisions about most effective way to implement redundant baseline calibration without having our results clouded by the irregularities and complexities of real experimental data.

2.4.1 Simulating the CHIME Pathfinder

As we are interested only in testing a calibration method and not in testing 21 cm hydrogen intensity detection, we use the software package COsmology in the RAdio Band (CORA) to generate a foreground map including the galaxy and point sources. Our map includes the actual values of point sources brighter than 4 Jy as well as a synthetic population of dimmer point sources (brighter than 0.1 Jy at 151 MHz) and a Gaussian realization representing even dimmer unresolved sources. The galactic component of the map is extrapolated from the Haslam map at 408 MHz with additional random fluctuations [16]. The map, with the colour bar scaled to make the galaxy visible, is shown in Figure 2.1.

After simulating a test sky, we generate test beam transfer matrices. While we use alternative beam models later, for this analysis we confine ourselves to the fiducial quasi-Gaussian beam introduced in [32]. This beam model varies slightly for the two polarizations present in CHIME, here designated x polarization (the feed dipole pointing east-west) and y polarization (the dipole pointing north-south). The fiducial model is not a direct attempt to solve for the beam pattern but rather

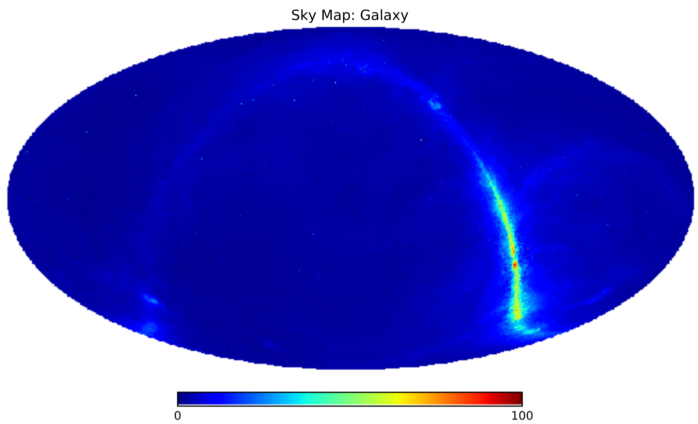

Figure 2.1: The simulated sky map used for simulations in Section 2.4. It contains both point sources and galaxy, but the point sources appear as isolated pixels and are therefore almost invisible at this size. The colour bar has been scaled to show the galaxy, as it is by default saturated by the brightest point sources. The map is created using a combination of the Haslam 408 MHz map, known bright point sources in the CHIME frequency band, and Gaussian realizations of dimmer point sources.

an approximation composed of the product of a function describing response in the EW direction and a function describing response in the N-S direction. We write the beam amplitude for an unfocused dipole, our model for the feed, as

$$A_D(\theta;\theta_W) = \exp\left(-\frac{\ln 2}{2}\frac{\tan\theta^2}{\tan{\theta_W}^2}\right), \tag{2.17}$$

with θ_W the beam's full width at half-power. Given the height of our feed relative to its conducting ground plane (i.e. the cylinder surface), we can calculate θ as $2\pi/3$ in the H-plane and $0.675\theta_h$ in the E-plane. We treat our problem then as a Fraunhofer diffraction problem with our given amplitude. Therefore, our cylinder has the amplitude

$$A_F(\theta,\theta_W,W) \propto \int_{-W/2}^{W/2} A_D\left(2\arctan\frac{2\pi}{W};\theta_W\right)e^{-ikx\sin\theta}dx \propto \int_{-1}^{1} e^{-\frac{\ln 2}{\tan{\theta_W}^2}\frac{u^2}{1-u^2}i\frac{\pi W}{\lambda}u\sin\theta}du \tag{2.18}$$

incorporating the fact that for a cylinder with f-ratio 1/4, a ray hitting a distance of x away from the cylinder's centre reflects at an angle $\theta = 2\arctan(2x/W) = 2\arctan u$ [32].

We can create our overall beam for each polarization using these two A functions and an addition function p_a, which is a unit vector in the polarization direction for a dipole in a given direction. Therefore, the x feed has a beam amplitude

$$A_a^X(\hat{\mathbf{n}}) = A_F(\arcsin(\hat{\mathbf{n}} \cdot \hat{\mathbf{x}}); \theta_E, W) \times A_D(\arcsin(\hat{\mathbf{n}} \cdot \hat{\mathbf{y}}); \theta_H)\, p_a(\hat{n}; \hat{x}) \tag{2.19}$$

and the y feed has a beam amplitude

$$A_a^Y(\hat{\mathbf{n}}) = A_F(\arcsin(\hat{\mathbf{n}} \cdot \hat{\mathbf{x}}); \theta_H, W) \times A_D(\arcsin(\hat{\mathbf{n}} \cdot \hat{\mathbf{y}}); \theta_E)\, p_a(\hat{n}; \hat{y}), \tag{2.20}$$

with $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ unit vectors pointing in the East (transverse) and North (parallel) directions respectively [32].

It is critical to note that the fiducial beam model is a significant simplification of the CHIME beam. It does not account for cylinder surface considerations or for any complicated beam structure or non-identicalities between individual feeds. This model is also not informed by knowledge of the instrument's specific characteristics.

The simulation pipeline creates a model telescope using the layout of the CHIME pathfinder at a provided time and the fiducial beam described above. It generates beam-transfer matrices then applies the input sky map to these beam transfer matrices to create a pseudo realistic realization of what the telescope might see.

The simulation pipeline then separately adds noise and complex gain to the simulated data. First, we add a constant amount of noise, corresponding to the receiver temperature, modeled as 50 K. Then, we generate complex gains for each input of the the simulated CHIME pathfinder. These gains have an average value of one, but include long-timescale random fluctuations. We do not expect CHIME's complex gains to change particularly rapidly (e.g. on few minute to hour timescales) but we do expect variations over several hours or a day. The gain fluctuations in our initial implementation are likely also larger amplitude than realistic gain fluctuations. Thus, we could describe this as a pessimistic gain scenario. The pipeline also calculates the appropriate combinations of feeds and applies these gains. Finally, we add a small amount of Gaussian random noise to the samples.

At this stage, our simulated data should look and behave exactly like real CHIME Pathfinder data, assuming that our initial model of the CHIME Pathfinder is correct. (We must particularly look out for inconsistencies between our assumed quasi-Gaussian beam model and the real telescope's more complicated beam model.)

2.4.2 First implementation

We first and most simply implement redundant baseline calibration assuming an identity noise covariance matrix. We do ensure that we remove feeds which are masked out in the simulation, as when they are left in the data set, the algorithm struggles to find a solution for those inputs which matches our imposed degeneracy fixing condition and therefore degrades the entire solution set. We also include only one polarization at a time and therefore run our redundant baseline analysis twice to extract all relevant underlying visibilities and gains. We do not attempt to recover the cross-polarization correlations, as they are not meaningful in this context.

Figure 2.2 shows the complete set of gains derived from redundant baseline calibration on our simulated CHIME pathfinder data over one six hour time stream file. While the plot is in many ways an overwhelming amount of information, there are a few results we can deduce from looking at this compilation.

First, notice the set of lines at exactly 1.0. These lines are the result of forcing the gain for masked out feeds and feeds from the second polarization to be 0. Redundant baseline calibration actually calculates the natural log of the gains, and we therefore exponentiate the entire data set before presenting it here. Therefore, our zeroed out gains become exactly 1.0. We can thus ignore those lines entirely.

Second, notice the improved precision at about 2000 seconds into the time stream, which is a simulated TauA transit. Redundant baseline calibration is in principle sky independent. However, in practice, bright sources dramatically improve the signal-to-noise ratio in CHIME Pathfinder data. When this happens, the redundant baseline algorithm is able to provide a noticeably more precise solution, as it is much less strongly influenced by sources of noise in the data. For example, when looking at an effectively blank sky, the 50 K receiver temperature is signif-

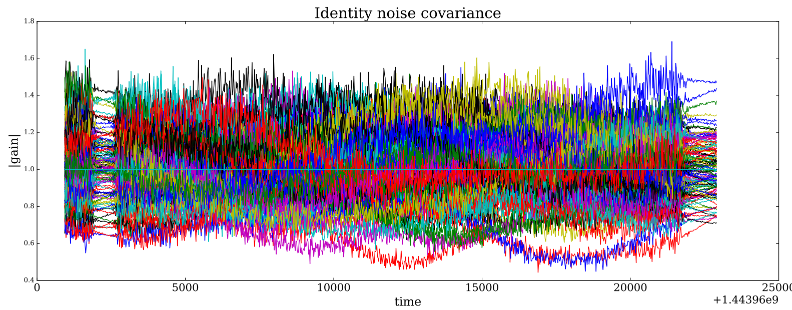

Figure 2.2: We see here the complete gain results for redundant baseline calibration implemented on a simulated CHIME Pathfinder data set, assuming identity noise covariance. Notice the set of lines at exactly 1.0, showing feeds which are masked out and therefore forcibly set to 0 (then exponentiated to become 1). Notice also the improved precision of results at around 2000 seconds after the beginning of the file, a result of improved signal-to-noise ratio during point source transit. The long-term wavy structure in the gains closely trace the input gain fluctuation.

icant. However, it is intuitively obvious that it is not nearly as significant when looking at a bright point source with a brightness temperature much greater than 50 K.

This is an interesting result, as it indicates that the assumption that redundant baseline is sky independent is true only in a limited sense. While redundant baseline calibration is not sensitive to the exact sky model, it is sensitive to signal to noise ratio variations. This is potentially valuable, as it may allow improved algorithms in the future which incorporate some sky info and therefore improve performance.

Third, notice the long-term wavy structure in the gains. Initially, this may seem concerning, as it seems that the solutions are floating around. However, when creating the simulation, we incorporated long-term gain fluctuations. Figure 2.2 is not the ideal figure to demonstrate that these variations are as expected, but the structure of the fluctuations in Figure 2.2 should be considered an encouraging sign. We will investigate deviation from the input fluctuations more exactly later in our analysis.

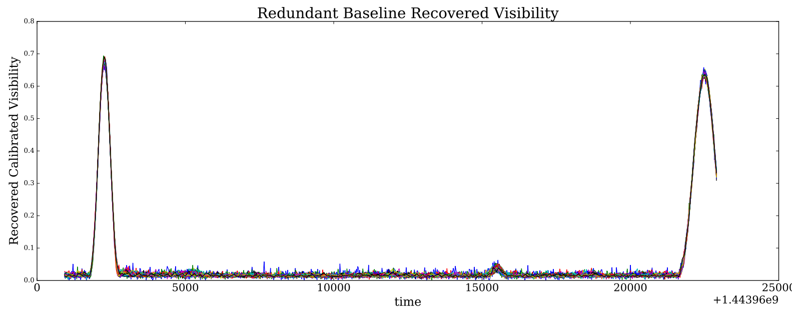

Figure 2.3: We see here the underlying visibilities derived from the redundant baseline calibration algorithm. We notice that the algorithm has correctly located the point source transits and some underlying features of the galaxy. The point source transits are the visible parabolic structures located at the beginning and end of the file, which corresponds to about 6 hours of time series data. We notice very little structure at times when the input simulated visibilities are predominantly noise, which is encouraging.

We can also generate a plot of the redundant baseline calibrated underlying visibilities. These visibilities should be those generated by Equation 2.5 whereas the visibilities input to the algorithm were of the form of Equation 2.7. Calibrated visibilities are, of course, desirable as our motivation for knowing the complex gains is ultimately to recover calibrated visibilities of the form of Equation 2.5. We notice that our recovered visibilities correctly locate our expected point source transits and a few underlying features of the galaxy. As those are the only two inputs to our map (see Figure 2.1), we consider that success. (Or at least conclude that any failures in the redundant baseline calibration will be visible only in the gain solutions.)

2.4.3 Determining Ideal Noise Covariance

While at a base level, we are pleased with the performance of the identity noise covariance version of the algorithm, we suspect that the performance could be improved by specifying the noise covariance more carefully. Therefore, we also tested variations of the noise covariance matrix. First, we calculate the noise covariance directly from the input data. Second, we use the radiometer noise test to determine the noise covariance. Note that while both of these methods attempt to represent the noise more accurately than the identity, both are constructed into diagonal matrices (just like the identity), representing a continued assumption that there is no correlated noise. This should be an accurate portrayal of our simulated data, where there genuinely is no correlated noise, but the assumptions may break down for real CHIME Pathfinder data. Summary results from these methods, equivalent to those presented in Figure 2.2 are presented in Figures 2.4 and 2.5.

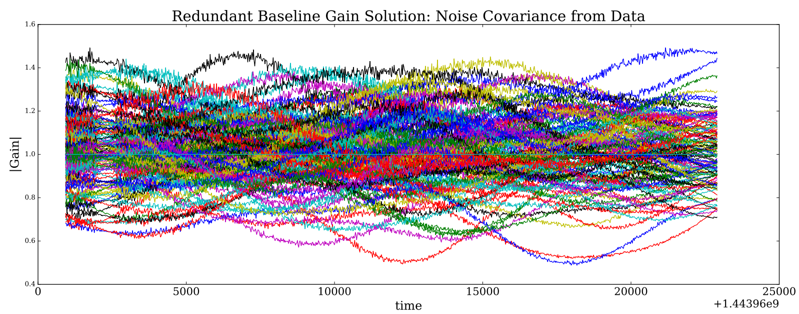

Figure 2.4: We show here the redundant baseline-derived gains, using a noise covariance calculated directly from the data. Compare to Figure 2.2 and Figure 2.5. We again see the congregation of turned off or opposite polarization inputs at 1, improved results at higher signal to noise, and a general correspondence with the shape of the gain fluctuations. However, we also see decreased noise in the solution relative to the identity noise covariance results.

All three variations generate reasonable results and thus we are not able to tell which is superior purely by looking at these summary plots. We therefore move forward with specific comparison plots. We take advantage of the fact that we are using simulated data to directly compare our redundant baseline outputs to our simulated inputs. We compare both in absolute difference and in percent difference. We find that on average the identity noise covariance gain results are within 15% of the input gains, while the radiometer and data derived noise covariances are within approximately 5% of the input gains. Figure 2.6 shows a sampling of percent dif-

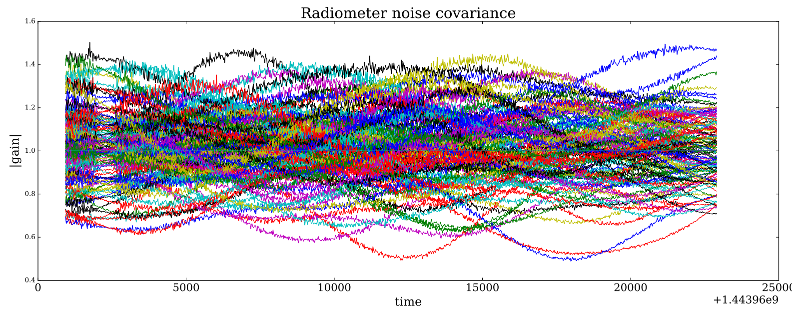

Figure 2.5: We show here the redundant baseline-derived gains, using a noise covariance calculated directly from the radiometer equation. Compare to Figure 2.2 and Figure 2.4. We again see the congregation of turned off or opposite polarization inputs at 1, improved results at higher signal to noise, and a general correspondence with the shape of the gain fluctuations. However, we also see decreased noise in the solution relative to the identity noise covariance results. This decreased noise is very similar to the covariance from data results in Figure 2.4.

ference results. We also look at absolute deviations, particularly for the radiometer and data-derived variations, see Figure 2.7 for a sampling. Finally, we find the average values and standard deviation of the gain differences for radiometer and data-derived results.

One important thing we notice is that the differences are consistently negative. We calculate our deviations as input minus recovered, and therefore deduce that our redundant baseline gain solution is consistently higher than our input gain set(for both data-derived and radiometer noise covariances). We suggest this is likely an effect of bias in redundant baseline calibration or an indication that our degeneracy fixing condition is failing, but also suggest that this deserves further examination in implementing redundant baseline calibration. One way to correct for this might be to relax the assumption that the noise is independent and thus to modify our noise covariance matrix to not be diagonal.

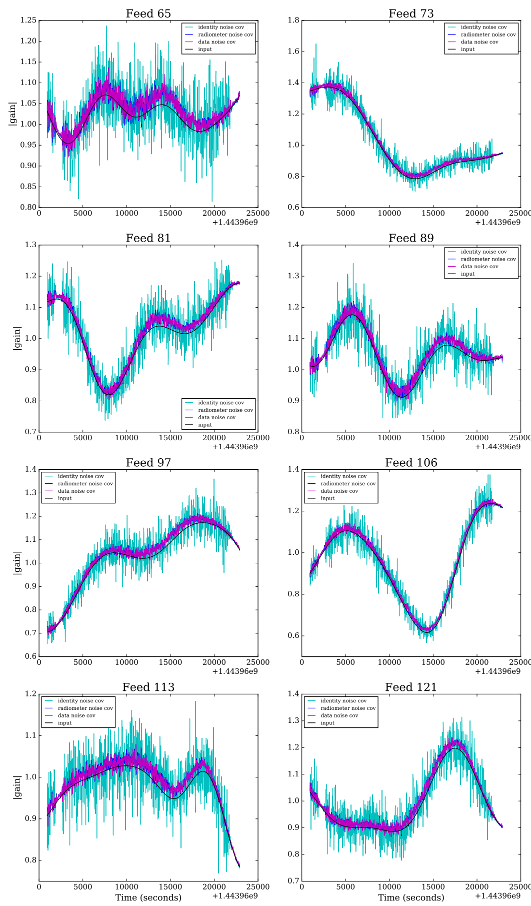

Figure 2.6: In this figure, we compare the input gains to the results for redundant baseline amplitude calibration for the identity noise covariance (cyan), data-derived noise covariance (magenta), and radiometer noise covariance (blue) for a sampling of inputs on the west cylinder. We see that the results trace the general structure of the gain at all times, that they are generally improved at times corresponding to point source transits, and that results for the data and radiometer noise covariances have less noisy solutions. Similar figures for other portions of the simulated array are shown in the Appendix.

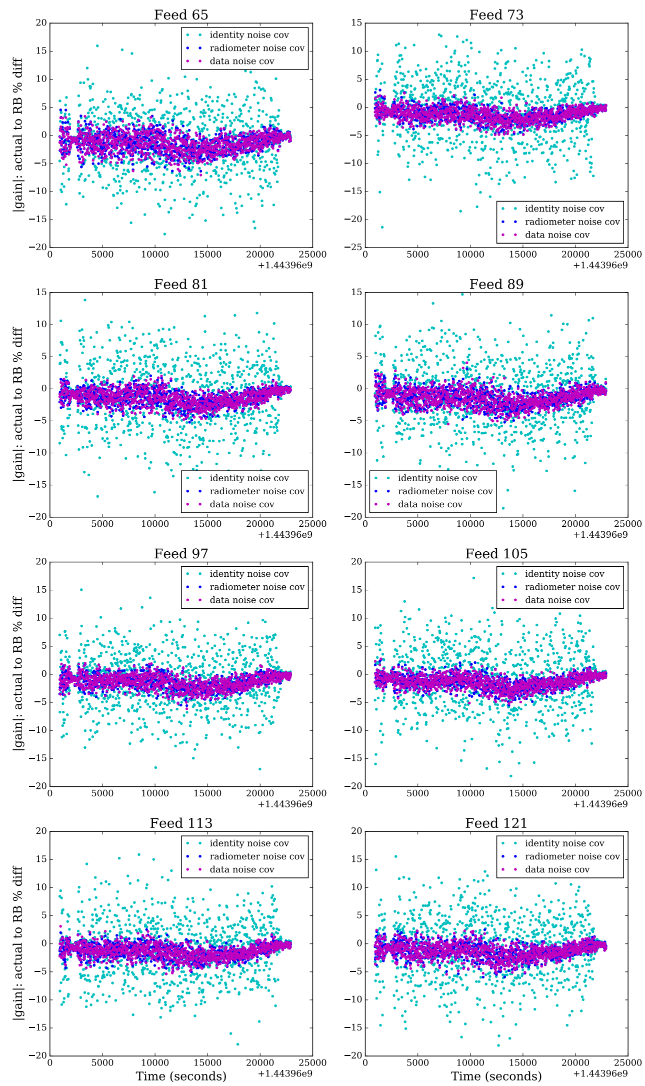

Figure 2.7: In this figure, we compare the percent deviation from input gains for the identity noise covariance (cyan), data-derived noise covariance (blue), and radiometer noise covariance(magenta) for a sampling of inputs on the west cylinder. The improved precision at higher signal to noise regions is less obvious here, but the smaller scatter in radiometer and data derived noise covariances is present clear.

