Improved Methods for Phased Array Feed Beamforming in Single Dish Radio Astronomy

Michael J. Elmer

A dissertation submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Brian D. Jeffs, Chair
Karl F. Warnick
David G. Long
Michael D. Rice
Wynn C. Stirling

Department of Electrical and Computer Engineering

Brigham Young University

August 2012

ABSTRACT

Improved Methods for Phased Array Feed Beamforming in Single Dish Radio Astronomy

Michael J. Elmer
Department of Electrical and Computer Engineering
Doctor of Philosophy

Among the research topics needing to be addressed to further the development of phased array feeds (PAFs) for radio astronomical use are challenges associated with calibration, beamforming, and imaging for single dish observations. This dissertation addresses these concerns by providing analysis and solutions that provide a clearer understanding of the effort required to implement PAFs for complex scientific research. It is shown that calibration data are relatively stable over a period of five days and may still be adequate after 70 days. A calibration update system is presented with the potential to refresh old calibrators. Direction-dependent variations have a much greater affect on calibration stability than temporal variations.

There is an inherent trade-off in beamformer design between achieving high sensitivity and maintaining beam pattern stability. A hybrid beamformer design is introduced which uses a numerical optimizer to balance the trade-off between these two conflicting goals to provide the greatest sensitivity for a desired amount of pattern control. Relative beam variations that occur when electronically steering beams in the field of view must be reduced in order for a PAF to be useful for source detection and imaging. A dual constraint beamformer is presented that has the ability to simultaneously achieve a uniform main beam gain and specified noise response across all beams. This alone does not reduce the beam variations but it eliminates one aspect of the problem. Incorporating spillover noise control through the use of rim calibrators is shown to reduce the variations between beams. Combining the dual constraint and rim constraint beamformers offers a beamforming option that provides both of these benefits.

Keywords: Arecibo telescope, array signal processing, beamformer design, beamforming, beam pattern, calibration stability, electronic drift, noise fields, phased array feeds, radio astronomy, receiver design, weak source imaging.

ACKNOWLEDGMENTS

It is a great pleasure to take this opportunity to formally thank those who have helped me with writing this dissertation. I am sincerely grateful to my advisor Dr. Brian Jeffs for believing in my potential from the beginning, for patiently guiding me toward the results I sought, and for offering many words of encouragement along the way. Thanks also to Dr. Karl Warnick for providing feedback, direction, and helpful insights that kept me headed in the right direction. I had the privilege of working with many wonderful individuals at BYU and I am grateful for the associations and friendships that have been developed over the last few years. My family has been a wonderful support throughout this process and I am grateful for their interest and influence in my life. I would like to specifically thank my grandmother Marilin, for believing in me and always making me feel like her favorite; my in-laws Hal and Connie, for selflessly giving of their time and energy; and my parents Chris and Diane, for their examples of love, kindness, sacrifice and appreciation, and for providing a way for this dream to become a reality. Finally, I would like to express my deepest appreciation to my wife Maria, who graciously agreed to take this journey with me and who has offered unwavering support and encouragement throughout; and to our wonderful children: Anna, Joe, Crista, Luke, Eliza, and Tessa, for their enthusiasm for life and learning.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Since the 1930s when Karl Jansky first detected radio signals from space and the field of radio astronomy was born, there has been a great interest among astronomers and other members of the scientific community to observe the invisible radio universe [1–3]. Through the use of a radio telescope, which typically consists of a large parabolic reflector and a horn feed, single pointing observations or full images are generated to reveal radio emissions from celestial sources throughout the universe. Such observations have helped to explain the interactions of celestial bodies and provide insight into the origins of the universe [4–6].

A current active area of radio astronomy research focuses on developing phased array feeds (PAFs) for use with a radio telescope. A PAF consists of a number of closely spaced antenna elements positioned in the focal plane of the telescope, with overlapping response patterns which can be combined to electronically steer a beam throughout a larger field of view (FOV) than is available for a single dish pointing with a fixed horn feed. In addition to the increased sky survey speed resulting from the expanded FOV, other potential benefits of such a system include active control of the beam pattern shape, nulling of interfering signals, and adaptation to changes in the observed noise field. These benefits, however, can only be realized after resolving a number of technical challenges, among which are the need for accurate calibration data, development of adequate beamforming techniques, and overcoming challenges associated with imaging the sky with electronically steered beams. These specific topics are addressed in this dissertation, which presents a study of long-term calibration stability, introduces a new beamformer designed for both fixed pattern control and high sensitivity, and develops methods for resolving the relative beam variations that inhibit weak source detection and imaging.