2.4.4 Conclusions from initial implementation

We find that our initial implementation is largely successful, as we are able to recreate to within 5% the input gain amplitudes. However, we recognize there are some concerns, such as the apparent bias of our results. Our first implementation shows some encouraging results, but also leaves room for improvement. One obvious spot for improvement is to incorporate a careful implementation of the redundant baseline phase calibration, which is currently only implemented for small test cases.

While we are encouraged by our success in recovering a simulated model, we also realize that this chapter's results are limited to cases that truly qualify as redundant baselines, which may not be realistic cases. Looking forward, in Chapters 3 and 4, we will apply redundant baseline calibration to a perturbed simulation and to real CHIME Pathfinder data.

Chapter 3

Using Redundant Baseline Calibration in CHIME Pathfinder Data

Our analysis in Chapter 2 suggests redundant baseline calibration is a model worth pursuing for a telescope such as CHIME. As our eventual goal is to have calibrated CHIME data for use in intensity mapping, it is important to demonstrate that calibration methods can be used with real data as well as with simulations. Therefore, we apply redundant baseline calibration to real CHIME pathfinder data.

However, our redundant baseline analysis in Section 3.1 indicates significant deviation from the expected gain results. This could in principle be due to unexpected but genuine structure in the telescope's complex gains, but other calibration methods applied to the CHIME Pathfinder do not support such a conclusion. Therefore, we argue that the deviation from expected gain results is attributable to deviations from redundancy in the CHIME Pathfinder. We then examine the extent to which nominally redundant baselines are actually redundant in a point-source calibrated CHIME Pathfinder data set, and find that deviation from redundancy is significant.

3.1 Redundant Baseline Calibration on CHIME Pathfinder Data

3.1.1 Modifying redundant baseline calibration for use on real data

Cross-Talk

Before beginning our analysis, we make a few modifications to our redundant baseline approach to account for known sources of error in the CHIME Pathfinder data. Most importantly, we know that the individual antenna in the CHIME Pathfinder suffer from cross-talk. This renders intra-cylinder baselines suspect up to some minimum separation. This effect likely on applies for very short intra-cylinder baselines (e.g. shortest or second shortest baselines), but we do not have a definite metric to establish what the minimum length necessary to avoid cross-talk is in the CHIME Pathfinder. Therefore, for this analysis we exclude intra-cylinder baselines.

This straightforward method for eliminating cross-talk, however, has implications for the rest of the analysis. In particular, removing intra-cylinder information introduces an additional degeneracy into the solution. Our usual degeneracy fixing condition sets $\sum \ln |g_i| = 0$, but in the absence of intra-cylinder baseline information, this condition fails to ensure that the algorithm does not actually set $\sum \ln |g_{i,\text{cylinder 1}}| = -1$ and $\sum \ln |g_{i,\text{cylinder 2}}| = 1$.

Therefore, we simply turn our overall degeneracy fixing conditions into two identical conditions:

$$\sum_{i,\text{ cylinder}} \ln |g_{i,\text{ cylinder}}| = 0, \tag{3.1}$$

one for each cylinder.

Using Transits

In the simulated data, we calculated redundant baseline solutions over an entire six hour period of data, including both point transits and quiet portions of the sky. However, when using real CHIME Pathfinder data, we choose to confine ourselves to point source transits. We saw in Chapter 2 that at higher signal to noise periods in the simulation, our redundant baseline solution was significantly less noisy and clearer than during quiet sections of the sky. CHIME Pathfinder data will necessarily be more complicated than simulated data, with less well-known system temperature and more complicated noise structure, so we choose to focus only on the clearest sections of the more complicated system.

We are also more confident as to the expected structure of the complex gains during the short period of time encompassed by a point source transit than we are over a longer period of time. While we expect time dependent gain fluctuations in CHIME Pathfinder data, we expect fluctuations during the tens of minutes encompassed in a point source transit to be relatively small.

Removing "Dead" Feeds

In real data, we are also confronted with the possibility of input channels which are absent or being used for non-CHIME antenna electronics. Including such anomalous channels will wreak havoc on the redundant baseline analysis, as all correlations including that non-CHIME Pathfinder channel will no longer be redundant with other correlations of the same baseline distance. For example, the RFI antenna will clearly not correlate with CHIME antennas at all like a CHIME antenna would.

Additionally, due to changes in experiment configuration or electronics malfunction, the CHIME Pathfinder experiences intermittent drop-outs: feeds which in principle should be CHIME antennas but are in fact either turned off or malfunctioning. Correlations involving such channels are also non-redundant.

If the number of inactive feeds is relatively small, the redundant baseline algorithm generally recognizes that these feeds are different and attempts to set their gains to 0 - effectively removing them from the final set. However, for best accuracy, we should remove these feeds ourselves before running the redundant baseline calibration algorithm.

In a small set of feeds, this can be done by manually checking each feed's autocorrelation and simply excluding them from the data set requested from the broader CHIME Pathfinder data repository. For a larger set of feeds (e.g. the CHIME Pathfinder), this method becomes laborious to follow. Therefore, we must develop a method to ensure we know which feeds should be excluded from our analysis and to exclude all correlations including such feeds.

The first step is fairly straightforward - we establish by loading and plotting a handful of working inputs' autocorrelations a threshold value (in correlator units) for a valid feed's autocorrelation then confirm that each autocorrelation in our data set remains above that threshold. (While we exclude autocorrelations from our redundant baseline analysis due to their noise properties, we can use them here as we are simply ascertaining whether the feed is functioning properly as a CHIME antenna.) We create a list of feeds which are not above this threshold and designate them "dead feeds." "Dead feeds" is a mild misnomer - they may be mechanically impaired, absent, or simply never present (i.e. the channel is designated for an alternative input such as a noise source, RFI antenna, or the John A. Galt 26 m telescope).

The second is slightly more complicated. At the most basic level, we can replace the columns in our degeneracy fixing corresponding to these feeds with zeros (instead of ones). This excludes those feeds from the degeneracy fixing condition, but may not fully solve our problem as the "dead feeds" are still in the data set. Thus, we also create a slightly more complicated, but more accurate version. We create a full coefficient matrix, with a degeneracy fixing condition that excludes dead feeds. We then create a vector of the length of the number of correlations which is one for all correlations we want to include and zero for all correlations involving "dead feeds." We multiply our coefficient matrix by this vector, zeroing out all rows corresponding to correlations involving "dead feeds." Thus, we force the algorithm to exclude those correlations from its calculation and our final solution excludes "dead feeds."

With these caveats in mind, namely our exclusion of intra-cylinder baselines with modified degeneracy fixing condition; our use of transits; and our exclusion of "dead feeds," we are ready to apply redundant baseline calibration to CHIME Pathfinder data.

3.1.2 Results

Besides the modifications above, the redundant baseline algorithm remains unchanged. Below we see the results for a Cygnus A transit, first the recovered underlying visibilities and second the recovered redundant baseline gains.

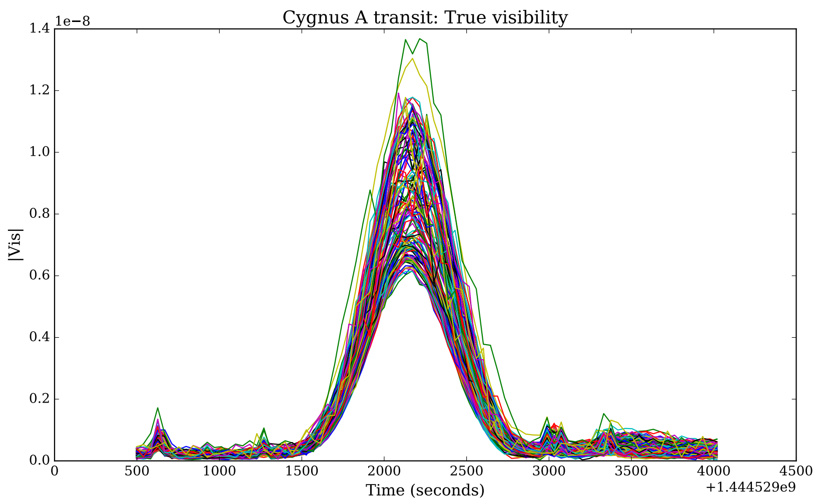

Figure 3.1: Redundant baseline true visibilities for a June 2015 CygA transit. In this analysis, we use only inter-cylinder baselines and exclude "dead feeds" using our more complete algorithm.The two polarizations are solved for separately. The recovered true visibilities certainly recover the existence of a point source transit and a reasonable shape estimate for it. We do observe some spikes through the solution - these may be attributable to the varying level of redundancy for short vs. long baselines. The units on the y-axis are correlator units, as redundant baseline is a purely relative calibrator. We could normalize this plot to know values for CygA should we prefer a plot in Jansky.

The true visibility plot resembles our expected result closely enough that we are reassured that our algorithm is likely functioning. However, the extent of structure in the plot of redundant baseline gains is cause for concern. It indicates something unexpected is occurring in our solution, something that justifies further analysis.

It is immediately apparent that the gains have significant slope. We attempt to quantify this amount of slope in Figure 3.3. We anticipate that during the relatively

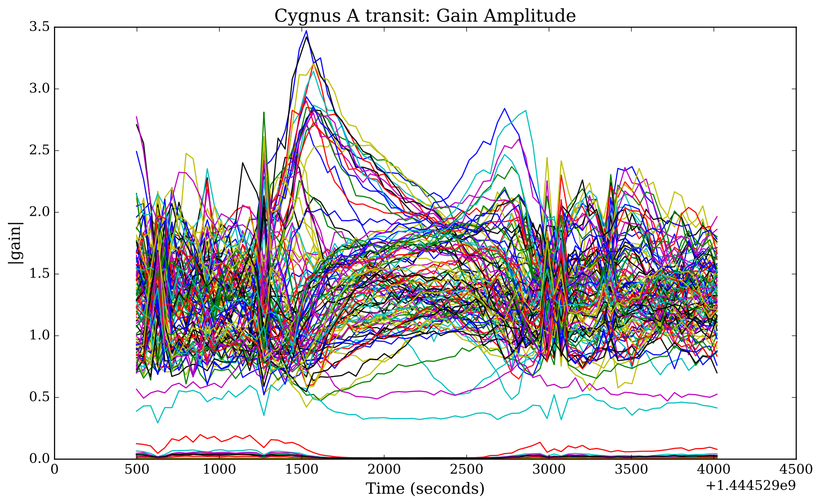

Figure 3.2: Redundant baseline gains calculated for a June 2015 CygA transit. In this analysis, we use only east-west baselines, and exclude "dead feeds." As described previously, we solve for East and South polarizations separately and recombine the results after our calculations. Notice the definite slope in the gain values during the point source transit. This is contrary to our expectation that gains would not vary much on a short, point source transit time scale.

brief timescale of the point source transit, the complex gains will change relatively little. Instead, we see dramatic changes, from e.g. 1.5 to 0.5. This could be do to purely mistakes in our algorithm, but this is unlikely because the true visibility graph was reasonable and the algorithm is almost totally identical to that used in the successful simulation analysis. It seems most likely that this represents a deviation from our assumptions of redundancy. Recall from Chapter 2 which showed that the visibility was determined by the beam functions A_i and A_j, the sky, and a phase factor dependent on baseline distance. The redundant baseline algorithm assumes this description is precisely correct and that $A_i = A_j$. If this is untrue, then we would see failures in the calibration.

Independent analysis within the CHIME collaboration suggests that deviations in the feed locations are relatively small. This diminishes the probability that the baseline distance is causing deviations in our calibration. Additionally, we have not assumed a sky value but know that it would be consistent for all feeds and therefore discount the sky as a possible source of deviation.

This indicates the beam function is likely the source of the problem. If the two beam functions A_i and A_j in any given visibility V_{ij} are not identical, then the redundant baseline algorithm would not operate as intended and could generate slopes in the gains such as those observed in Figure 3.2. This could also account for the spikes and deviations in Figure 3.1. If this is the case, we would expect relatively small deviations from day to day for a given point source.

While beam values might be very different at different points on the sky, they should not be significantly time dependent. Therefore, we plot a sampling of redundant baseline gain amplitudes over two consecutive CygA transits in Figure 3.4. Notice that for a given feed, the two curves have similar structure, though for different feeds, the curves show a variety of behaviours. This result is also consistent with our hypothesis regarding beams, which would not significantly vary from day to day. The absolute values vary, but this is not a cause for concern as the redundant baseline solution is only relative, so the absolute levels on two separate sets of data may vary freely.

3.2 Examining Redundancy in the CHIME Pathfinder

Evidence from the section 3.1.2 suggests variation from redundancy in the CHIME Pathfinder, likely due to feed-to-feed beam variations. If that conclusion is accurate, deviations from redundancy should be visible in other analyses. Directly comparing measured visibilities to check for redundancy is not a useful technique, as those measured visibilities contain complex gain information and are therefore obviously non-redundant. We therefore require a set of calibrated data, with complex gains removed.

Thanks to the efforts of the CHIME Pipeline Processing Team, such a set of data exists. Members of the CHIME collaboration created a set of fully processed, point source calibrated results for CHIME Pathfinder data from the fall of 2015 and the spring of 2016 and selected consistently good quality frequencies to be used as test frequencies. The analysis that follows uses a subset of that data, specifically

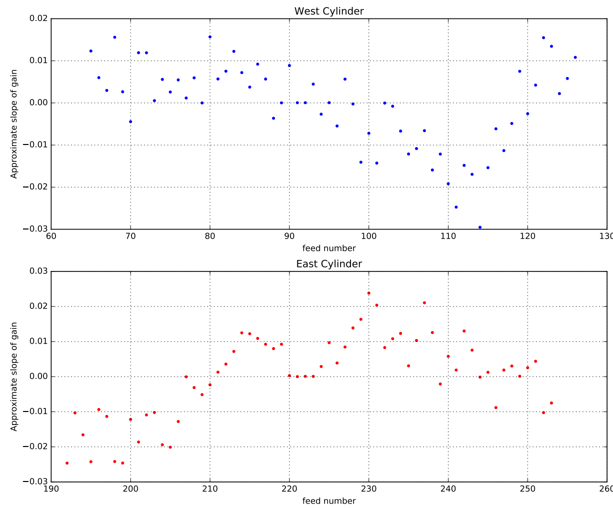

Figure 3.3: In this figure, we calculate a rough measure of the slope of each gain, obtained by taking the rise over run for an individual feed's redundant baseline gain results between samples 60 and 100. The 40 sample range represents approximately one standard deviation around transit, which is the period of time in which we are most confident in our results. This figure indicates that the slope effect observed by eye does appear to be significant.

pass_1p, which occurred from October 9-22, 2015. This data set was calibrated to CygA.

We began our analysis with the simplest possible process: selecting pairs of baselines which should be redundant and plotting them during a point source transit. Results from this process for four of the eight test frequencies and for all instances of a short inter-cylinder baseline during a Cassiopeia A transit are shown in Figure 3.5. Each of these plots shows all instances of the same baseline, which should be redundant. Therefore, we would naively expect the graph to show one single trace, repeatedly overplotted. Perhaps accounting for small deviations or shifts in gain from CygA to CasA, there would be some small range around an average value. However, that is not what is recorded in Figure 3.5. We see instead

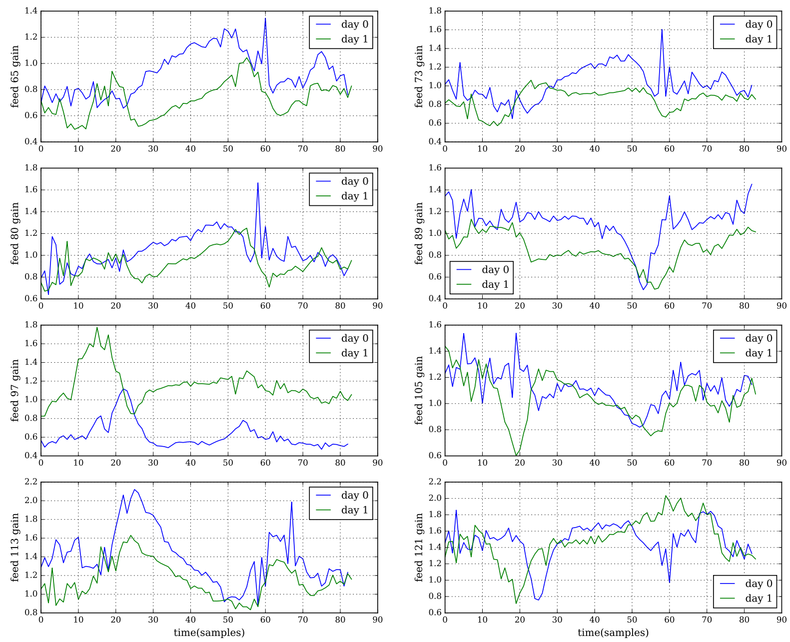

Figure 3.4: We compare redundant baseline gain results for selected feeds on the west cylinder calculated over two consecutive CygA transits. We notice that there is very little deviation between the two days and regard this as evidence that the cause of the slope in the gain solution is not strongly time dependent and is likely a property of the array. An identical figure showing the east cylinder is included in the Appendix.

significant scatter in the values. Discounting very high or very low curves which appear to be outliers, there seems to be about a factor of two variation between instances of the same nominally redundant baselines.

We continued our analysis by viewing a slice of the visibilities at the peak of a CasA transit, as seen in Figure 3.6. This allows us to compare the relative amplitude of the transits in each nominally redundant instance and quickly examine whether there is obvious north-south structure in the value of the nominally redundant baselines. We suggest that there is not a significant pattern along the cylinder,

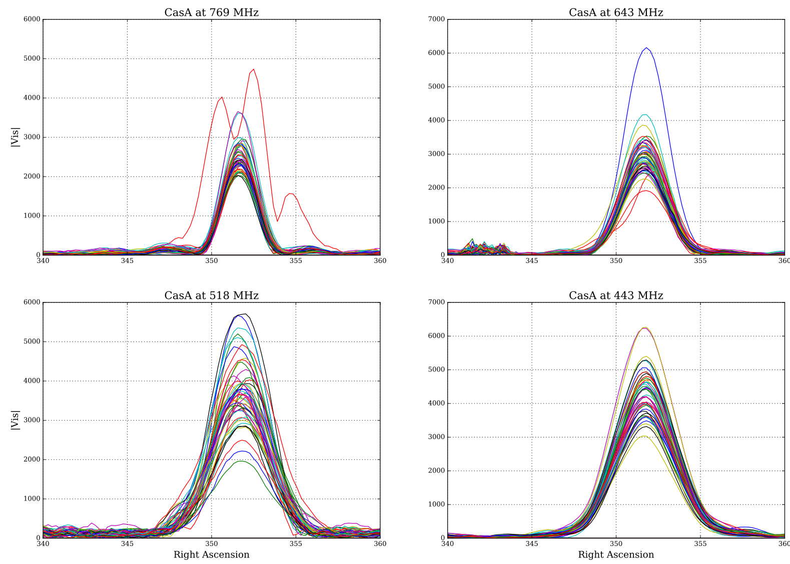

Figure 3.5: We compile redundancy comparisons for a short intercylinder baseline during a CasA transit in pass_1p for four test frequencies. Nominally, these are instances of the same redundant baseline, and we would therefore expect there to be little to no variation between curves in a given frequency. Instead we observe significant deviation. We propose that this deviation is largely derived from variations in the beam pattern between feeds.

though there are regions that are systematically higher or lower.

We also examined the effect of changing to a different point source, to understand the declination dependence of the deviation from redundancy. By examining a Taurus A transit in Figure 3.7, identical to the transit plot from CasA in Figure 3.5, we found that the variation from redundancy is not significantly declination dependent. The exact ordering and amplitude of deviations from redundancy vary somewhat between CasA and TauA, leading us to believe there may be a slight declination dependence in deviation from redundancy. However, the existence of the phenomenon is consistent across different declinations, leading us to believe

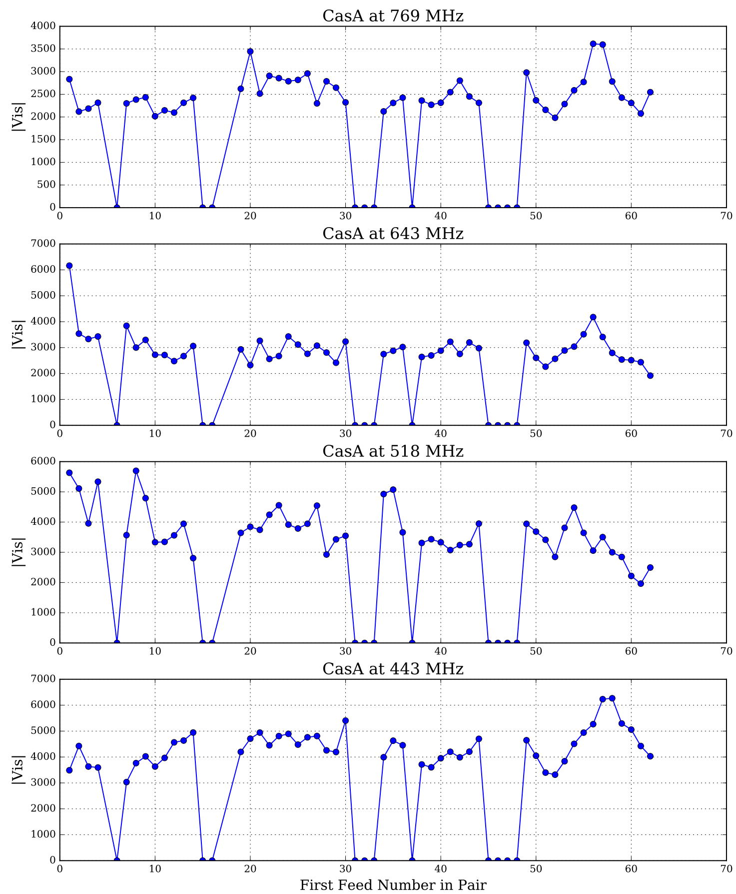

Figure 3.6: We compile slices at CasA transit peak for each nominally redundant instance of a short inter-cylinder baseline at each of four test frequencies. We notice that there is not a defined pattern in visibility value based on feed location, which indicates that the effect causing the deviation from redundancy does not vary in a systematic way along the CHIME Pathfnder cylinder.

.

that is not the primary factor.

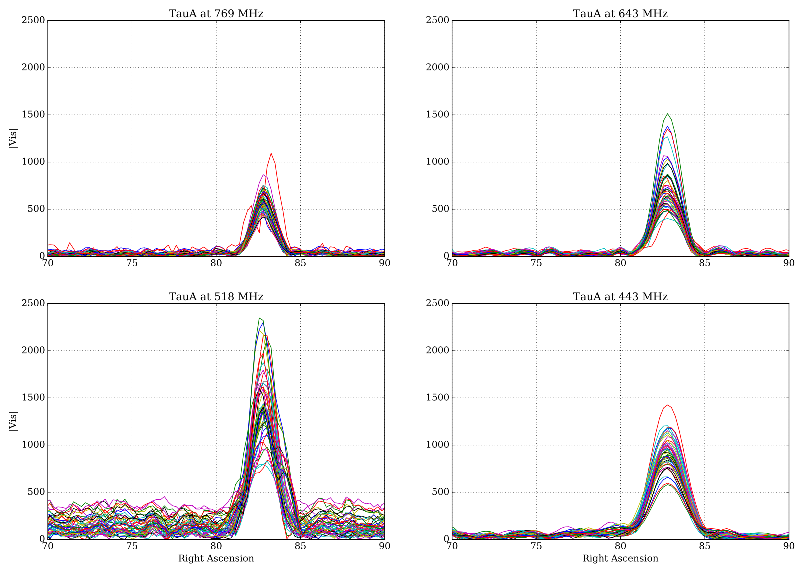

Figure 3.7: We compile redundancy comparisons for a short intercylinder baseline during a TauA transit in pass_1p in the same manner as Figure 3.5. We note that while the exact magnitude of deviations from redundancy may differ, the general structure of the deviation is similar to that present in the CasA data, indicating the existence of such deviations is not declination-dependent although the values may be.

Until this point, the redundancy analysis has focused on sidereal stacks and therefore taken into account each day of data in the stack. One potential cause for the deviation from redundancy could be deviations across days. Perhaps one very deviant day influenced the rest of the stack and created the appearance of deviations from redundancy.

Therefore, we plotted comparisons between two consecutive sidereal days, CHIME Sidereal Day (CSD) 693 and CSD 694. However, each individual day shows a similar pattern to the overall sidereal stack. Comparing the difference between the two days, it is clear that the deviation between the two days is significantly smaller than the deviation within a given baseline on a given day. (It is in fact about 10% of the spread in each individual day).

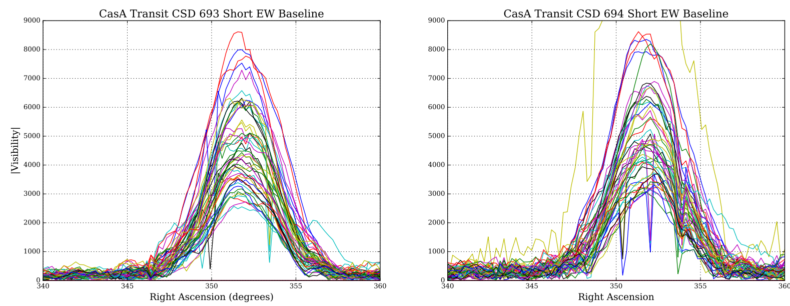

Figure 3.8: In this figure, we observe all instances of a nominally redundant, short inter-cylinder baseline at frequency 518 MHz, compared between CSD 693 and CSD694. Notice that both individual days are significantly non-redundant, lessening the likelihood that the deviations from redundancy present in the sidereal stack for pass1_p are caused by a deviant day included in the pass. It appears further that deviation from redundancy is not strongly time-dependent.

Based on several different pathways of analysis, we conclude that there are significant deviations from redundancy in point source calibrated CHIME Pathfinder visibilities. Based on the consistent existence of these deviations across different declinations and times, we assert that these deviations are probably due to feed-to-feed beam variations. This conclusion is strengthened when combined with evidence from our redundant baseline test, which suggested similar non-time dependent deviations from redundancy.

We hypothesize that this lack of redundancy is arising from feed-to-feed beam variations, as that would explain both the presence of the deviations across time and declination as well as the deviations in redundant baseline results. Armed with that knowledge, we postulate that the redundant baseline algorithm may be able to probe these feed-to-feed beam variations for periods where the complex gain is relatively constant (e.g. a point source transit). We will investigate these possibilities further in Chapters 4 and 5.

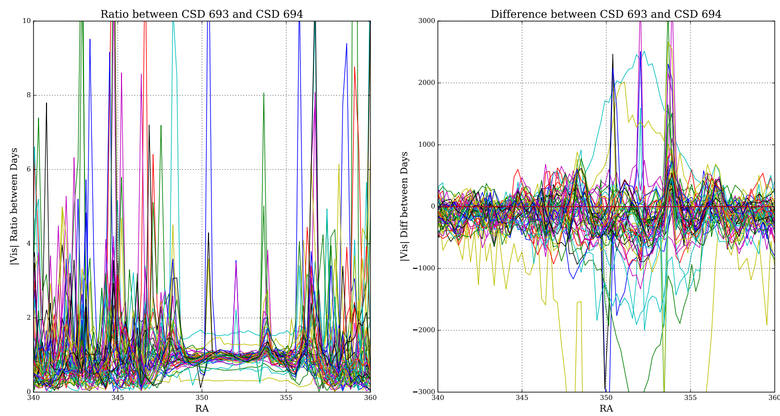

Figure 3.9: The right hand panel shows the ratio between the two panels of Figure 3.9 and the left hand panel shows the difference between them. Each is a short inter-cylinder baseline at frequency 518 MHz on CSD 693 and CSD694. The deviation between days is on the order of 100, while the spread within a day is on the order of 1000, meaning the deviation between instances is much larger. Though the ratio is relatively large for areas outside of the central transit, at the transit peak, it is approximately 1.

Chapter 4

Redundant Baseline Calibration with Perturbed Beams

4.1 Creating a Simulation with Beam Perturbations

As we saw in Chapter 3, the actual CHIME Pathfinder instrument shows significant perturbations in beams, and it is therefore valuable to have a method of simulating such beam perturbations to test analysis and calibration techniques in a controlled, realistic setting. We have therefore constructed an extension to the existing CHIME simulation pipeline which can flexibly incorporate beam perturbations. The design of the beam perturbation simulation is described in greater detail in the Appendix.

4.1.1 Design of beam perturbation

We treat the beam perturbation or perturbations as additional parameters in the expression for beam amplitude, A. We then Taylor expand to first order and move forward with a first order representation of the beam.

In other words, we transform the beam function from $A(\hat{\mathbf{n}};\phi)$ to $A(\hat{\mathbf{n}};\phi;\alpha)$, where α is some perturbation to a parameter of the beam, such as full width at half maximum or pointing. We could also use this same structure to introduce an arbitrary number of perturbations, designating each by a Greek letter: α, β, δ, etc.

Our expression for the visibility is then

$$V_{ij} = \int d^2n \, A_i\left(\hat{\mathbf{n}};\phi;\alpha_i\right) A_j^*\left(\hat{\mathbf{n}};\phi;\alpha_j\right) \exp\left(2\pi\hat{\mathbf{n}}\cdot\mathbf{u_{ij}}\right) S, \tag{4.1}$$

where α_i and α_j are perturbations unique to each feed i and j, and S represents the sky.

One way to generate such visibilities would be to uniquely determine beam amplitudes $A\left(\hat{\mathbf{n}};\phi;\alpha\right)$. However, directly incorporating unique beam functions would be computationally expensive, as we would have to generate beam transfer matrices for each individual input (256 for a CHIME Pathfinder sized simulation) and then would have no redundancy present in our modified pipeline, requiring almost an entirely separate pipeline.

We do not require an exact representation of the individual beam functions, but only a reasonable approximation. Provided the beams are different from one another in a way we can recreate, we have achieved our goal. Therefore, we Taylor expand the beam functions as a function of α and keep only terms to first order. Thus, we expand each A as

$$A_i\left(\hat{\mathbf{n}};\phi;\alpha\right) \approx A^{(0)}\left(\hat{\mathbf{n}};\phi;0\right) + A^{(1)}\left(\hat{\mathbf{n}};\phi;0\right)\alpha + O\left(\alpha^2\right) \tag{4.2}$$

and rewrite the visibility expression as

$$\begin{aligned} V_{ij} = \int d^2n \, & \left(A_i(0)\left(\hat{\mathbf{n}};\phi\right) + A_i^{(1)}\left(\hat{\mathbf{n}};\phi\right)\alpha_i\right) \\ & \times \left(A_j^{*(0)}\left(\hat{\mathbf{n}};\phi\right) + A_j^{*(1)}\left(\hat{\mathbf{n}};\phi\right)\alpha_j\right) \\ & \times \exp\left(-2\pi\, i'\, \hat{\mathbf{n}}\cdot\mathbf{u_{ij}}\right) S. \end{aligned} \tag{4.3}$$

Keeping only terms to first order, we have

$$\begin{aligned} V_{ij} = \int d^2n \, & \left[A_i^{(0)}\left(\hat{\mathbf{n}};\phi\right)A_j^{*(0)}\left(\hat{\mathbf{n}};\phi\right) + A_i^{(0)}\left(\hat{\mathbf{n}};\phi\right)A_j^{*(1)}\left(\hat{\mathbf{n}};\phi\right)\alpha_j + A_i^{(1)}\left(\hat{\mathbf{n}};\phi\right)A_j^{*(0)}\left(\hat{\mathbf{n}};\phi\right)\alpha_i\right] \\ & \times \exp\left(-2\pi\, i\, \hat{\mathbf{n}}\cdot\mathbf{u_{ij}}\right) S. \end{aligned} \tag{4.4}$$

We can think of this visibility then as being composed of three separate components: an unperturbed visibility with beam $A_i^{(0)}A_j^{*(0)}$ and two combinations of an unperturbed and perturbed beam, i.e. with beams $A_i^{(1)}A_j^{*(0)}\alpha_i$ and $A_i^{(0)}A_j^{*(1)}\alpha_j$.

Should we desire more than one perturbation per polarization, we would simply repeat the Taylor expansion and recombination process for each additional parameter β, γ, δ, etc.

A way to avoid creating a full N set of beam transfer matrices then becomes apparent. Ordinarily, we would generate a set of beam transfer matrices for all A. Now, we generate a set for all $A^{(0)}$ and $A^{(1)}$ for each feed. We then use these beam transfer matrices to create a sidereal stream as in the usual simulation code. However, the usual code creates three times the correct number of products, as it creates separate products for each combination of beam amplitudes and derivatives. Until this point, we have not had to alter the simulation pipeline code, merely the input. Following this point, we want to combine perturbed and unperturbed components, so we must alter the pipeline code to accommodate the new product structure.

4.2 Redundant Baseline Calibration Results

We previously supposed that a deviation expected results in redundant baseline amplitude calibration could arise from feed-to-feed beam variations. In the redundant baseline algorithm, one of our key suppositions is that the beam function for each feed is identical. When that assumption fails and there is feed-dependent beam information in the input visibilities, the redundant baseline algorithm will attempt to incorporate that information into the only solely feed-dependent output it has: the gain values. Therefore, we expect that a perturbed beam's beam structure will be visible in redundant baseline gain amplitude results. (This conclusion is supported by a comparison with holography data in Chapter 5.)

It can be slightly complicated to establish a direct probe of the simulated beam's structure as we do not have the ability to do e.g. simulated holography measurements. Plotting slices of the beam map is also complicated by the fact that the perturbation values α are added not in the initial telescope definition but in the "ExpandPerturbedProducts" step. For the moment, then, we use a slightly indirect probe. We run the redundant baseline algorithm on both the final simulated data set with gains, receiver temperature, and noise and an intermediate timestream product, prior to the addition of gains. This intermediate product has no feed-to-feed variation from complex gains, so the gain amplitude values that the redundant baseline algorithm generates are solely driven by feed-to-feed beam variations.

Our first analysis step is the simplest: we simply plot the beam-only redundant baseline results, the full redundant baseline results results, and the input gains. We notice that the full analysis mostly tracks the input gains, but has noticeable deviation from them, particularly near point source transit. We hypothesize that the full result minus input gains quantity should be equivalent to the beam only result (up to some constant offset). We plot these results and find this to be the case.

Correspondingly, we find that if we take the ratio of the redundant baseline solution with gains and beam variations to the redundant baseline solution with only beam variations, we recover the input gain variations as well as in the unperturbed case. This is also highly encouraging, as it indicates that the addition of beam variations does not irreparably corrupt the gain solutions, provided the information from beam and gain variations can be separated.

4.3 Solving for Beam Perturbation Values

Given our conclusion that the redundant baseline gain solutions in the presence of beam perturbations trace the per-feed beam perturbation structure, we would like use those solutions to solve for quantitative information about the beam perturbations from the redundant baseline gain information. Here, we attempt to find a method to quantify the information redundant baseline calibration preserves and/or to develop a scheme to solve for beam perturbations (especially width perturbations) using the nominal redundancy of an instrument like the CHIME Pathfinder.

This analysis assumes two very important things. First, we assume that the beam pattern can be modeled as a fiducial beam and an arbitrary number of perturbations to the fiducial model. Second, we assume that the telescope in question has an alternative calibration method for complex gains such as a noise source or a carefully constructed thermal model.

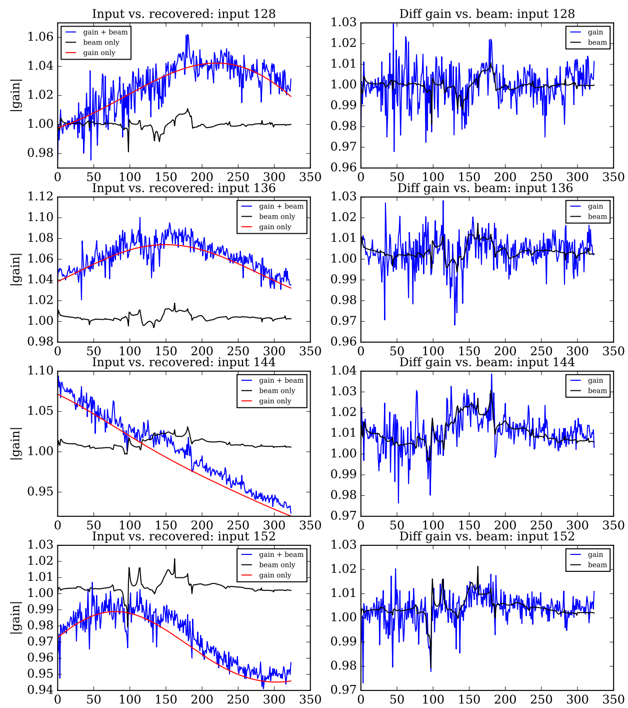

Figure 4.1: The left panel displays perturbed beam redundant baseline gain amplitude results, for selected feeds of a given polarization in a simulated perturbed beam telescope, compared with the input gains and a redundant baseline analysis conducted on data without gains added (i.e. an analysis that detects only beam effects). The right panel compares the beam only analysis to the full analysis with redundant baseline gains subtracted. This figures shows only a small sampling of inputs; more are shown in the Appendix. We notice that the full redundant baseline solutions deviate from the input gain solution near the peak of the beam only solutions. We infer that this deviation is caused by the beam perturbations, and the right hand panel confirms this.

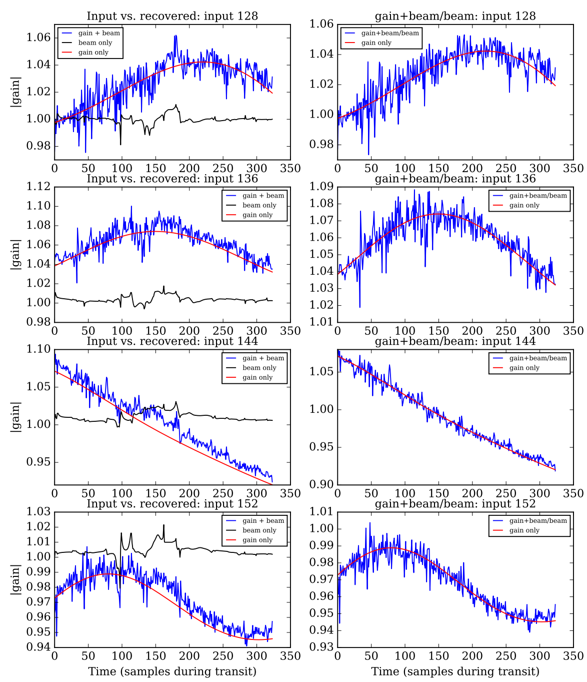

Figure 4.2: The left panel shows perturbed beam redundant baseline gain amplitude results, for selected feeds of a given polarization in a simulated perturbed beam telescope, compared with the input gains and a redundant baseline analysis conducted on data without gains added (i.e. an analysis that detects only beam effects. The right panel compares the ratio of the beam and gain to beam only analysis and the input gain variations. the full analysis with redundant baseline gains subtracted, showing that the ratio recovers the correct input gain. As in Figure 4.1, results for more inputs are shown in the appendix.

4.3.1 Solving with One Perturbation of Known Structure

Before we begin, we apply redundant baseline calibration to data simulated with these beam structures and subtract the known gains from the redundant baseline results. Alternatively (and in practice, preferably) we can just run redundant baseline calibration on a simulation with no gains added. Thus all the gain deviation from 1 must be due to beam effects. We now have recovered visibilities, which are created by recombining the recovered redundant baseline gains and true visibilities. This, mathematically, is

$$\mathbf{V_{ij}^{rec}} = \exp(\mathbf{M\hat{x}}), \tag{4.5}$$

where $\mathbf{M}$ is the coefficient matrix in the redundant baseline problem as in Chapter 2 and $\mathbf{\hat{x}}$ is the vector of recovered "gains" and true visibilities.