1.1 Progression of Radio Astronomy Instrumentation Development

Single dish radio astronomy, as its name suggests and in contrast to a synthesis array [7–12], uses a single radio telescope to observe the sky. Operating a single telescope offers a considerable reduction in complexity over radio synthesis imaging and is useful for spectroscopy, observing pulsars, hydrogen emissions, redshift, and interstellar molecules, and for fast surveying of a region of the sky [13–17]. In contrast to the typical single horn feed of a radio telescope, focal plane arrays (FPA), also called multibeam receivers, consist of several horn feeds placed near the focus of the dish allowing for a decrease in required observation time proportional to the number of antennas in the array. Instead of projecting a single beam to a single point in the sky, multiple independent, non-overlapping beams are projected to multiple individual look directions simultaneously. Several FPA systems have been commissioned, including the NAIC ALFA array at Arecibo and the Australian (ATNF/CSIRO) Parkes 21 cm multibeam receiver [18–20].

With the hope of reducing single dish sky survey time further, significant effort is currently being devoted to the development of phased array feeds (PAFs) for radio astronomical use. Moving toward the use of an array of closely spaced antennas in the focal plane provides several potential benefits, the most important of which is the ability to increase the instantaneous field of view by forming multiple simultaneously steered beams. PAFs can also form dense grids of optimally spaced overlapping beams, and improve interference mitigation. Contributors to this worldwide effort include ASTRON in the Netherlands, DRAO in Canada, NAOC in China, CSIRO in Australia, NAIC at Arecibo, Puerto Rico and a collaboration between the Radio Astronomy Systems Research (RASR) group at Brigham Young University (BYU) and the National Radio Astronomy Observatory (NRAO) in the USA [21–30]. Although a number of prototypes and partial systems are currently deployed, including the WSRT [31] and the Australian SKA Pathfinder (ASKAP) project [25, 32], no complete PAF-based system has yet to be fully commissioned for scientific observations. There is still a large ongoing research and development effort working toward this goal.

With this significant transition from traditional single-pixel radio telescopes to PAF systems, it is apparent that promised performance improvements can be achieved only with increased system complexity and after resolving a number of technical challenges. Challenges include mutual coupling among the closely spaced array elements [33–35]; hardware requirements for many (e.g., 40 to 200) identical analog and digital receiver chains and correlator/beamforming processing; and the need for new high performance beamformer design methodologies suitable for the PAF environment.

1.2 Problem Statement: Radio Astronomical Phased Array Feed Development

Before PAFs are ready for scientific observations there are several key topics that need to be investigated and understood. Among them are the topics of calibration methods, beamforming techniques, and imaging capabilities, which are addressed in this dissertation. While the development of these topics is mainly concerned with their application to single dish radio astronomy, the results can be easily extended to synthesis arrays that incorporate PAFs, such as the ASKAP array.

Obtaining accurate calibration data is a necessary first step in working with a PAF. In radio astronomy the word “calibration” has several different meanings. For the context of this dissertation, the term refers to array calibration, which means to measure the relative responses between each array element for a far-field point source at different angles of arrival. Because of the large size of radio telescopes and the dynamic noise environment in which they operate, calibration data must be obtained with the PAF mounted on the telescope using an available bright astronomical source. Electronic gain drift and other time-dependent sources of noise contribute to degradation of the calibrators over time. It has been suggested that these variations may be accounted for with the use of a calibration update system, such as the one which is presented in Chapter 3, but it is shown that the performance of the calibrators is much more dependent on telescope pointing direction than temporal drift. Until the directionally dependent variations are addressed there is no need to correct for temporal drift since it appears to be negligible. It is shown that PAF sensitivity and beam pattern structure experience very little change, relative to other variations, over periods of days and weeks.