We can expand $\mathbf{V_{ij}}^{\mathrm{rec}}$ in equation 4.5 as

$$\begin{aligned}\mathbf{V_{ij}^{rec}} = \mathbf{V_{ij}}^0 &+ \sum_{(ij)} \alpha_i \int A_i^1 A_j^{*0} e^{-(2\pi i \mathbf{\hat{n}} \cdot \mathbf{u_{ij}})} T(\mathbf{\hat{n}}) d^2\mathbf{\hat{n}} \\ &+ \sum_{(ij)} \alpha_j \int A_j^{*1} A_i^0 e^{(-2\pi i \mathbf{\hat{n}} \cdot \mathbf{u_{ij}})} T(\mathbf{\hat{n}}) d^2\mathbf{\hat{n}}\end{aligned} \tag{4.6}$$

We rewrite this into an entirely matrix equation setup

$$\mathbf{V_{ij}^{rec}} = \mathbf{V_{ij}}^0 + \mathbf{W_{ij}}\alpha. \tag{4.7}$$

Here, $\mathbf{V_{ij}^{rec}}$ is a vector of the recovered visibilities, $\mathbf{V_{ij}}^0$ is a vector of the 0th order (unperturbed) visibilities, α is a vector of beam perturbation values, and $\mathbf{W_{ij}}$ is a matrix containing the appropriate V_{ij}^1 terms. The $\mathbf{W}$ matrices must be designed to encapsulate the structure of the beam perturbations and be able to be combined with the only desired perturbation value. While $\mathbf{W}$ is similar to $\mathbf{M}$ in that it generates combinations of different feeds, it is different in that its entries are either 0 or a perturbation structure not 0 or 1. $\mathbf{W}$ is $N(N+1)/2$ rows by N columns, where N is the number of feeds in a given telescope. Each row corresponds to a correlation ij. In a given row W_{ij}, the ith element is of the form $V_{ij,i}^1 = \int A_i^1 \, A_j^{*0} \; T \; e^{-2\pi i \mathbf{\hat{n}} \cdot \mathbf{u_{ij}}} \; d^2\mathbf{\hat{n}}$ and the jth element is of the form $V_{ij,j}^1 = \int A_i^0 \, A_j^{*1} \; T \; e^{-2\pi i \mathbf{\hat{n}} \cdot \mathbf{u_{ij}}} \; d^2\mathbf{\hat{n}}$. Then, α is free to be a vector of length N where each element is the perturbation corresponding to a given feed. For example, in a four feed one cylinder telescope, **W** would be

$$\mathbf{W} = \begin{pmatrix} V^1_{ij,i} & V^1_{ij,j} & 0 & 0 \\ V^1_{ij,i} & 0 & V^1_{ij,j} & 0 \\ V^1_{ij,i} & 0 & 0 & V^1_{ij,j} \\ 0 & V^1_{ij,i} & V^1_{ij,j} & 0 \\ 0 & V^1_{ij,i} & 0 & V^1_{ij,j} \\ 0 & 0 & V^1_{ij,i} & V^1_{ij,j} \end{pmatrix} \tag{4.8}$$

Then, if we presume we understand the structure of the beam (but not the perturbation value), we know $\mathbf{V^0}$ and $\mathbf{W}$. We can then solve our matrix equation for α

$$\mathbf{W}^+ \left(\mathbf{V}^{\text{rec}}_{\mathbf{ij}} - \mathbf{V}^{\mathbf{0}}_{\mathbf{ij}}\right) = \alpha. \tag{4.9}$$

This method works well if we are able to provide full information (phase and amplitude) about the recovered visibilities. However, we would like to attempt this solution with just information about the recovered redundant baseline amplitudes.

4.3.2 Solving for Beam Perturbations with only Amplitude Information

In our current redundant baseline solution, we do not have full phase and amplitude information about recovered visibilities or "gains," but only have amplitude. Therefore, instead of evaluating Equation 4.5, we need to evaluate

$$\ln|\mathbf{V_{ij}}^{\text{rec}}| = \ln|\mathbf{V}^{\mathbf{0}}_{\mathbf{ij}} + \mathbf{W_{ij}}\alpha|. \tag{4.10}$$

We can evaluate the absolute value function by multiplying $\mathbf{V_{ij}}^0 + \mathbf{W_{ij}}\alpha$ by its complex conjugate, giving us

$$\ln|\mathbf{V_{ij}}^{\text{rec}}| = \frac{1}{2}\ln\left(\left(\mathbf{V}^{\mathbf{0}}_{\mathbf{ij}} + \mathbf{W_{ij}}\alpha\right)\left(\mathbf{V}^{\mathbf{0}}_{\mathbf{ij}} + \mathbf{W_{ij}}\alpha\right)^*\right) \tag{4.11}$$

We multiply out and Taylor expand to first order and are left with

$$\ln|\mathbf{V}^{\text{rec}}_{\mathbf{ij}}| = \ln|\mathbf{V}^{\mathbf{0}}_{\mathbf{ij}}| + \frac{1}{2}\alpha\left(\frac{\mathbf{V}^{\mathbf{1}*}_{\mathbf{ij}}}{\mathbf{V}^{\mathbf{0}*}_{\mathbf{ij}}} + \frac{\mathbf{V}^{\mathbf{1}}_{\mathbf{ij}}}{\mathbf{V}^{\mathbf{0}}_{\mathbf{ij}}}\right) \tag{4.12}$$

We then redefine the non-zero elements of $|W_{ij}|$ to be

$$|\mathbf{W_{ij}}| \equiv \frac{1}{2}\left(\frac{\mathbf{V^{1*}_{ij}}}{\mathbf{V^{0*}_{ij}}} + \frac{\mathbf{V^{1}_{ij}}}{\mathbf{V^{0}_{ij}}}\right) \tag{4.13}$$

and revert to the simpler expression

$$\ln|\mathbf{V^{rec}_{ij}}| = \ln|\mathbf{V^{0}_{ij}}| + \alpha|\mathbf{W_{ij}}|. \tag{4.14}$$

Finally, we solve for α:

$$\alpha = |\mathbf{W_{ij}}|^{+}\left(\ln|\mathbf{V^{rec}_{ij}}| - \ln|\mathbf{V^{0}_{ij}}|\right). \tag{4.15}$$

If we assume that we know V^0 and V^1, this procedure produces a consistent solution. See Figure 4.3 for an example with a small two cylinder, sixteen input telescope. We test this solution both using “input visibilities” - the actual perturbed visibilities generated by the simulation - and the “recovered visibilities” - the visibilities recovered using the redundant baseline solution. For the “input” case we find near perfect agreement. The absolute deviations from the correct solution remain approximately the same size for all perturbation values. Due to the range of actual values for α, from about 0.01 to about 0.001, percent differences corresponding to feeds with smaller perturbation values due have muh larger percent differences. For the “recovered” case, we find the agreement is relatively poor, but is correct to the first significant digit or better for all feeds. This is not surprising, as any variation in the redundant baseline solution propagates to the beam perturbation solution. The process of applying redundant baseline calibration does result in a loss of information, which cannot be regained for the perturbation solution.

This method solves for alpha uniquely at each time point. This means we can learn additional information by solving for a time series of α values. Based on the design of our simulation, α should be constant over all time, so we would be able to improve our estimate by considering possible variation in the recovered α values. We notice that α solution are not actually constant, but they seem to improve in constancy in regions of higher signal to noise. As in the individual time point solution, we find that the beam perturbation values from the “recovered” data

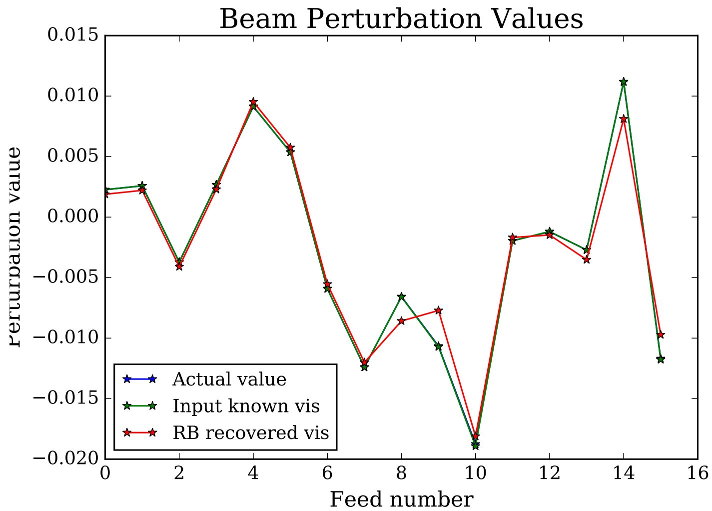

Figure 4.3: This figure shows beam perturbation solution for a 16 feed total telescope with random perturbations α applied in the beam width of all feeds has a few noticeable features. The actual input α values used to create the simulation are plotted in blue, but are almost exactly overplotted by the green values. The green values represent the result when the output of the simulation is used as the "recovered visibility," and we therefore expect this close correspondence. The red curve represents the results using the redundant baseline calibration results as the "recovered visibility" and is noticeably less accurate than the green.

set deviate slightly from the exact input α values used in the simulation, but this remains consistent with loss of information in the redundant baseline calibration algorithm. One apparently odd characteristic is that the ordering of α values from largest to smallest is not necessarily constant at all times. This arises from the independent solutions at each time point. This allows imprecise solutions for α values which are close in magnitude to cross one another at different time points.

4.3.3 Relaxing Assumptions about V_{ij}^0 and V_{ij}^1

In our analysis thus far, we have assumed that we know the unperturbed first order perturbed components of the visibility, V_{ij}^0 and V_{ij}^1. However, this assumes that we

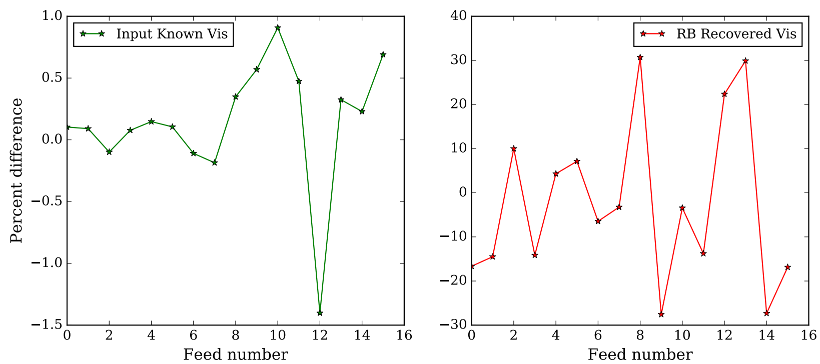

Figure 4.4: We see here the percent difference between the actual simulated α beam perturbation values and the recovered α values for both the actual input visibility and the recovered redundant baseline solution as input visibility. We observe that the input visibility has quite good agreement with the actual perturbation values, but the recovered redundant baseline solution is useful only for order of magnitude approximations.

know the convolution of the zeroth and first order components of the beam with the sky at all times. This is necessarily not the case for real data, so we would like a method to replace these parameters.

One way forward is to take advantage of the near-redundancy of the baselines. Each instance of a baseline is composed of a zeroth order part, which is redundant with other baselines of the same length, and a first order part which is not. Therefore, presuming the perturbation values are scattered about zero, we can take an average of all instances of a given nominally redundant baseline and substitute this average for V^0_{ij}. In other words,

$$\overline{V_{i-j}} \approx V^0_{i-j} = \frac{1}{m}\sum_{k=(ij)_0}^{(ij)_m} V_k^{meas}, \tag{4.16}$$

where m is the number of instances of a given nominally redundant baseline and I am using the convention that V_{ij} is an individual correlation between feeds i and j and V_{i-j} is the redundant correlation for all i's and j's with the baseline spacing i-j.

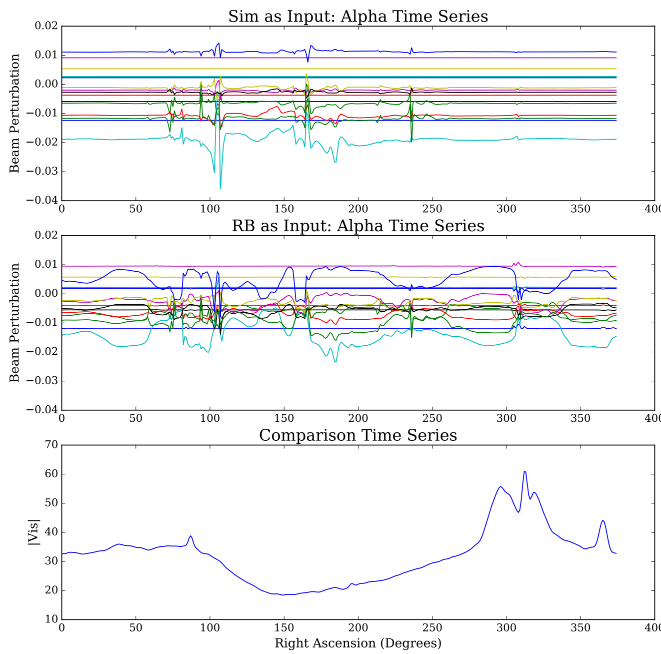

Figure 4.5: This figure examines the beam perturbation solution for a 16 feed total telescope with all feed perturbed for a sidereal day. Each time point is solved independently, but time dependent features are consistent with the redundant baseline solution more generally, e.g. that solution improves with improved signal to noise ratio. The top panel shows the result for the solution using the simulation output as the "recovered visibility," the middle panel shows the result for a solution using the redundant baseline calibration results as the "recovered visibility," and the final panel shows a single visibility's time series during this sidereal day.

Then, we can subtract the average $\overline{V_{i-j}}$ from each actual measured V_{ij} and are left with an approximation of the perturbed components, which we will call ΔV_{ij}, for each instance of the baseline.We know in principle that if $\overline{V_{i-j}}$ is a good approximation for V_{ij}^0, ΔV_{ij} should be the perturbed portion of V_{ij} or rather. In summary,

$$\Delta V_{ij} \equiv V_{ij}^{\text{meas}} - \overline{V_{i-j}} \approx \alpha_i V_{ij}^{01} + \alpha_j V_{ij}^{10}. \tag{4.17}$$

However, we do not have enough information to separate ΔV_{ij} into i and j perturbed portions, so we consider the entire perturbed portion as being

$$\Delta V_{ij} = W_{i-j}\alpha, \tag{4.18}$$

where W_{i-j} is a matrix which in some way encodes the structure of the perturbed beam and α is the perturbation values as before. It is important to note that this W_{i-j} is *not* identical to $\mathbf{W}$ referenced previously. That $\mathbf{W}$ has independent entries for each component of the perturbation structure, whereas this W_{i-j} is a single value per unique baseline.

At first glance, we cannot uniquely solve for all values of W_{i-j} and α, but if we can solve for each α value in terms of one of the α values. In this case, we solve for $(\alpha_j + \alpha_0)$ for each α_j. We can also solve for a combination of $(\alpha_j + \alpha_0)$ and $(\alpha_j - \alpha_0)$, and use the additional information from the $(\alpha_j - \alpha_0)$ solutions to determine the sign of each α_j value. However, we focus on here on the $(\alpha_j + \alpha_0)$ solutions because instances of $(\alpha_j - \alpha_0)$ do not exist for all α_j.

We begin by writing ΔV_{0i} for all i. These equations take the form

$$\begin{aligned}
\Delta V_{01} &= W_{0-1}(\alpha_0 + \alpha_1) \\
\Delta V_{02} &= W_{0-2}(\alpha_0 + \alpha_2) \\
&\vdots \\
\Delta V_{0n} &= W_{0-n}(\alpha_0 + \alpha_n).
\end{aligned} \tag{4.19}$$

We will then chain together instances of these known $0-i$ baselines to create a system of equations which can be linearized and solved in much the same way as the redundant baseline problem. This chaining process is simplest for the shortest baseline, $0-1$. For example, for V_{12}, we write

$$\Delta V_{12} - \Delta V_{01} = W_{0-1}(\alpha_2 + \alpha_1) - W_{0-1}(\alpha_1 + \alpha_0) = W_{0-1}\left(\alpha_2 - \alpha_0\right), \qquad (4.20)$$

and for V_{13} we write

$$\begin{aligned}\Delta V_{13} - \Delta V_{12} + \Delta V_{01} = W_{0-1}(\alpha_3 + \alpha_2) - W_{0-1}(\alpha_2 + \alpha_1) + W_{0-1}(\alpha_1 + \alpha_0) \\ = W_{0-1}\left(\alpha_3 + \alpha_0\right).\end{aligned} \qquad (4.21)$$

In chaining instances of the shortest baseline, we create a combination of $(\alpha_j + \alpha_0)$ terms and $(\alpha_j - \alpha_0)$ terms. We solve for both terms simultaneously, but focus on the $(\alpha_j + \alpha_0)$ terms.

For this shortest north-south baseline, we create this type of combination for each subsequent instance by using the formula

$$W_{0-1}\ (\alpha_m \pm \alpha_0) = \sum_{i=1}^{m} \Delta V_{i,i-1}(-1)^{(i+m)}, \qquad (4.22)$$

where m is the feed we are isolating. Equations with $-1^m = -1$ are $(\alpha_j - \alpha_0)$ equations, while equations with $-1^m = 1$ are $(\alpha_j + \alpha_0)$ equations.

We can write a general relationship analogous to Equation 4.22, but we must be slightly more creative. For baselines longer than the shortest north-south baseline, we will not be able to chain every instance, but will only be able to use a limited set. Our more general relationship is

$$W_\delta \alpha_{m\cdot\delta} = \sum_{i=0}^{m} \Delta V_{i\delta,(i+1)\delta}(-1)^{i+m}, \qquad (4.23)$$

where δ is the difference in feed indices, m is the instance number, and ΔV is as before. As before, odd m values lead to $(\alpha_j - \alpha_0)$ equations and even m values lead to $(\alpha_j + \alpha_0)$ equations.

For each baseline we will have a different number of instances, varying inversely with the length of the baseline. For example, in our 32 feed one cylinder telescope, we will have 31 ($m = 0$ to $m = 31$) instances of the shortest baseline, but only 15 instances ($m = 0$ to $m = 15$) of the second shortest baseline and 10 instances ($m = 0$ to $m = 10$) for the third shortest baseline. This does imposes a relatively large minimum size on the telescope. For a one cylinder telescope, we require 32 inputs per polarization to solve for both $(\alpha_j + \alpha_0)$ and $(\alpha_j - \alpha_0)$ values.

We have now developed a system of equations of the form $\Delta V_{ij} = W_{i-j}\alpha_j$. We linearize these equations in the same manner as in the redundant baseline problem outlined in Chapter 2, by taking the logarithm. Then, our equations are of the form

$$\log|V_{ij}| = \log|W_{i-j}| + \log|\alpha_j|. \tag{4.24}$$

Continuing as in the standard redundant baseline problem, we construct a matrix $\mathbf{M}_{\alpha W}$ reminiscent of the $\mathbf{M}$ matrix of Chapter 2 and re-write our problem as a vector equation

$$\Delta\mathbf{V} = \mathbf{M}_{\alpha W}\hat{\mathbf{x}}, \tag{4.25}$$

and solving for $\hat{\mathbf{x}}$, a vector of α and W_{i-j} values.

This solution is underdetermined for our small test problems, meaning that the result represents a minimum norm solution and not a unique determination of the values for $(\alpha_j + \alpha_0)$, $(\alpha_j - \alpha_0)$, and W_{i-j}. Therefore, we do not necessarily expect to see perfect recovery of the actual values for $(\alpha_j + \alpha_0)$, $(\alpha_j - \alpha_0)$, and W_{i-j}. By examining the null space of the $\mathbf{M}_{\alpha W}$, we can make more precise determinations of which values may deviate from expected. In future analysis, applying additional independent constraints may enable us to determine values more accurately, but this underdetermined solution offers a starting point. A larger telescope model would create an over-determined problem instead of an underdetermined one, but is left for future work.

By determining the ratio of $(\alpha_j + \alpha_0)$ to a particular $(\alpha_j + \alpha_0)$ value (and the same for $(\alpha_j - \alpha_0)$) for both the original simulated α values and the recovered $(\alpha_j + \alpha_0)$ and $(\alpha_j - \alpha_0)$ results, we can begin to assess the effectiveness of this approach. Figures 4.6 and 4.7 show a first result, which correlates strongly with the known values determined from the actual simulations.

Figure 4.6: We examine the ratio $(\alpha_j + \alpha_0)/(\alpha_2 + \alpha_0)$ for both the recovered $(\alpha_j + \alpha_0)$ values and the actual $(\alpha_j + \alpha_0)$ values for a 32 feed one cylinder telescope. The solution values are close to the expected values, except at noticeable outlier feed 16. The imperfect correspondence to the correct answers, in spite of the absence of noise, is due to the underdetermined nature of the problem.

4.4 Conclusions

We simulated a CHIME Pathfinder like telescope with a beam perturbation in full width half maximum, then applied the redundant baseline calibration method to it. By applying the redundant baseline calibration algorithm as outlined in Chapter 2, we have confirmed our hypothesis that in non-ideal cases, redundant baseline calibration is sensitive to the effects of non-identical beams.

We also examined possible methods for extracting information about beam perturbations from either gain-calibrated data or from gain-calibrated, redundant baseline calibrated data. However, our approach is limited in one case by the assumption of a known beam and sky model at all times. In both cases, our approach is limited by the assumption of noiseless data and perfect calibration.

However, we conclude that there is strong evidence for the utility of redundant baseline calibration and related methods in characterizing feed-to-feed beam perturbations. In Chapter 5, we will demonstrate the usefulness of redundant baseline

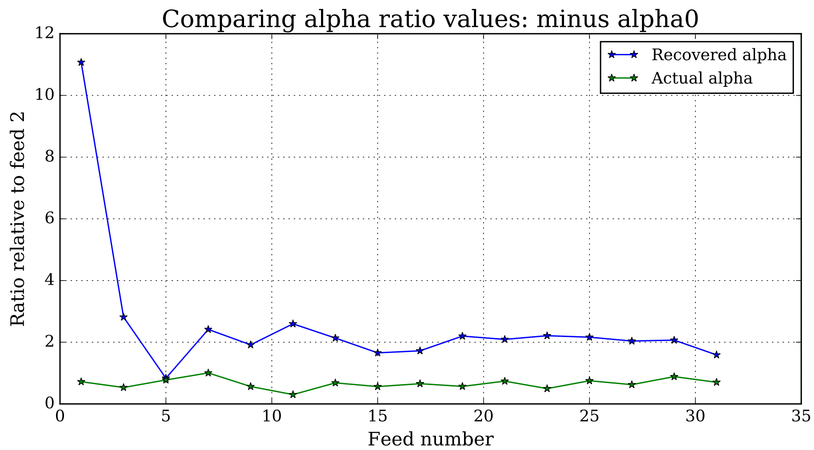

Figure 4.7: We examine the ratio $(\alpha_j - \alpha_0)/(\alpha_2 - \alpha_0)$ for both the recovered $(\alpha_j - \alpha_0)$ values and the actual $(\alpha_j - \alpha_0)$ values for a 32 feed one cylinder telescope. Unlike the $(\alpha_j + \alpha_0)/(\alpha_2 + \alpha_0)$ solution, the $(\alpha_j - \alpha_0)/(\alpha_2 - \alpha_0)$ solution exists only for approximately every other input. As these results derive from the same underdetermined problem, they too deviate from the expected values in ways that can be examined more carefully using the null space.

calibration in predicting beam perturbations in actual data by comparing to point source holography, a well established beam mapping approach.

Chapter 5

Holography and Redundant Baseline Calibration as Beam Probes

Though results from simulations are powerful evidence for the efficacy of redundant baseline calibration in finding feed-to-feed beam variations in the CHIME Pathfinder, this result is essentially worthless if it cannot be applied to real CHIME Pathfinder data. However, conducting such an analysis on real CHIME Pathfinder data is challenging as we do not know the exact beam pattern of the telescope or the actual complex gain. We must find an additional test to check the validity of our redundant baseline analysis. We choose to compare our redundant baseline analysis to holographic beam measurements conducted using the CHIME Pathfinder and the John A. Galt 26 m telescope at DRAO (hereafter the 26 m telescope). As both methods should measure beam amplitudes, if our redundant baseline analysis is working as we believe it is, the two methods will agree.

5.1 Holography: A method for probing CHIME beams

The importance of precise beam measurements for an instrument like CHIME has long been known, and therefore the CHIME collaboration has been engaged in mapping the full two-dimensional primary beam of each feed with point-source holography during the operation of the CHIME Pathfinder [6].

Holographic beam mapping is a recognized technique to make beam measurements in radio telescopes. The maps are made by tracking a bright point source with one reference telescope while correlating that telescope with the CHIME Pathfinder. During the transit, we measure the source's track through the stationary CHIME Pathfinder beam and are thus measuring the east-west track of the CHIME Pathfinder beam. Specifically, we introduce the 26 m telescope as an additional input, correlating it as if it were a CHIME pathfinder channel. Thus, we create a set of visibilities

$$V_{26\,i} \propto A_{26}\, A_i^*\left(\hat{n}_{ps};\phi\right) T\left(\hat{n}_{ps}\right) \exp\left[-2\pi i \hat{n}_{ps}\cdot \mathbf{u_{ij}}\right] \tag{5.1}$$

As the 26 m telescope can only point at one point source at a time, we require multiple measurements at different declinations to properly measure the north-south shape of the beam. Holography efforts for the CHIME Pathfinder to date use Cygnus A, Taurus A, Virgo A, Hercules A, Hydra A, Perseus B, and 3C_295 [6].

5.2 Comparing redundant baseline beam measurements with holography data

To evaluate our redundant baseline beam probe's efficacy on real data, we used data from the July 2015 Cygnus A holography run combined with a redundant baseline solution for the same transit, applying our best practices, inter-cylinder only redundant baseline algorithm (as described in Chapters 2 and 3) to the uncalibrated CHIME pathfinder data from the same transit.

The holography analysis generates absolute measurements of the beam pattern, whereas our redundant baseline results are both only relative and only applicable to feed-to-feed variations (e.g. not a complete trace of the beam). Therefore, we must be a bit creative in comparing the two data sets. Specifically, both should record relative, feed-to-feed variations. Therefore, we primarily analyze the data by looking at ratios between each CHIME Pathfinder feed and an arbitrarily selected reference CHIME feed. In order to orient ourselves in analyzing the holography data, we do a brief comparison by eye.

5.2.1 Comparing by eye

In the simplest method for comparing redundant baseline transits with holography transits, we simply plot both the redundant baseline gains and the holography transits on the same axes and looks for correlation between the slope of the gain (positive or negative) and the shift in the peak of the transit (right or left) from the expected point. To do this comparison, we must scale up the holographic beam measurements, as their value in correlator units is of the order of 10^{-9} while the gains are of the order 1 by design.

Figures 5.1 and 5.2 show a summary of these plots, showing every eighth feed on the CHIME Pathfinder, with the goal of acquiring a sampling of the cylinder. There does appear to be a correlation between the slope of the gains and the shift in the peak from the average peak sample, but it is not easily quantified in this “by eye” analysis.

5.2.2 Ratio analysis

We anticipate that the redundant baseline gains include information both about the complex gain and beam of each receiver. We further approximate that during a bright point source transit, each feed sees approximately the same very bright sky, dominated by a source at one declination. Therefore, we can roughly approximate the visibility as

$$V_{ij} = g_i\, g_j^*\, A_i\, A_j^*\, T, \tag{5.2}$$

where T is a constant sky value. We further approximate that CHIME Pathfinder gains are constant on the timescale of a single point source transit. Thus, relative shape variations between two redundant baseline gain solutions necessarily arise from beam variations (since the sky and the gain are taken to be constant).

We can simply represent the variations between each feed’s beam pattern by taking the ratio between each calculated redundant baseline gain and an arbitrarily chose reference gain (in this case, that of feed 65). In other words the ratio R_{RB} we

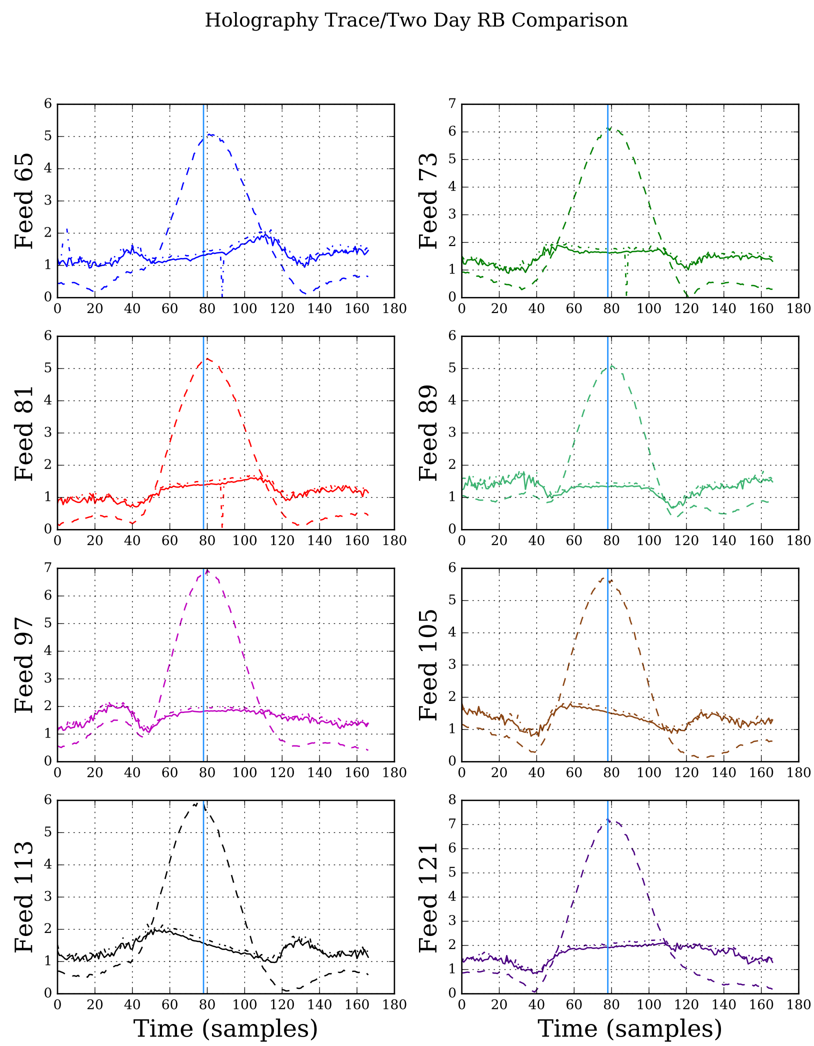

Figure 5.1: Comparison between holography transit and calculated redundant baseline gain for selected feeds on the west cylinder, looking only at results from the east west polarization. The two gain traces represent solutions for CygA transits on consecutive days. The average transit peak time is marked by the blue vertical line, and the shift to before or after the average peak time appears to correlate with the slope of the redundant baseline gain.

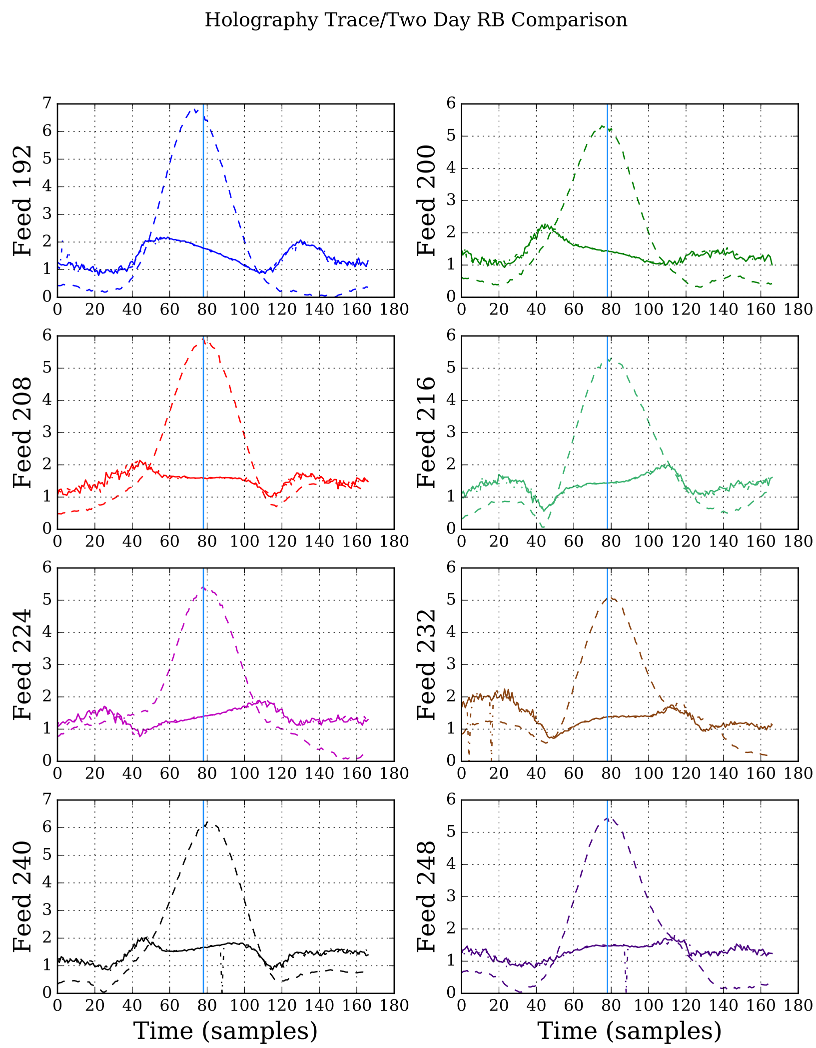

Figure 5.2: Comparison between holography transit and calculated redundant baseline gain for the selected feeds on the east cylinder, EW polarization. The average transit peak time is marked by the blue vertical line, and the shift to before or after the average peak time appears to correlate with the slope of the redundant baseline gain.

are interested in is

$$R_{\text{RB}} = \frac{g_{\text{RB},i}}{g_{\text{RB},65}} \tag{5.3}$$

for each feed i. The absolute level of each ratio is irrelevant (and approximately one), but the shape of the ratio traces the ratio of the feed's beam pattern to the reference feed's beam pattern. In other words, we examine both

We also replicate this ratio using the holography data. We take the ratio between each channel's holography cross-correlation and a reference channel's holography cross-correlation. In other words the ratio R_{holo} we are interested in is

$$R_{\text{holo}} = \frac{V_{\text{holo},i}}{V_{\text{holo, }65}}, \tag{5.4}$$

for each feed i. In this ratio, we cancel out the common portions of the beam pattern and are also left with a ratio describing feed-to-feed beam variations. Therefore, we can now make quantitative, head-to-head comparison between gain ratios and holography ratios to see whether the two show consistent beam variations.

If the two methods of analyzing the beams are consistent, we would expect the difference between them to be 0 during transit. It is however important to note that we would not necessarily expect consistency between the two methods before or after the point source transit. Our approximation in Equation 5.2 would no longer apply, and the comparison would no longer be meaningful.

For context, we first plot the redundant baseline measure of feed-to-feed beam variation independently of the holography data. This is calculated by dividing each redundant baseline gain by the value of a reference gain, from feed 65, at each time point. If our reference feed has anomalous structure in its gain, this is a dangerous method as it will propagate that structure into each ratio. Therefore, before calculating these ratios, we established that feed 65 had no significant anomalies relative to other feeds. The only traces at or near zero are those of "dead feeds," which are input channels that are either not functioning or not connected to CHIME channels.

We then compared the two ratios by taking the difference between them. This gives us a sense of the deviation between feed-to-feed beam variation found by

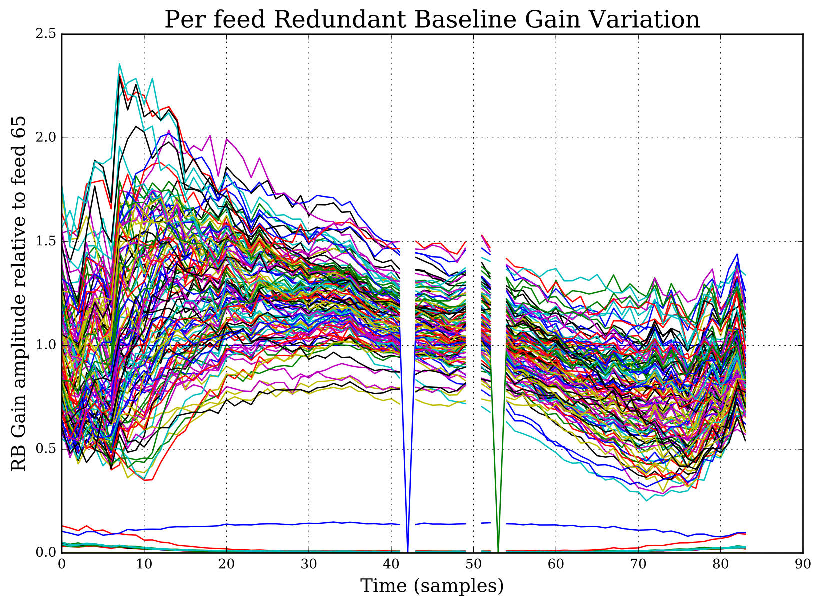

Figure 5.3: Redundant baseline gain amplitude results relative to a reference redundant baseline gain solution (from feed 65). The variations in this plot are expected to trace feed-to-feed beam amplitude variations. These feed-to-feed beam amplitude variations will then be compared to holography results to verify the correspondence.

holograph and by the redundant baseline gains. In other words, we examine

$$\Delta_{\text{RB-holo}} = \frac{V_{\text{holo},i}}{V_{\text{holo},65}} - \frac{g_{\text{RB},i}}{g_{\text{RB},65}}. \tag{5.5}$$

We quantified these results by we taking averages of the differences, both for each feed during transit and for each time over all feeds. Figure 5.4 shows the average difference between redundant baseline and holography beam estimates at each time point. Recall that the transit peaks at approximately sample 80 and that it reaches half its maximum value approximately 20 samples before and after transit, so the prime transit region is samples 60-100, with the region worth considering stretching between samples 40 and 120. Notice that, while areas outside of the transit in Figure 5.4 are quite erratic, areas in the sample 40-120 range are quite stable and close to 0, as expected.

We also found the average difference between ratios for each feed during the reasonable region (samples 40-120); these are plotted in Figure 5.5. Feeds toward the centre of the cylinder do seem to have slightly lower deviation from 0, but there is not a definitive pattern. Notice that during transit, all the feeds average differences of less than 0.1, and most average differences less than ± 0.05. We note that the average deviations in the difference between the ratios, is approximately 10% of the actual range of the ratios. In other words, the difference between redundant baseline beam analysis and holography beam analysis is significantly smaller than the variation in either method. We consider this to be strong evidence that the two methods are equivalent.

5.2.3 Conclusions

Based on our comparisons, we find that redundant baseline gains and holographic beam mapping report feed-to-feed beam variations during a point source transit which are in agreement. Though there are definite deviations between the two results, we see by comparing Figure 5.3 to Figure 5.4 that these deviations from each other are small relative to the size of the feed-to-feed beam variations.

We regard this analysis as observational evidence of the effect we observed in simulations in Chapter 4. It does seem that deviation from redundancy in CHIME Pathfinder data is generated by beam variations and these beam variations can be determined with redundant baseline calibration.

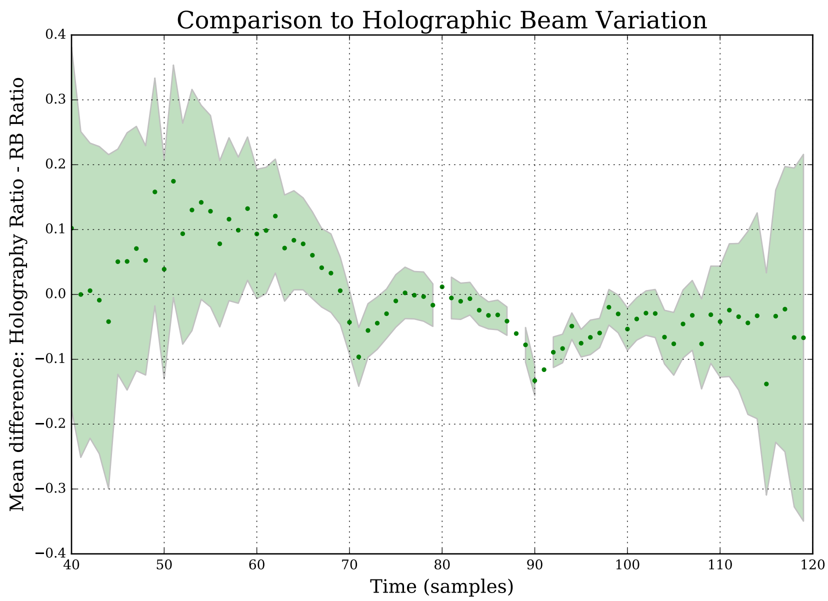

Figure 5.4: In this figure, the points represent the average ratio difference across feeds at a given time point, while the shaded band represents one standard deviation in the ratio difference. Notice that the deviation values (and especially the standard deviation) are small nearest the transit (which peaks at about sample 100) and larger near the edges. Additionally, notice that the average deviation at points nearest the transit (the center of the plot) are significantly smaller than the range in values within the set.