A beamformer that provides a high signal to noise ratio (SNR) seems ideal for PAF use. The obvious beamforming choice to provide this ability is the max-SNR (maximum sensitivity) beamformer (described in Chapter 2), which by definition offers the highest

SNR for a given noise field. It has the ability to minimize noise from all sources collectively (including spillover, main beam sky, receiver, and antenna noise) and detect the weak sources of interest. The disadvantage of this beamformer is that it has an unpredictable beam pattern that varies its response with changes in the noise environment. This unpredictability may cause inaccuracies when observing weak sources in the presence of other much stronger sources. Unknown high side lobe levels may be present which is generally a poor design feature. A deterministic beamformer is required to fully control the beam pattern shape. However, while a deterministic beam pattern design offers the desired known structure, it may lack the ability to provide the necessary sensitivity for radio astronomical observations. In this dissertation a hybrid approach is presented that strives to balance the tradeoff between these two extreme approaches to offer a beamformer that gives the greatest amount of sensitivity for a specified amount of pattern control. In Chapter 4 this method is shown to be a viable option for PAF beamformer design.

The ability of a PAF to increase sky survey/imaging speed is only beneficial if the amount of additional noise introduced by steering different beams is minimized, or at least tolerable. Each beam typically provides a different noise response which appears as variations in the field of view that may obscure a weak source in the image. The variations are reduced by normalization techniques but relative beam power measurements still differ from their actual values, since noise field normalization introduces relative differences in the gains of each beam. These differences, whether present in the noise response or the main beam gain, present a challenge when imaging a wide-field astronomical source since the pixel-to-pixel power values may be inaccurately scaled. To remedy this problem, a dual constraint beamformer has been designed, which has the ability to meet both a specified main beam gain and a desired noise response to eliminate this level of variation (see Chapter 5). Other sources of noise variation arise from differences in dish illumination spillover patterns between beams. As the radio telescope changes pointing direction, each beam observes different structures in the spillover noise field, resulting in increased variability. Introduction of a rim constraint beamformer, which incorporates eigenvector constraints based on rim calibrator data, is shown to provide a reduction in beam-to-beam variations as it constrains the array illumination pattern to equally reduce the amount of spillover noise in each beam. Further, a combination of these two approaches is shown to offer a beamformer that has reduced beam variations while meeting the gain and noise response constraints.

1.3 Related Work

Phased antenna aperture arrays (i.e., bare arrays with no large reflector) have been used for 70 years in applications of wireless communication, radar, sonar, and remote sensing, providing benefits such as improved direction finding, spatial interference canceling, rapid beam steering, forming multiple simultaneous beams, gain optimization, etc. [36–39]. Placing a phased array at the focal point of a reflector dish is a relatively new concept that also has application beyond the field of radio astronomy. For example, recent work in the area of satellite communications has included discussion on space platform orbital PAF-fed reflectors to provide adaptation to a changing radio environment [40,41], similar to the application in radio astronomy, but on a smaller scale. Though phased array antenna theory and beamformer design are relatively mature fields, the stringent demands of radio astronomical observation, including observing at very low SNRs and the need for extreme pattern and gain stability, have until very recently kept these techniques from being used in feed designs for the large dish instruments. This section contains a review of some of the relevant PAF work being done at BYU and by other radio astronomy PAF research groups, as well as related work for other array antenna applications.

Calibration Stability

Stable calibrators are required to compute accurate PAF beamformers for measured data, but temporal variations in the noise environment introduce a source of calibration drift. The main contributor to this drift is electronic gain variations and a proposed solution suggests placing calibrator reference antennas at the vertex and around the rim of the reflector dish [22, 42]. A single reference signal can be used to correct electronic gain drift. Multiple reference signals may be used to identify and correct additional sources of drift also. Electronic gain drift and its potential affect on PAF beamformers is discussed in [43].

The use of a reference antenna for electronic gain drift calibration correction is not unique to radio astronomy. Several reference horns were placed around the radar array described in [44] in order to compute both a thorough initial calibration and later to correct for electronic drift to update the calibration data. Corrections were shown to be beneficial for electronic component gain drift and for identification of a failed component. In [45] a reference signal is used to update calibration measurements of a general-use phased array antenna system using a maximum likelihood estimate to identify the change in response of each array element. By continually tracking the changes the time required to recalibrate the array was significantly decreased.

To the present day, calibration update systems and algorithms for PAFs have been relatively limited in complexity. In [46], a noise source was mounted at the vertex of the radio telescope. It served as a reference to correct for relative drifts in electronic gains with temperature and time, or even component failure. A noise signal was turned on and measurements were coupled with the signal and fed into a correlator in order to compute a phase reference from which to measure the drift. The purpose of the study was to demonstrate PAF beamforming on a prototype platform and not to analyze the gain correction system. For this reason very little information was given concerning the gain correction component.

In [47] a single reference noise source was also placed at the vertex of a PAF-fed reflector dish and was switched on periodically to measure the short-term variations between receiver channels as a function of time. Throughout a three-hour observation the relative receiver channel gain variations were compensated for but the results showed that this had very little effect on the beam patterns because pattern variations were dominated by mechanical pointing errors.

For the dual polarization studies of [48] a single horn calibrator was placed at the vertex of the reflector dish. The calibrator was fed by a CW signal generator. It is suggested that the measurement be used to update beamformers in response to changes in the amplitude and phase of the receiver gains. The authors suggest that the time between such updates should be on the order of hours. Analysis was performed using simulated perturbations of the real data and there is no evidence that true receiver gain drift was measured and corrected.

This dissertation includes the first known study of long-term calibration stability, presented in Chapter 3. A complete calibration update system including vertex and rim-mounted calibrator sources is employed to measure and correct for calibrator drift. It is shown that calibration data are surprisingly stable over periods of weeks and even months. Additionally, correction of electronic drift is found to be unnecessary at this time since larger, directionally dependent variations are observed to dominate the calibration drift.

Beamformer Design

Developments in applied PAF beamformer design have been relatively limited, although there have been a number of promising results based on simulation models. To date, the data-dependent max-SNR beamformer [49], which was simultaneously introduced for astronomical PAF use by the ASTRON and BYU-NRAO teams, has been the only one successfully applied to create images of experimental PAF data [28, 50–55]. Conjugate field match beamforming has also been attempted, but has been found to be unsuitable due to the inability to control beam pattern shape when there is significant gain variation across the sensor array for far-field sources [56–59].

The conjugate field match, max-SNR, and linearly constrained minimum variance (LCMV) beamformers were jointly considered for PAF use in [60]. Conjugate field match was shown to perform very poorly since it is designed to maximize received power in a given direction without regard to noise. The max-SNR beamformer is shown to provide the greatest sensitivity but the inability to control the beam pattern is emphasized and for this reason the LCMV beamformer is presented as the optimal choice. The LCMV beamformer provides the ability to constrain the beam pattern at various points, but there are drawbacks to this approach and an improved method is introduced in Chapter 4 of this dissertation.

In [48] beamforming for a dual polarized PAF is considered. A beamformer is calculated for each polarization as the eigenvectors corresponding to the two largest eigenvalues of the signal covariance matrix. This was the conjugate field match method and was likely used for convenience. The purpose of the study was to show the benefit of a dual polarized PAF over a single polarized version. The paper did not specifically focus on beamformer design techniques.

Using a simulated PAF, eigenbeams have been effective in reducing data transfer and storage requirements [61]. A set of basis weight vectors is obtained from the eigenvectors corresponding to a subset of the largest eigenvalues of the signal correlation matrix. These orthogonal basis weight vectors are shown to contain most of the information needed to compute the desired beams. Another simulated result consists of numerically optimized Gaussian beams that are steered around the FOV without distortion, while accounting for polarization effects [62]. The capabilities of the simulation model are highlighted and the beam patterns are impressive, but real data experimental results have not been reported.

Broadband PAF beamforming is considered in [63] for the CSIRO ASKAP array. There is no discussion about specific beamformers that are to be used, but there is a general discussion about plans to implement a beamformer. A separate beamformer is to be calculated for every 1 MHz band throughout the passband.

In [64] there is a comparison of single-ended and differential beamforming approaches. No specific beamforming method is presented. The focus of the article is to examine the computational savings associated with a differential beamformer and the associated reduction in performance.

The results of Chapter 4 show the difficulty in applying simulation model beamformers to real measured data, where much of the variability in the FOV is random and cannot be accurately modeled. The need for beamformers that provide both high sensitivity and beam pattern shape control is shown to be satisfied by a new hybrid beamforming approach. The hybrid beamformer offers the greatest sensitivity for a given amount of beam pattern control. Only narrowband, single polarized PAF beamformer design is considered.

Beam Variations in Single Dish Imaging

Several groups have produced single dish radio astronomical PAF images of relatively bright sources but it is unclear what normalization procedures have been used [28,53,54,65]. The results in [28,53] show some artifacts of beam-to-beam variations. In these publications some mosaic images are presented combining the imaged FOVs for several pointings of the telescope. Hard transition lines are observed between the separate FOVs because of the variations between the beams.