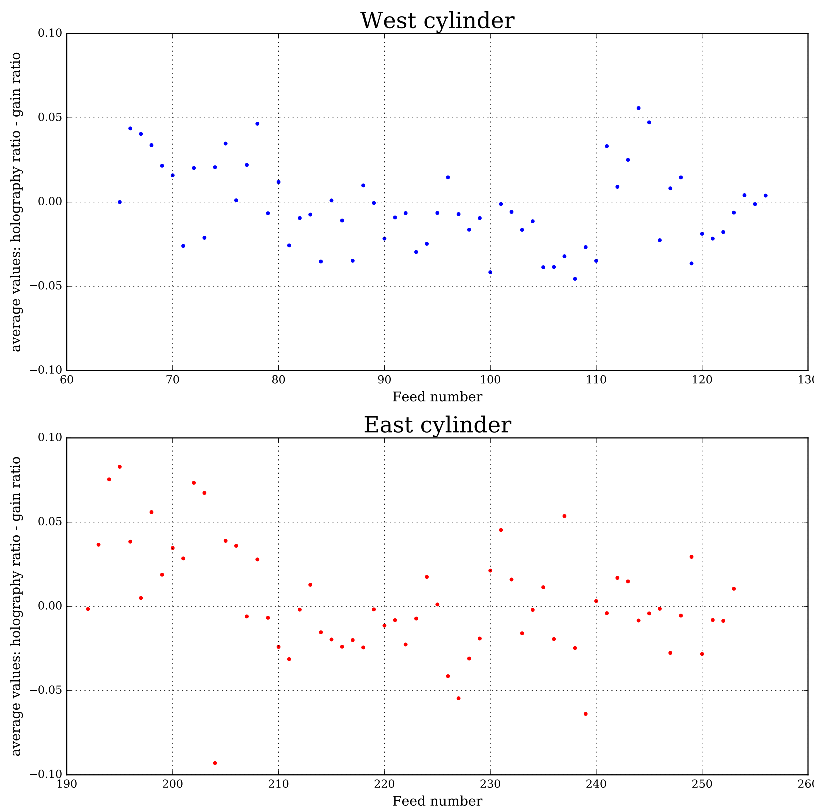

Figure 5.5: In this figure, we plot the average deviation between holography and redundant baseline gains for each feed in our sample, averaged over the 80 time samples closest to the transit peak (about 25 minutes of data). Notice that for all feeds, the deviation from zero is less than ± 0.1, much smaller than the values of the redundant baseline gain ratios themselves.

Chapter 6

Conclusions and Further Work

6.1 Conclusions

In this thesis, we have examined the potential for redundant baseline calibration in CHIME, based on both simulations of and data from the CHIME Pathfinder. We find that, for an ideal version of a CHIME-like telescope, redundant baseline amplitude calibration is very effective.

However, we observe that data from the CHIME Pathfinder is not well calibrated by redundant baseline amplitude calibration, in a manner that is consistent between days. This analysis also led us to discover significant non-redundancy in point source calibrated CHIME Pathfinder results. As this non-redundancy is also stable over time, we conclude both problems likely arise from beam variations between feeds.

We therefore develop and implement a perturbed beam simulation, and show that it displays similarly deviant redundant baseline results. However, we also show that redundant baseline amplitude calibration on such a system recovers information about feed-to-feed beam variations, both as a general trend and quantitatively for limited, noise-less cases.

Finally, we compare redundant baseline amplitude results from real CHIME Pathfinder data to holographic beam measurements, showing that both reproduce the same feed-to-feed beam variations. Given the importance of beam calibration for successful CHIME analysis, this is encouraging.

6.2 Future Work

In many ways, though, this work raises more questions than it answers. Of particular concern is the significant deviation from redundancy in point source calibrated visibilities. Full N^2 data for the final CHIME instrument cannot be stored, meaning the data must be compressed. At first glance, the logical way to do this is by collapsing data across redundant baselines. If these baselines are not actually redundant, this may lead to significant signal loss in compressed data and decreased ability to attain CHIME's science goals.

Besides new questions regarding redundancy, there is still work to be done in implementing the redundant baseline algorithm, in particular implementing phase calibration on a CHIME Pathfinder sized telescope. Several paths forward present themselves based on the work discussed here. These can be classed into three main categories. The first is extensions of the existing redundant baseline algorithm, the second is improvements of redundant baseline calibration based on our findings, and the third is further examination of redundant baselines as a probe of beam structure.

The first category is the most obvious, although not necessarily the simplest in practice. The current work does not include implementations of the phase component of the redundant baseline calibration algorithm in the CHIME Pathfinder. This is an important step if we would like to seriously use redundant baseline calibration to calibrate CHIME Pathfinder data. Additionally, we have not implemented any component of redundant baseline calibration for a full CHIME scale telescope. Because CHIME is a significant increase in scale beyond the CHIME Pathfinder, we will have to make this transition mindfully to avoid devouring computational resources. Finally, we have done very little examination of redundant baseline calibration across multiple frequencies. Preliminary examinations have not shown anything surprising, but a careful analysis should be done for the sake of completeness.

The second category, improvements to the redundant baseline calibration algorithm, is likely the richest and certainly the most open-ended. We have definitively shown that redundant baseline calibration is affected by feed-to-feed beam variations and that such beam variations are present in the CHIME Pathfinder. Preliminary examinations of CHIME feeds and structure indicate that they will remain a factor in CHIME. Therefore, if we would like to use redundant baseline calibration as a CHIME calibration method, we must find ways to use the algorithm in the absence of identical beams. While it is too early to assert a particular method as the definite path forward, approaches might either incorporate information from sky maps via iterative Gibbs sampling or might solve for additional parameters as basis functions of the non-identical beams.

The final category, further use of redundant baselines as a beam probe, diverges in two directions, from the work in solving for beam perturbations in simulations and from the work comparing to holography. The beam perturbation solutions presented here are somewhat limited in that they are done on simulated data with no gain error or noise incorporated. It is obviously important to ascertain whether the methods put forward here are possible in more realistic scenarios, by testing them with simulations with gain errors and noise. If it seems that the methods in Chapter 4 are feasible for realistic scenarios, we should also try to apply them to real CHIME Pathfinder data. Additionally, the results from Chapter 5's comparison with holography were quite encouraging. Further work can include comparing with more holographic measurements and perhaps examining feed-to-feed beam variations via redundant baseline calibration at times when the 26 m telescope is not available and holographic measurements cannot be conducted.

Redundant baseline calibration and its offshoots will be fruitful areas of study not just for the moment, but into the future of CHIME and 21 centimeter hydrogen intensity mapping experiments at large.

Bibliography

[1] L. Abbott. The mystery of the cosmological constant. *Scientific American*, 258(5):82–88, 1988. → pages 2

[2] P. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. Banday, R. Barreiro, J. Bartlett, N. Bartolo, et al. Planck 2015 results-XIII. Cosmological parameters. *Astronomy and Astrophysics*, 594: A13, 2016. → pages 8

[3] A. Albrecht, G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox, et al. Report of the dark energy task force. *arXiv preprint astro-ph/0609591*, 2006. → pages 7, 8, 9, 10

[4] L. Anderson, É. Aubourg, S. Bailey, F. Beutler, V. Bhardwaj, M. Blanton, A. S. Bolton, J. Brinkmann, J. R. Brownstein, A. Burden, C.-H. Chuang, A. J. Cuesta, K. S. Dawson, D. J. Eisenstein, S. Escoffier, J. E. Gunn, H. Guo, S. Ho, K. Honscheid, C. Howlett, D. Kirkby, R. H. Lupton, M. Manera, C. Maraston, C. K. McBride, O. Mena, F. Montesano, R. C. Nichol, S. E. Nuza, M. D. Olmstead, N. Padmanabhan, N. Palanque-Delabrouille, J. Parejko, W. J. Percival, P. Petitjean, F. Prada, A. M. Price-Whelan, B. Reid, N. A. Roe, A. J. Ross, N. P. Ross, C. G. Sabiu, S. Saito, L. Samushia, A. G. Sánchez, D. J. Schlegel, D. P. Schneider, C. G. Scoccola, H.-J. Seo, R. A. Skibba, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, L. Verde, D. A. Wake, B. A. Weaver, D. H. Weinberg, M. White, X. Xu, C. Yèche, I. Zehavi, and G.-B. Zhao. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Releases 10 and 11 Galaxy samples. *Monthly Notices of the Royal Astronomical Society*, 441:24–62, June 2014. doi:10.1093/mnras/stu523. → pages 8

[5] K. Bandura, G. E. Addison, M. Amiri, J. R. Bond, D. Campbell-Wilson, L. Connor, J.-F. Cliche, G. Davis, M. Deng, N. Denman, et al. Canadian

hydrogen intensity mapping experiment (CHIME) pathfinder. In *SPIE Astronomical Telescopes+ Instrumentation*, pages 914522–914522. International Society for Optics and Photonics, 2014. → pages 15

[6] P. Berger, L. B. Newburgh, M. Amiri, K. Bandura, J.-F. Cliche, L. Connor, M. Deng, N. Denman, M. Dobbs, M. Fandino, A. J. Gilbert, D. Good, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, C. Hofer, A. M. Johnson, T. L. Landecker, K. W. Masui, J. Mena Parra, N. Oppermann, U.-L. Pen, J. B. Peterson, A. Recnik, T. Robishaw, J. R. Shaw, S. Siegel, K. Sigurdson, K. Smith, E. Storer, I. Tretyakov, K. Van Gassen, K. Vanderlinde, and D. Wiebe. Holographic beam mapping of the CHIME pathfinder array. *Proc. SPIE*, 9906:99060D–99060D–16, 2016. doi:10.1117/12.2233782. URL http://dx.doi.org/10.1117/12.2233782. → pages v, 70

[7] F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell, Q. Parker, W. Saunders, and F. Watson. The 6dF Galaxy Survey: baryon acoustic oscillations and the local Hubble constant. *Monthly Notices of the Royal Astronomical Society*, 416:3017–3032, Oct. 2011. doi:10.1111/j.1365-2966.2011.19250.x. → pages 8

[8] S. Carroll, W. Press, and E. Turner. The cosmological constant. *A.R.A&A*, 30:499–542, 1992. doi:10.1146/annurev.aa.30.090192.002435. → pages 3

[9] T.-C. Chang, U.-L. Pen, J. B. Peterson, and P. McDonald. Baryon acoustic oscillation intensity mapping of dark energy. *Physical Review Letters*, 100 (9):091303, 2008. → pages ix, 11, 12, 13

[10] T.-C. Chang, U.-L. Pen, K. Bandura, and J. B. Peterson. An intensity map of hydrogen 21-cm emission at redshift z [thinsp][ap][thinsp] 0.8. *Nature*, 466 (7305):463–465, 2010. → pages 14

[11] A. Conley, J. Guy, M. Sullivan, N. Regnault, P. Astier, C. Balland, S. Basa, R. G. Carlberg, D. Fouchez, D. Hardin, I. M. Hook, D. A. Howell, R. Pain, N. Palanque-Delabrouille, K. M. Perrett, C. J. Pritchet, J. Rich, V. Ruhlmann-Kleider, D. Balam, S. Baumont, R. S. Ellis, S. Fabbro, H. K. Fakhouri, N. Fourmanoit, S. González-Gaitán, M. L. Graham, M. J. Hudson, E. Hsiao, T. Kronborg, C. Lidman, A. M. Mourao, J. D. Neill, S. Perlmutter, P. Ripoche, N. Suzuki, and E. S. Walker. Supernova Constraints and Systematic Uncertainties from the First Three Years of the Supernova Legacy Survey. *ApJS*, 192:1, Jan. 2011. doi:10.1088/0067-0049/192/1/1. → pages 6

[12] S. Dodelson. *Modern Cosmology*. Elsevier Science, 2003. ISBN 9780080511979. URL https://books.google.ca/books?id=zk5e6XsY6nsC. → pages 4, 5, 6

[13] J. A. Frieman, B. Bassett, A. Becker, C. Choi, D. Cinabro, F. DeJongh, D. L. Depoy, B. Dilday, M. Doi, P. M. Garnavich, et al. The Sloan Digital Sky Survey-II Supernova Survey: Technical Summary. *The Astronomical Journal*, 135(1):338, 2007. → pages 6

[14] S. R. Furlanetto, S. P. Oh, and F. H. Briggs. Cosmology at low frequencies: The 21cm transition and the high-redshift universe. *Physics Reports*, 433(4): 181–301, 2006. → pages 10, 11

[15] J. Guy, M. Sullivan, A. Conley, N. Regnault, P. Astier, C. Balland, S. Basa, R. Carlberg, D. Fouchez, D. Hardin, et al. The Supernova Legacy Survey 3-year sample: Type Ia supernovae photometric distances and cosmological constraints. *Astronomy & Astrophysics*, 523:A7, 2010. → pages 6

[16] C. Haslam, U. Klein, C. Salter, H. Stoffel, W. Wilson, M. Cleary, D. Cooke, and P. Thomasson. A 408 MHz all-sky continuum survey. I-Observations at southern declinations and for the North Polar region. *Astronomy and Astrophysics*, 100:209–219, 1981. → pages 27

[17] E. A. Kazin, J. Koda, C. Blake, N. Padmanabhan, S. Brough, M. Colless, C. Contreras, W. Couch, S. Croom, D. J. Croton, T. M. Davis, M. J. Drinkwater, K. Forster, D. Gilbank, M. Gladders, K. Glazebrook, B. Jelliffe, R. J. Jurek, I.-h. Li, B. Madore, D. C. Martin, K. Pimbblet, G. B. Poole, M. Pracy, R. Sharp, E. Wisnioski, D. Woods, T. K. Wyder, and H. K. C. Yee. The WiggleZ Dark Energy Survey: improved distance measurements to z = 1 with reconstruction of the baryonic acoustic feature. *Monthly Notices of the Royal Astronomical Society*, 441:3524–3542, July 2014. doi:10.1093/mnras/stu778. → pages 8

[18] A. Liu and M. Tegmark. How well can we measure and understand foregrounds with 21-cm experiments? *Monthly Notices of the Royal Astronomical Society*, 419(4):3491–3504, 2012. → pages 11, 19

[19] A. Liu, M. Tegmark, S. Morrison, A. Lutomirski, and M. Zaldarriaga. Precision calibration of radio interferometers using redundant baselines. *Monthly Notices of the Royal Astronomical Society*, 408(2):1029–1050, 2010. doi:10.1111/j.1365-2966.2010.17174.x. URL http://mnras.oxfordjournals.org/content/408/2/1029.abstract. → pages 22

[20] C. J. Lonsdale, R. J. Cappallo, M. F. Morales, F. H. Briggs, L. Benkevitch, J. D. Bowman, J. D. Bunton, S. Burns, B. E. Corey, L. Desouza, S. S. Doeleman, M. Derome, A. Deshpande, M. R. Gopala, L. J. Greenhill, D. E. Herne, J. N. Hewitt, P. A. Kamini, J. C. Kasper, B. B. Kincaid, J. Kocz, E. Kowald, E. Kratzenberg, D. Kumar, M. J. Lynch, S. Madhavi, M. Matejek, D. A. Mitchell, E. Morgan, D. Oberoi, S. Ord, J. Pathikulangara, T. Prabu, A. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. U. Shankar, K. S. Srivani, J. Stevens, S. Tingay, A. Vaccarella, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, and C. Williams. The Murchison Widefield Array: Design Overview. *IEEE Proceedings*, 97:1497–1506, Aug. 2009. doi:10.1109/JPROC.2009.2017564. → pages 15

[21] L. B. Newburgh, G. E. Addison, M. Amiri, K. Bandura, J. R. Bond, L. Connor, J.-F. Cliche, G. Davis, M. Deng, N. Denman, M. Dobbs, M. Fandino, H. Fong, K. Gibbs, A. Gilbert, E. Griffin, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, C. Hofer, P. Klages, T. Landecker, K. Masui, J. M. Parra, U.-L. Pen, J. Peterson, A. Recnik, J. R. Shaw, K. Sigurdson, M. Sitwell, G. Smecher, R. Smegal, K. Vanderlinde, and D. Wiebe. Calibrating CHIME: a new radio interferometer to probe dark energy. *Proc. SPIE*, 9145:91454V–91454V–18, 2014. doi:10.1117/12.2056962. URL http://dx.doi.org/10.1117/12.2056962. → pages 17, 18, 19

[22] G. Paciga, J. G. Albert, K. Bandura, T.-C. Chang, Y. Gupta, C. Hirata, J. Odegova, U.-L. Pen, J. B. Peterson, J. Roy, J. R. Shaw, K. Sigurdson, and T. Voytek. A simulation-calibrated limit on the H I power spectrum from the GMRT Epoch of Reionization experiment. *Monthly Notices of the Royal Astronomical Society*, 433:639–647, July 2013. doi:10.1093/mnras/stt753. → pages 15

[23] A. R. Parsons, A. Liu, J. E. Aguirre, Z. S. Ali, R. F. Bradley, C. L. Carilli, D. R. DeBoer, M. R. Dexter, N. E. Gugliucci, D. C. Jacobs, P. Klima, D. H. E. MacMahon, J. R. Manley, D. F. Moore, J. C. Pober, I. I. Stefan, and W. P. Walbrugh. New Limits on 21 cm Epoch of Reionization from PAPER-32 Consistent with an X-Ray Heated Intergalactic Medium at z = 7.7. *ApJ*, 788:106, June 2014. doi:10.1088/0004-637X/788/2/106. → pages 15

[24] P. J. E. Peebles and J. T. Yu. Primeval Adiabatic Perturbation in an

Expanding Universe. *ApJ*, 162:815, Dec. 1970. doi:10.1086/150713. → pages 7

[25] U.-L. Pen, L. Staveley-Smith, J. B. Peterson, and T.-C. Chang. First detection of cosmic structure in the 21-cm intensity field. *Monthly Notices of the Royal Astronomical Society: Letters*, 394(1):L6–L10, 2009. → pages 14

[26] S. Perlmutter, G. Aldering, G. Goldhaber, R. Knop, P. Nugent, P. Castro, S. Deustua, S. Fabbro, A. Goobar, D. Groom, et al. Measurements of ω and λ from 42 high-redshift supernovae. *The Astrophysical Journal*, 517(2):565, 1999. → pages 6

[27] J. B. Peterson, K. Bandura, and U. L. Pen. The Hubble sphere hydrogen survey. *arXiv preprint astro-ph/0606104*, 2006. → pages 14

[28] J. R. Pritchard and A. Loeb. 21 cm cosmology in the 21st century. *Reports on Progress in Physics*, 75(8):086901, Aug. 2012. doi:10.1088/0034-4885/75/8/086901. → pages 10, 11

[29] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116(3):1009, 1998. → pages 6

[30] A. J. Ross, L. Samushia, C. Howlett, W. J. Percival, A. Burden, and M. Manera. The clustering of the SDSS DR7 main Galaxy sample - I. A 4 per cent distance measure at z = 0.15. *Monthly Notices of the Royal Astronomical Society*, 449:835–847, May 2015. doi:10.1093/mnras/stv154. → pages 8

[31] H.-J. Seo and D. J. Eisenstein. Probing dark energy with baryonic acoustic oscillations from future large galaxy redshift surveys. *The Astrophysical Journal*, 598(2):720, 2003. → pages 7, 8

[32] J. R. Shaw, K. Sigurdson, M. Sitwell, A. Stebbins, and U.-L. Pen. Coaxing cosmic 21 cm fluctuations from the polarized sky using *m*-mode analysis. *Phys. Rev. D*, 91:083514, Apr 2015. doi:10.1103/PhysRevD.91.083514. URL http://link.aps.org/doi/10.1103/PhysRevD.91.083514. → pages 17, 19, 21, 27, 29, 89