In [43] simulated electronic receiver drift is shown to cause distortion in observed images. The issue of dynamic range is considered, where the ability to image weak sources is inhibited by the presence of stronger sources or beam variations due to electronic gain drift. Electronic gain drift is discussed in Chapter 3 of this dissertation but it is found to be relatively minimal and is not considered in the context of imaging.

Each of the beams of a multibeam receiver forms a different far-field pattern on the sky, with different gains and noise responses, which must be accounted for when processing the data and creating images. In [66] the differing gains in each pixel of an image from the ALFA array are smoothed by making several passes of the same part of the sky after adding a slight offset to the array pointing. Averaging multiple overlapping observations is also discussed in [67], where data from the Parkes 21-cm multibeam array is used to conduct a neutral hydrogen sky survey. While the processing of multibeam receiver data is quite different from that of a PAF, these examples confirm the need to account for differences in beam gains in order to decrease the beam-to-beam variations that appear in PAF images. Performing multiple passes over the same region of sky is a possibility for PAFs but this is undesirable since it increases survey time.

The issue of relative beam differences is addressed in Chapter 5, where several techniques are discussed for decreasing the variations. A dual constraint beamformer is introduced, which offers uniform main beam gain and noise response across all the beams in the FOV to eliminate one aspect of the variability. Additionally, an LCMV beamformer with eigenvector constraints derived from rim calibrator source data is shown to effectively control response to spillover noise. These techniques are shown to improve PAF imaging capabilities.

1.4 Summary of Contributions

The following list summarizes the contributions made by the research presented in this dissertation.

- Developed a model for elevation-dependent atmospheric noise that was incorporated into the modeling of spillover and main beam sky noise.
- Introduced eigenvalue scaling of unit normalized response vectors to maintain relative gain differences between directionally-dependent beam responses.

- Instituted a calibration vector quality metric based on the minimum description length principle, which examines the eigenvalues of a pre-whitened covariance matrix to detect the presence of the calibrator source.

- Developed a hybrid beamforming approach to balance the demands for high sensitivity with a known and stable beam pattern structure.

- Demonstrated that fixed adaptive beamforming performs nearly as well as fully adaptive beamforming in response to elevation-dependent noise variations.

- Studied the effects of elevation-dependent noise using electromagnetic models and experimental data. Analyzed the contributions of the various noise sources and the response of an adaptive beamformer to changes in the noise environment. The results led to the practice of obtaining noise measurements at the same elevation as the source of interest.

- Played a lead role in planning, preparing, and executing a large-scale experiment on the world's largest radio telescope in Arecibo, Puerto Rico. Worked with members of the RASR group and the NAIC team in Arecibo to accomplish the desired experimental goals while introducing a new array, new receiver system, and new data acquisition system at an unfamiliar test location.

- Assisted with the design and development of the first PCB version of the PAF analog receiver, which was largely based on the original plug-in component version. Then designed an upgraded version of the board to work with the new FPGA-based system with higher sample rate and larger bandwidth requirements.

- Analyzed data collected at the anechoic chamber at the University of Utah to determine the temporal stability of the electronic receiver and data acquisition systems.

- Analyzed the long-term calibration stability of PAFs with regard to sensitivity and beam pattern degradation. Found that calibration data remains fairly stable over a significant period of time and that directionally-dependent variations dominate. Variations due to calibrator drift are relatively small in comparison.

- Constructed a reflector rim calibration system to update calibrators over time. The system proved to be unnecessary for this purpose but was later used to provide data for the rim constraint beamformer. The prototype system will serve as the basis for a larger scale version that will be used to increase control over the measured spillover response.

- Developed a calibration update strategy for multiple sources. A single vertex calibration signal can be used to correct temporal gain drift. The formula for doing so was extended to incorporate multiple calibration sources to reduce additional effects.

- Introduced a beam pattern stability metric that computes the average temporal pattern variation at a given radial distance from, and with respect to, a unity gain main lobe peak.

- Developed a dual constraint beamformer capable of satisfying both a specified main beam gain constraint and a desired noise response constraint.

- Developed a rim constraint beamformer that uses calibration sources placed near the rim of the reflector dish to limit the spillover noise response through the use of eigenvector constraints.

1.5 Dissertation Outline

This dissertation presents analysis and solutions to some of the current challenges being researched in radio astronomy PAF development. The presentation of this work is as outlined below.