[33] M. Sitwell. *Models and Probes of the Early and Dark Universe: Inflation and 21-cm Radiation in Cosmology*. PhD thesis, University of British Columbia, 2014. → pages 25, 26, 27

[34] N. Suzuki, D. Rubin, C. Lidman, G. Aldering, R. Amanullah, K. Barbary, L. F. Barrientos, J. Botyanszki, M. Brodwin, N. Connolly, K. S. Dawson, A. Dey, M. Doi, M. Donahue, S. Deustua, P. Eisenhardt, E. Ellingson, L. Faccioli, V. Fadeyev, H. K. Fakhouri, A. S. Fruchter, D. G. Gilbank, M. D. Gladders, G. Goldhaber, A. H. Gonzalez, A. Goobar, A. Gude, T. Hattori, H. Hoekstra, E. Hsiao, X. Huang, Y. Ihara, M. J. Jee, D. Johnston, N. Kashikawa, B. Koester, K. Konishi, M. Kowalski, E. V. Linder, L. Lubin, J. Melbourne, J. Meyers, T. Morokuma, F. Munshi, C. Mullis, T. Oda, N. Panagia, S. Perlmutter, M. Postman, T. Pritchard, J. Rhodes, P. Ripoche, P. Rosati, D. J. Schlegel, A. Spadafora, S. A. Stanford, V. Stanishev, D. Stern, M. Strovink, N. Takanashi, K. Tokita, M. Wagner, L. Wang, N. Yasuda, H. K. C. Yee, and T. S. C. Project. The Hubble Space Telescope Cluster Supernova Survey. V. Improving the Dark-energy Constraints above $z > 1$ and Building an Early-type-hosted Supernova Sample. *The Astrophysical Journal*, 746(1):85, 2012. URL http://stacks.iop.org/0004-637X/746/i=1/a=85. → pages 6

[35] M. Tegmark and M. Zaldarriaga. Fast fourier transform telescope. *Physical Review D*, 79(8):083530, 2009. → pages 14

[36] M. Tegmark and M. Zaldarriaga. Omniscopes: Large area telescope arrays with only $n \log N$ computational cost. *Physical Review D*, 82(10):103501, 2010. → pages 14

[37] D. H. Weinberg, M. J. Mortonson, D. J. Eisenstein, C. Hirata, A. G. Riess, and E. Rozo. Observational probes of cosmic acceleration. *Physics Reports*, 530(2):87–255, 2013. → pages 6, 7, 8, 9, 10

[38] S. Weinberg. The cosmological constant problem. *Rev. Mod. Phys.*, 61: 1–23, Jan 1989. doi:10.1103/RevModPhys.61.1. URL http://link.aps.org/doi/10.1103/RevModPhys.61.1. → pages 2

[39] W. M. Wood-Vasey, G. Miknaitis, C. Stubbs, S. Jha, A. Riess, P. M. Garnavich, R. P. Kirshner, C. Aguilera, A. C. Becker, J. Blackman, et al. Observational constraints on the nature of dark energy: first cosmological results from the essence supernova survey. *The Astrophysical Journal*, 666 (2):694, 2007. → pages 6

[40] S. Yatawatta, A. G. de Bruyn, M. A. Brentjens, P. Labropoulos, V. N. Pandey, S. Kazemi, S. Zaroubi, L. V. E. Koopmans, A. R. Offringa, V. Jelić, O. Martinez Rubi, V. Veligatla, S. J. Wijnholds, W. N. Brouw, G. Bernardi, B. Ciardi, S. Daiboo, G. Harker, G. Mellema, J. Schaye, R. Thomas,

H. Vedantham, E. Chapman, F. B. Abdalla, A. Alexov, J. Anderson, I. M. Avruch, F. Batejat, M. E. Bell, M. R. Bell, M. Bentum, P. Best, A. Bonafede, J. Bregman, F. Breitling, R. H. van de Brink, J. W. Broderick, M. Brüggen, J. Conway, F. de Gasperin, E. de Geus, S. Duscha, H. Falcke, R. A. Fallows, C. Ferrari, W. Frieswijk, M. A. Garrett, J. M. Griessmeier, A. W. Gunst, T. E. Hassall, J. W. T. Hessels, M. Hoeft, M. Iacobelli, E. Juette, A. Karastergiou, V. I. Kondratiev, M. Kramer, M. Kuniyoshi, G. Kuper, J. van Leeuwen, P. Maat, G. Mann, J. P. McKean, M. Mevius, J. D. Mol, H. Munk, R. Nijboer, J. E. Noordam, M. J. Norden, E. Orru, H. Paas, M. Pandey-Pommier, R. Pizzo, A. G. Polatidis, W. Reich, H. J. A. Röttgering, J. Sluman, O. Smirnov, B. Stappers, M. Steinmetz, M. Tagger, Y. Tang, C. Tasse, S. ter Veen, R. Vermeulen, R. J. van Weeren, M. Wise, O. Wucknitz, and P. Zarka. Initial deep LOFAR observations of epoch of reionization windows. I. The north celestial pole. *Astronomy and Astrophysics*, 550:A136, Feb. 2013. doi:10.1051/0004-6361/201220874. → pages 15

[41] H. Zheng, M. Tegmark, V. Buza, J. S. Dillon, H. Gharibyan, J. Hickish, E. Kunz, A. Liu, J. Losh, A. Lutomirski, et al. MITEoR: a scalable interferometer for precision 21 cm cosmology. *Monthly Notices of the Royal Astronomical Society*, 445(2):1084–1103, 2014. → pages

[42] Q. Zheng, X.-P. Wu, J.-H. Gu, J. Wang, and H. Xu. A Method to Extract the Angular Power Spectrum of the Epoch of Reionization from Low-frequency Radio Interferometers. *APJ*, 758:L24, Oct. 2012. doi:10.1088/2041-8205/758/1/L24. → pages 15

Appendices

Appendix A

CHIME simulation pipeline

The CHIME simulation pipeline is composed of several git repositories (cora, caput, draco, driftscan, and ch-util), many of which are publicly available on github.

The progression through the pipeline outlined in Chapter 2 is demonstrated in Figure A.1.

A.1 Modification to pipeline code

A.1.1 Use of CylinderPerturbed Telescope Object

To create a telescope with perturbed beams, we must modify our starting point. The standard CHIME Pathfinder model as used in the previous chapter assumes a quasi-Gaussian beam which is identical for all inputs. (In other words, it lacks the additional term α used in the beam model presented in 4.1.1.) To incorporate beam perturbations, we use and later modify an existing class, PerturbedCylinder, written by Richard Shaw for the simulations discussed in [32].

The class begins with the same quasi-Gaussian beam used in the standard CHIME Pathfinder simulation, but adds the structure for a perturbation in the width of the beam, using the structure outlined in 4.1.1. PerturbedCylinder uses a finite difference method to generate the first derivatives necessary for the perturbed beam

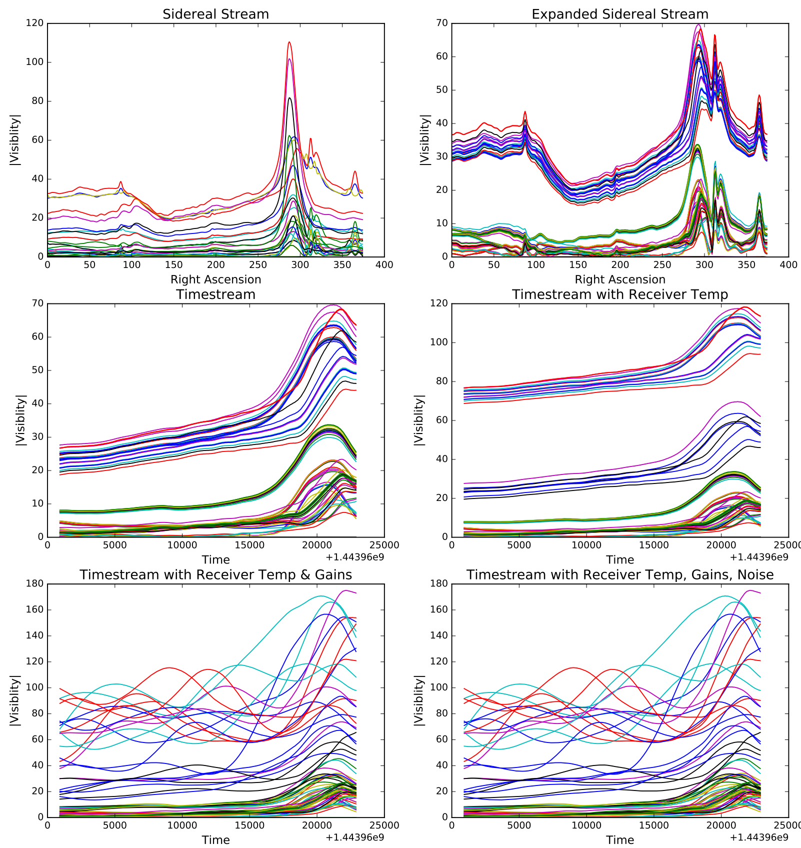

Figure A.1: In this figure, we see the progression of simulated data through the simulation pipeline. First, a sidereal stream is generated for each unique baseline. Then, the sidereal stream is expanded so that there is a representation of each individual baseline. At this stage of the pipeline process, these products do not incorporate complex gains or noise and are therefore perfectly redundant. Third, we expand from sidereal streams to individual time streams, with "20 second" samples, mimicking the actual CHIME data. Fourth, we add a constant receiver temperature to our timestream. Fifth, we add complex gains to the timestream, and finally ,we add sample noise.

structure. This class can also easily be modified to account for other perturbations.

A.1.2 ExpandPerturbedProducts

After creating a sidereal stream, the simulation pipeline generally executes the task "ExpandProducts," which expands the sidereal stream from having one instance of each unique product to having all redundant instances of that unique product. To incorporate our perturbed beam, we insert one entirely new task and also modify this task to both create all instances of a unique baseline and to combine the unperturbed component of the visibility with the two perturbed components.

We first generate the actual perturbation values α_i, which we disregarded in creating the initial, pre-expansion sidereal stream, using a new task GeneratePerturbations. This task, accompanied by the new container BeamPerturbations, uses the numpy standard normal random number generator to generate small, random multipliers which are scaled down to ensure this is in fact a small perturbation. The scaling factor can be input as a parameter when setting up the simulation or is set to a default value of 0.01. (This is a somewhat arbitrary, but seems to satisfy our goal of having a small perturbation.)

After generating perturbations, we turn our attention to modifying"ExpandProducts" to accommodate our perturbed simulation. We accomplish this by modifying the existing perturbation expansion code. The unperturbed "ExpandProducts" matches a product index p_i with the individual pairs of inputs f_i and f_j, then loops over each $[p_i,(f_i,f_j)]$ set. This is equivalent to forming the $A_i^{(0)}A_j^{*(0)}$ term in our perturbed simulation. For each product p_i, we follow a similar process to create the $A_i^{(0)}A_j^{(*1)}$ and $A_i^{(1)}A_j^{*(0)}$ terms. At this time, we also re-incorporate the previously calculated values, α, multiplying each product of beam functions by the appropriate α value. Finally, we add the three terms together (the zeroth order term as well as the first order terms in f_i and f_j.) Should we prefer to assume the perturbations are not very small, we can also add the double perturbed term here $\alpha_i\alpha_j\,A_i^{(1)}A_j^{*(1)}$; this term is normally neglected but the information for it is all present at this point.

The initial sidereal stream (pre-expansion) created with the perturbed beam appears to have a different number of inputs than the telescope actually has, $N_{\text{input}}(1+N_{\text{pert}})$ rather N_{input}. Therefore, in the expansion step, we must not only write out a new set of visibilities but must also write out a modified input map, which reflects the physical number of inputs in the telescope. (Failure to do this will confound later steps in the pipeline, such as adding random gains.)

Once we have output our new expanded stream, we are able to follow the simulation pipeline as in the unperturbed case, creating time streams and adding a receiver temperature, complex gains, and noise.

A.2 Example results

Results from a small scale example, with 16 total inputs, are shown in the figures below. This telescope size is of course less interesting for further applications of simulations with perturbed beam models, but it represents a good test set and is easier to compare between perturbed and unperturbed results.

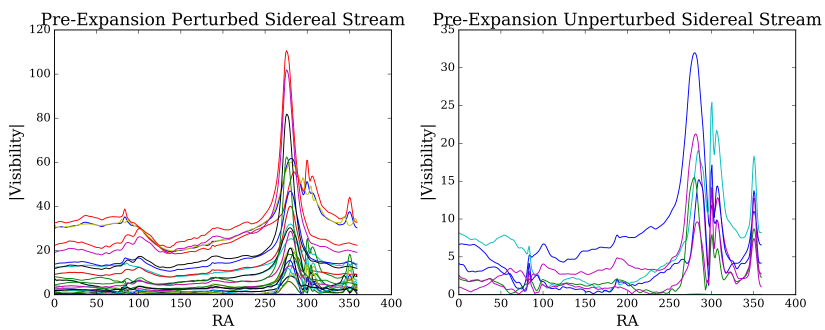

Figure A.2: Sample sidereal stream prior to expansion of redundant baselines and re-combination of perturbed components.

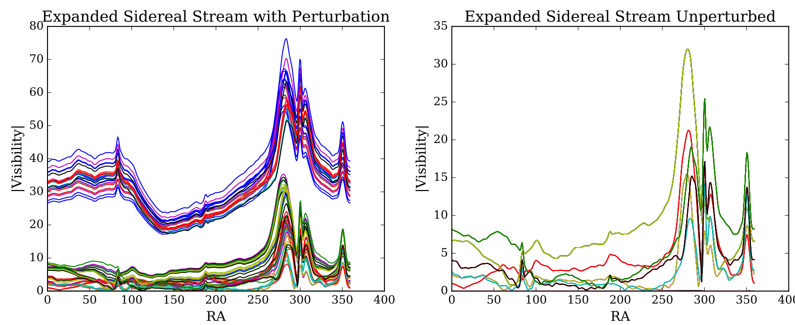

Figure A.3: The left hand panel shows sample sidereal stream after to expansion of redundant baselines, for an identical, unperturbed simulation. The right hand panel shows a sample perturbed sidereal stream after the expansion of redundant baselines and re-combination of perturbed components. There are now as many different baselines as real CHIME Pathfinder data and redundancy has been broken by the addition of per-feed beam perturbations. However, at this stage, the simulated data does not include instrumental gains or any noise estimate and is therefore not generally suitable for analysis tasks. In this particular figure, the perturbation is turned up to approximately 0.1 (from approximately 0.01) to make its existence more obvious. Without the perturbation, the two panels would be identical, as the underlying telescope configuration is the same as is the input sky map.

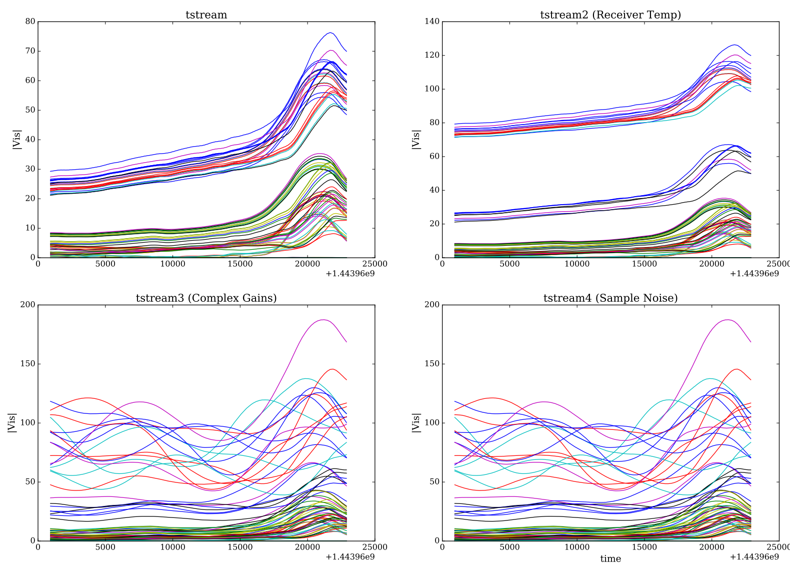

Figure A.4: After the product expansion stage, the simulation pipeline should proceed as in a standard unperturbed version. First, we add the receiver temperature, then time dependent complex gains, then we add sample noise. After adding the complex gains and sample noise, would-be redundant baselines are no longer redundant in either the perturbed or unperturbed case, but do resemble raw CHIME Pathfinder data

Appendix B

Supplemental Figures

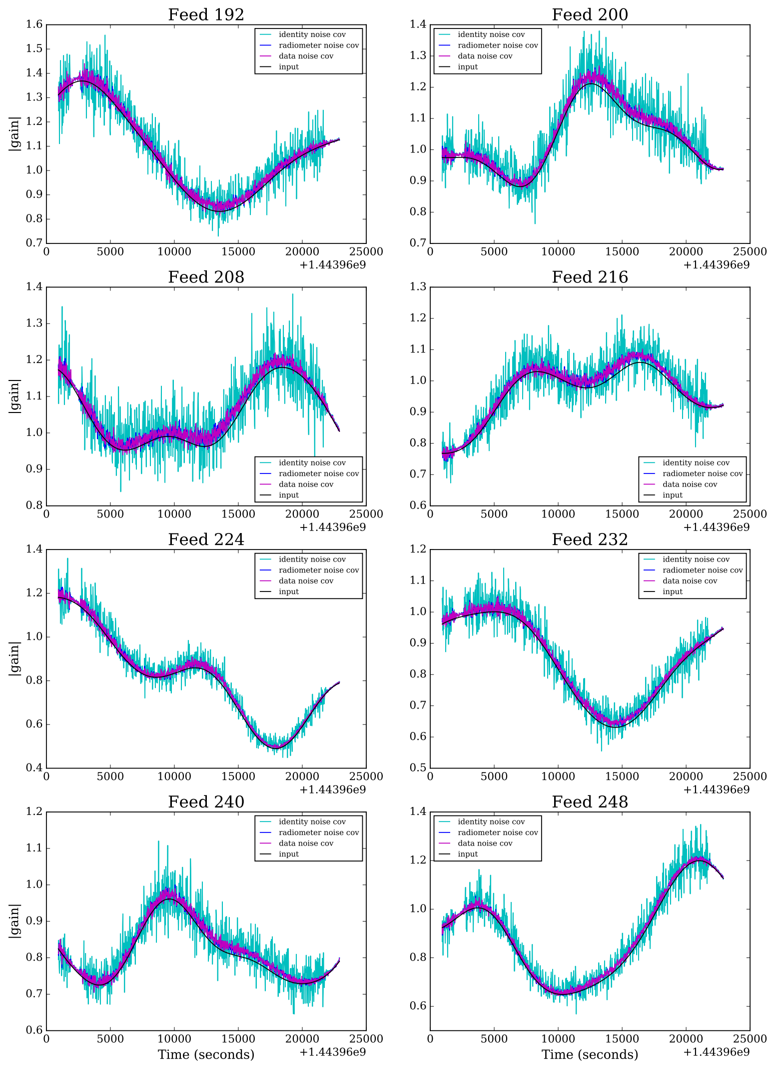

Figure B.1: From Chapter 2. In this figure, we compare the percent deviation from input gains for the identity noise covariance, data-derived noise covariance, and radiometer noise covariance for a sampling of inputs on the east cylinder.

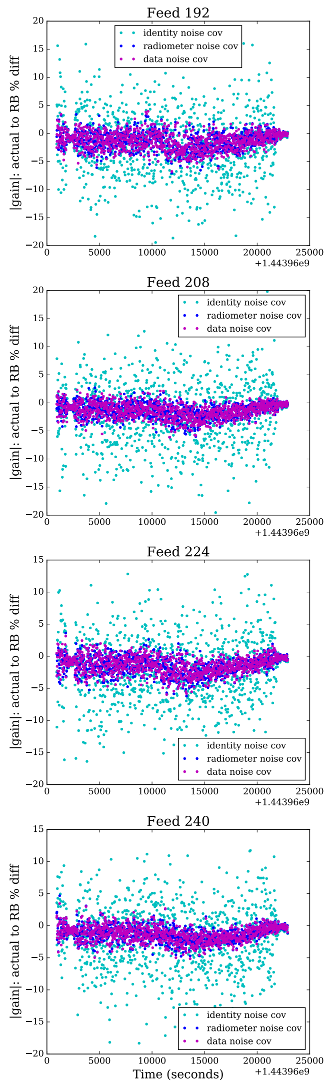

Figure B.2: From Chapter 2. In this figure, we compare the absolute deviation from input gains for the identity noise covariance, data-derived noise covariance, and radiometer noise covariance for a sampling of inputs on the east cylidner.