Chapter 2 provides background information relevant to the work presented in the remainder of the dissertation. It includes a description of the signal and noise models, the calibration procedure, an overview of beamforming methods, and introduces applicable performance metrics.

Chapter 3 discusses the issue of long-term calibration stability. A calibrator quality metric is presented, which is based on the minimum description length principle. A strategy for updating old calibration data based on measurements from a calibration update system is presented. Receiver electronics are shown to be fairly stable over a five minute period, based on measurements obtained in an anechoic chamber. Sensitivity and beam pattern measurements are provided to show the effect of temporal gain variations.

Chapter 4 compares PAF beamforming approaches. The trade-off between the two conflicting goals of high sensitivity and beam pattern shape control is balanced by implementing a hybrid beamforming approach that offers the best sensitivity for a given amount of beam pattern control. Beamformers designed in simulation require a transformation step before they can be applied to real measured data, but it is shown that PAF beamformer design is best completed using measured calibrators.

Chapter 5 considers relative beam variations that can obscure a weak signal of interest. These variations are addressed by implementing a dual constraint beamformer that simultaneously provides uniform main beam gain and desired noise response, and a rim constraint beamformer that implements eigenvector constraints based on reflector rim calibration data to control spillover response.

A discussion of conclusions and possible topics for related future work is presented in Chapter 6.

Experimental platform development and data collection field studies on world class telescopes are documented in Appendix B. These experiments and data motivated and provided the raw source material for the work presented in previous chapters.

Chapter 2

Background

This chapter presents the mathematical notation, signal models, algorithms, and basic theory necessary to understand the remainder of this dissertation. It is not meant to be an exhaustive reference source but to provide some basic background that is helpful to establish a starting point from which later work has been developed. A thorough introduction to radio astronomy principles, practice and theory can be found in [9, 68]. For more information on beamforming fundamentals, advanced beamforming theory, digital filter design, and array signal processing refer to [38, 39, 69–72].

2.1 Signal Model

As depicted in Figure 2.1, assuming narrowband operation of a P element PAF, the $P \times 1$ complex basebanded data vector at time sample n is given as

$$\mathbf{x}[n] = \mathbf{v}_{\mathrm{s}} s[n] + \eta[n] \tag{2.1}$$

where $\mathbf{v}_{\mathrm{s}}$ is the unit length array response to a unit amplitude signal in the far field arriving from the direction of a point source signal of interest (SOI) $s[n]$, and $\eta[n]$ is the array noise vector. Signals $s[n]$ and $\eta[n]$ are assumed to be independent zero mean random processes, statistically stationary across the L samples obtained during the specified observation time. The array covariance matrix is defined as

$$\mathbf{R}_{\mathrm{x}} = E\{\mathbf{x}\mathbf{x}^{H}\} \tag{2.2}$$

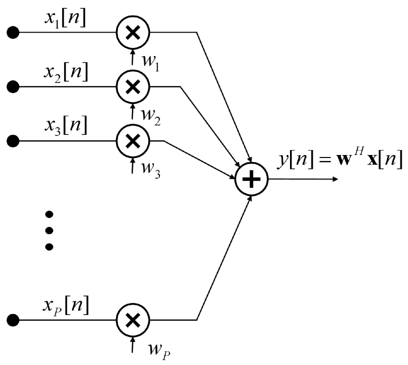

Figure 2.1: Block diagram for signal processing of a narrowband PAF.

where $E[\cdot]$ denotes expected value and superscript H is complex conjugate transpose (Hermitian transpose). Assuming the SOI and noise are statistically independent we have

$$\mathbf{R}_{\mathrm{x}} = \mathbf{R}_{\mathrm{s}} + \mathbf{R}_{\eta}. \tag{2.3}$$

Making the simplifying assumptions that the SOI is a point source and noise is white then

$$\mathbf{R}_{\mathrm{x}} = \sigma_{\mathrm{s}}^2 \mathbf{v}_{\mathrm{s}} \mathbf{v}_{\mathrm{s}}^H + \sigma_{\eta}^2 \mathbf{I} \tag{2.4}$$

where σ_{s}^2 and σ_{η}^2 are the power in the SOI and noise respectively, and $\mathbf{I}$ is the identity matrix. Because of the non-isotropic distribution of spillover ground noise seen by the array, and mutual coupling between the antenna elements, PAF system noise is in fact correlated, and this simplified model is rarely used (although it provides a reference for comparing performance results and may serve as a starting point for more sophisticated modeling).