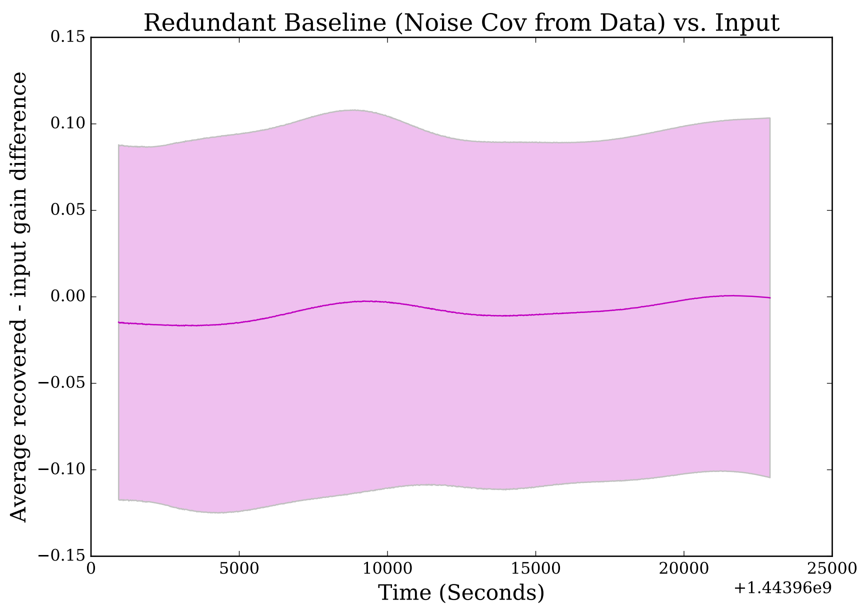

Figure B.3: From Chapter 2. Here, we view the average deviation between the input and recovered gain. The pink band represents one standard deviation above and one standard deviation below the average.

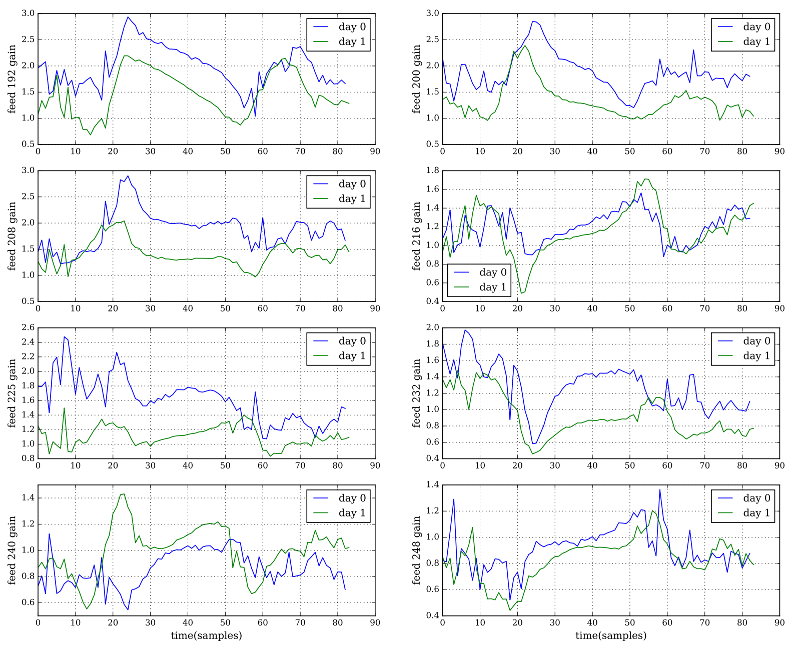

Figure B.4: From Chapter 3. We compare redundant baseline gain results for selected feeds on the east cylinder calculated over two consecutive CygA transits. We notice that there is very little deviation between the two days and regard this as evidence that the cause of the slope in the gain solution is not strongly time dependent and is likely a property of the array.

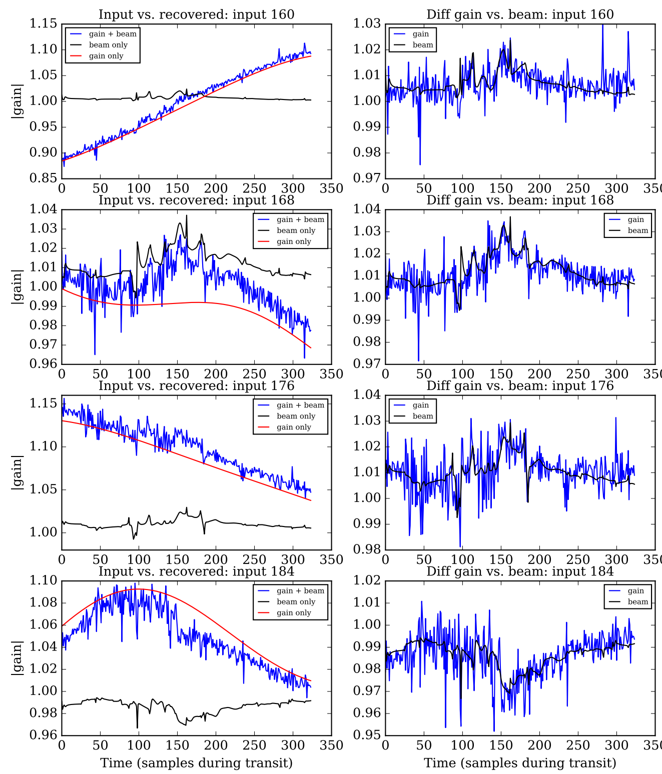

Figure B.5: From Chapter 4. A continuation of Figure 4.1

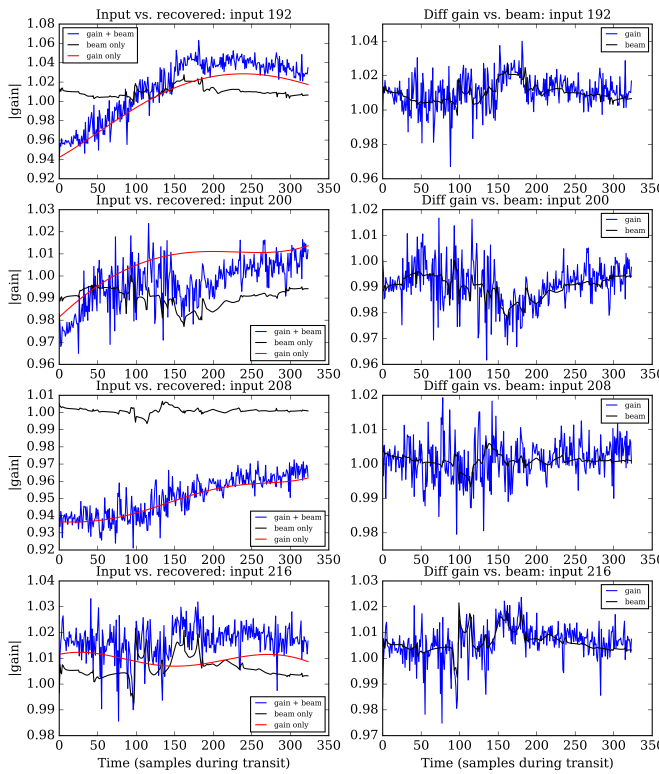

Figure B.6: From Chapter 4. A continuation of Figure 4.1

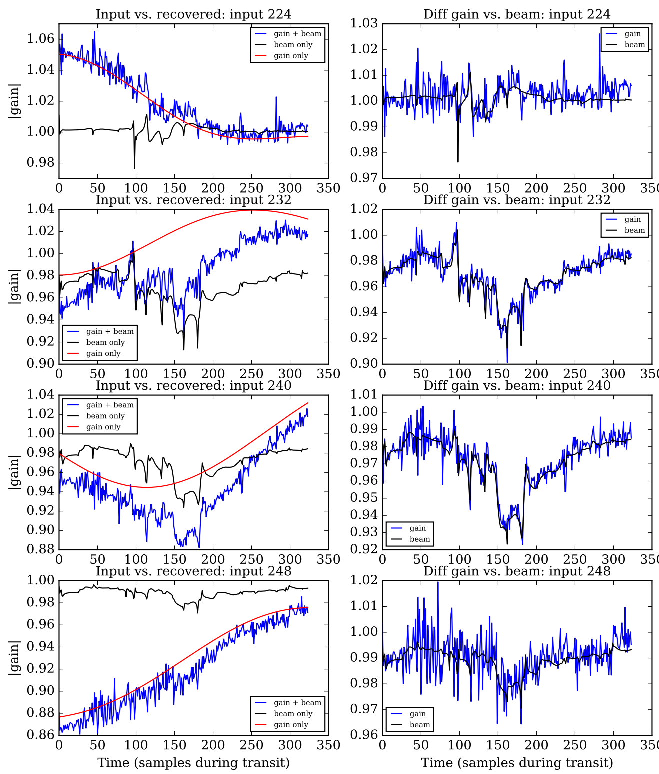

Figure B.7: From Chapter 4. A continuation of Figure 4.1

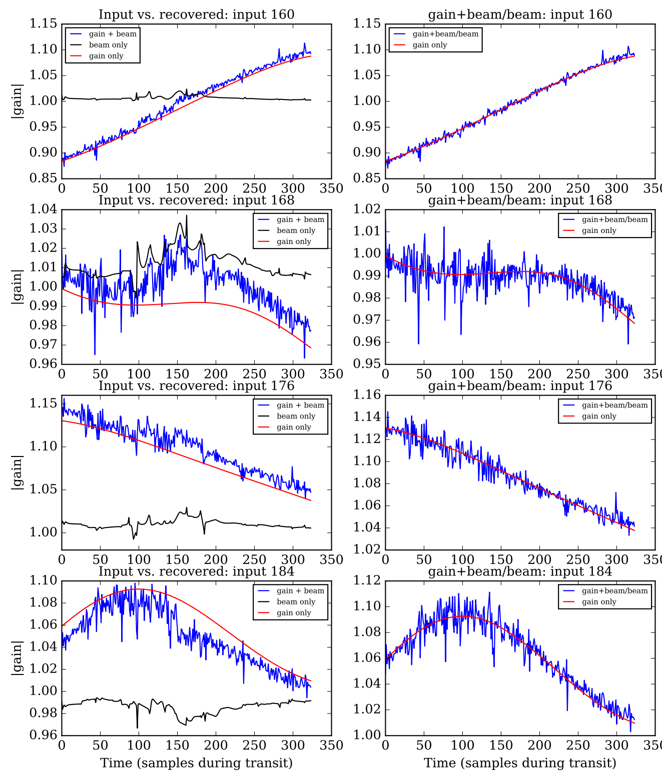

Figure B.8: From Chapter 4. A continuation of Figure 4.2

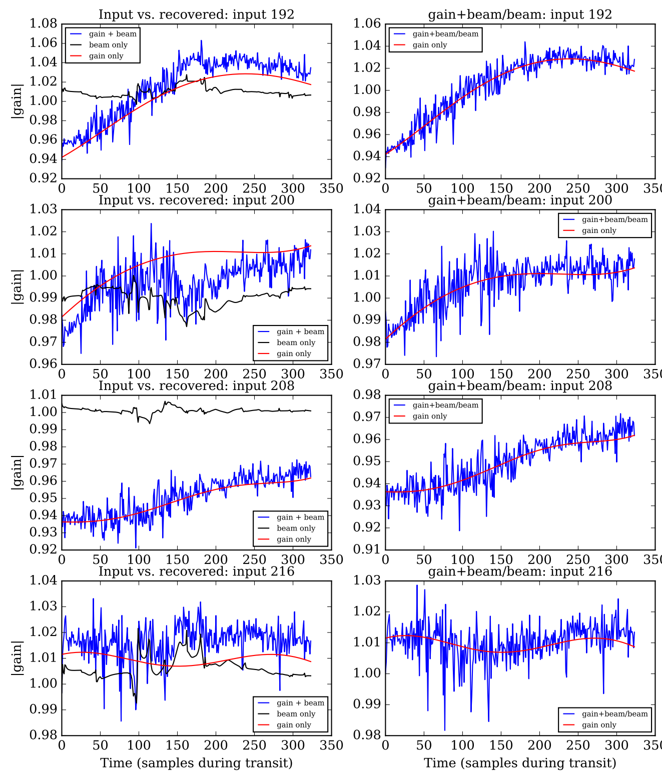

Figure B.9: From Chapter 4. A continuation of Figure 4.2

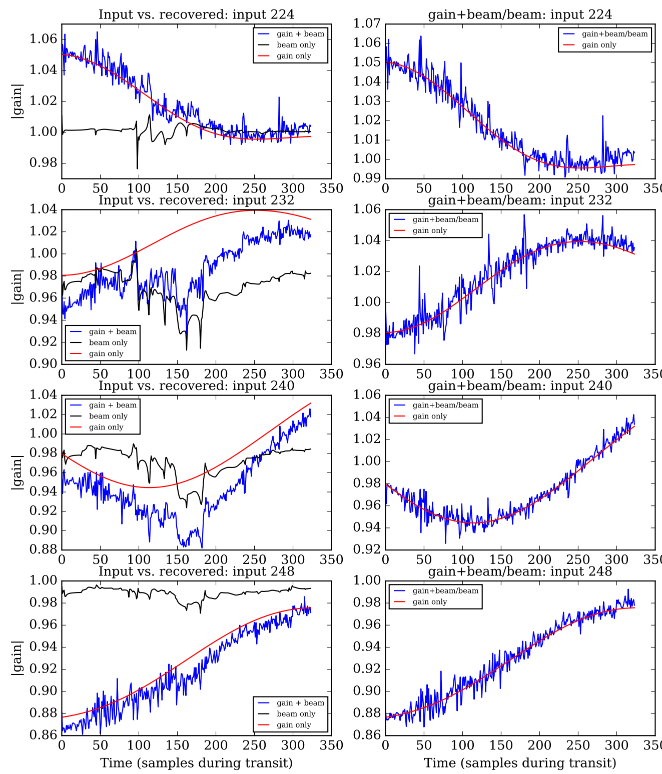

Figure B.10: From Chapter 4. A continuation of Figure 4.2

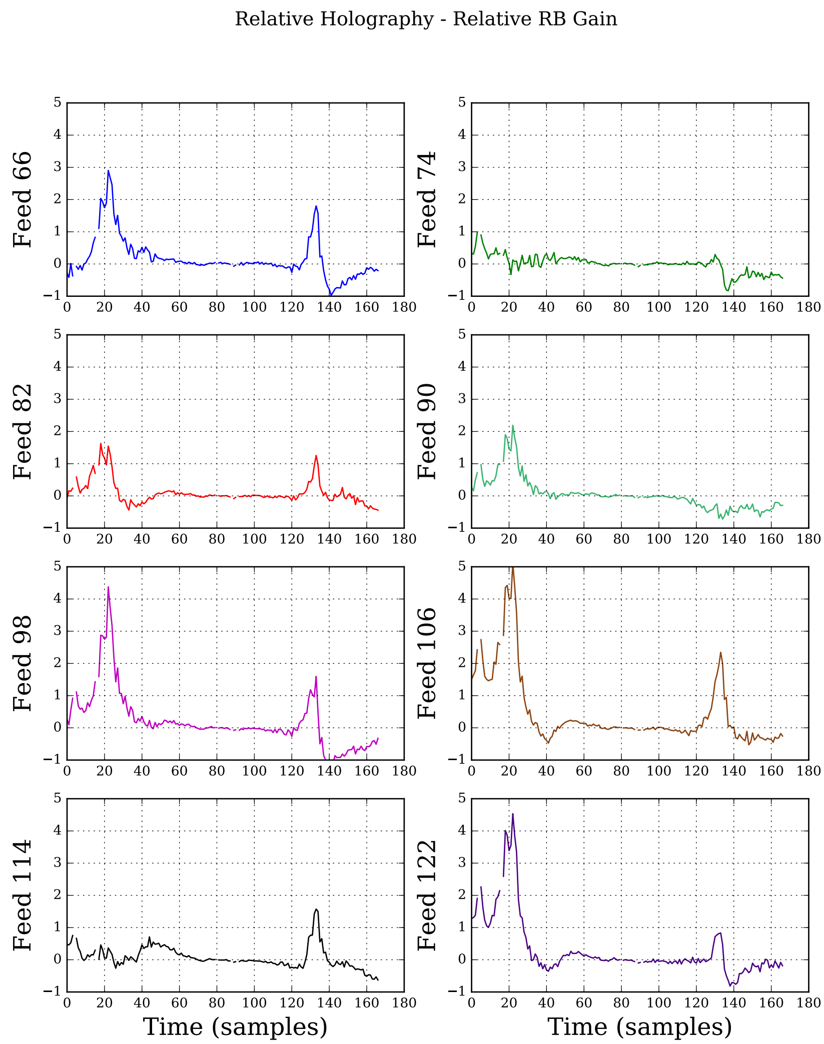

Figure B.11: From Chapter 5. Differences between redundant baseline gain ratios and ChIME-26 m cross correlation ratios; indicative of difference between beam estimates for a given feed on the west cylinder.

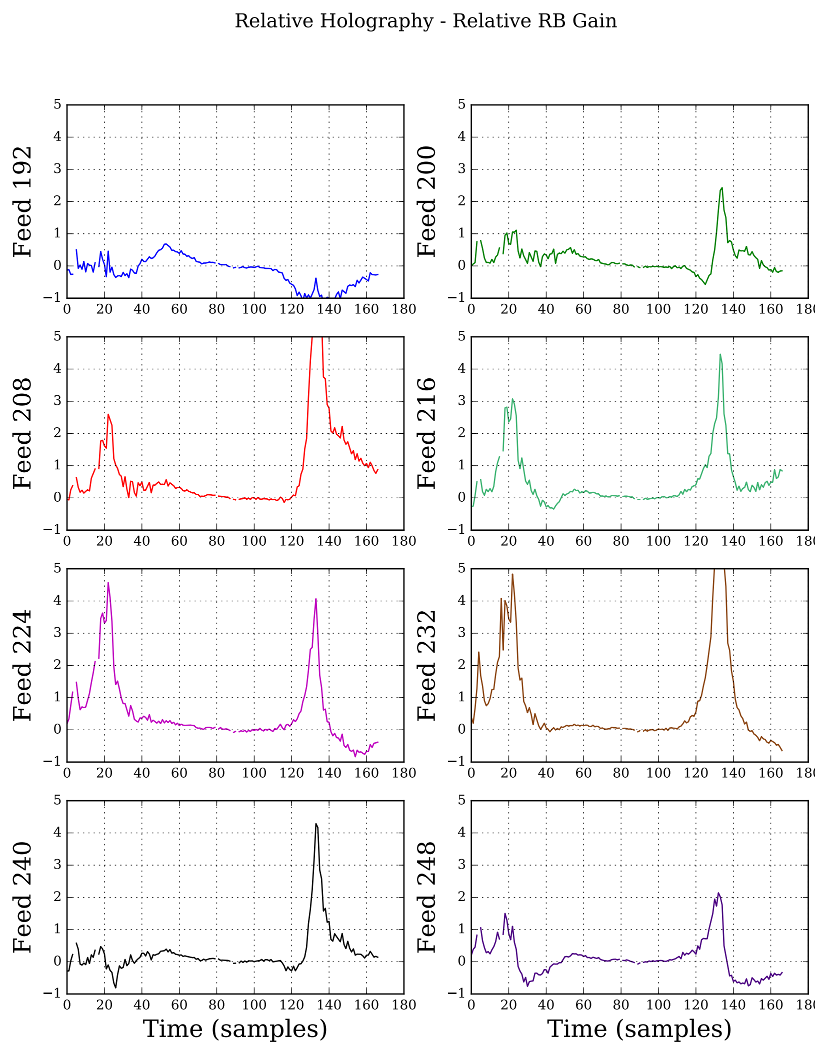

Figure B.12: From Chapter 5. Differences between redundant baseline gain ratios and CHIME-26 m cross correlation ratios; indicative of difference between beam estimates for a given feed on the east cylinder.