An estimate of $\mathbf{R}_{\mathrm{x}}$ for an L-sample short-term integration (STI) window is calculated from observed data samples as

$$\hat{\mathbf{R}}_{\mathrm{x}} = \frac{1}{L} \sum_{n=0}^{L-1} \mathbf{x}[n]\mathbf{x}^H[n] = \frac{1}{L}\mathbf{X}\mathbf{X}^H, \tag{2.5}$$

$$\mathbf{X} = \big[\mathbf{x}[0], \mathbf{x}[1], \cdots, \mathbf{x}[L-1]\big].$$

The beamformer output is the weighted sum of the signals received by each array element and is computed as

$$y_k[n] = \mathbf{w}_k^H \mathbf{x}[n] \tag{2.6}$$

where $0 \leq k \leq K$ indexes one of K main lobe beam steering angles and $\mathbf{w}_k$ is the kth beamformer complex weight vector. Narrowband beamformer operation is assumed in this notation. For the broadband case the signal is decomposed into many frequency channels and separate beamformers with distinct weights $\mathbf{w}_k$ are computed as in (2.6) for each channel.

2.2 Noise Model

The weak astronomical signals sought for in radio astronomy are often dominated by relatively large noise sources. For a radio telescope operating with a PAF, these noise sources include cross-coupling between closely packed array elements, spillover noise seen by the array from beyond the edges of the reflector dish, sky noise collected by the main beam of the antenna radiation pattern, and additional noise due to resistive losses in antennas, cables, and connectors before the low-noise amplifier (LNA).

The spillover, receiver, main beam sky, and loss noise signals are each independent of each other. We compute a system noise covariance matrix $\mathbf{R}_\eta$ by adding the covariance matrices of the individual noise sources such that

$$\mathbf{R}_\eta = \mathbf{R}_{\text{rec}} + \mathbf{R}_{\text{sp}} + \mathbf{R}_{\text{sky}} + \mathbf{R}_{\text{loss}} \tag{2.7}$$

where $\mathbf{R}_{\text{rec}}$, $\mathbf{R}_{\text{sp}}$, $\mathbf{R}_{\text{sky}}$, and $\mathbf{R}_{\text{loss}}$ are defined in (2.9), (2.12), (2.14), and (2.15) respectively.

The beam equivalent system noise temperature is [28, 73]

$$\begin{aligned} T_{\text{sys}} &= T_{\text{rec}} + T_{\text{sp}} + T_{\text{sky}} + T_{\text{loss}} \\ &= T_{\text{iso}} \frac{\mathbf{w}^H(\mathbf{R}_{\text{rec}} + \mathbf{R}_{\text{sp}} + \mathbf{R}_{\text{sky}} + \mathbf{R}_{\text{loss}})\mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\text{iso}} \mathbf{w}} \\ &= T_{\text{iso}} \frac{\mathbf{w}^H \mathbf{R}_\eta \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\text{iso}} \mathbf{w}} \end{aligned} \tag{2.8}$$

where $\mathbf{R}_{\text{iso}}$ is the array covariance matrix due to an isotropic source at temperature T_{iso}. In practice, $\mathbf{R}_{\text{iso}}$ may be derived from array Y-factor measurements such as those described in [28].

A detailed explanation of the noise model is given in [74] with the exception of the main beam sky noise, which is presented in this section. A brief summary of each of the noise components and additional details relating to the spillover noise are also included here for convenience. The following models were used in all simulation experiments reported in this dissertation.

2.2.1 Receiver Noise

Due to the use of a high gain LNA at the beginning of the receiver chain, the dominant component of $\mathbf{R}_{\text{rec}}$ is LNA noise. A simplified model for $\mathbf{R}_{\text{rec}}$ assumes i.i.d. noise. This implies no correlation between channels and leads to a diagonal receiver noise covariance matrix. However, since the antenna elements of a PAF are situated closely together, there are significant LNA input noise cross-coupling effects, as described in [33], that are ignored by such a model. Including these effects the receiver noise covariance matrix $\mathbf{R}_{\text{rec}}$ is computed according to the model described in [58] as

$$\mathbf{R}_{\text{rec}} = 2B\mathbf{Q}[\mathbf{V}_{\eta,\text{R}}^2 + \mathbf{Z}_A\mathbf{Y}_{\text{c}}\mathbf{V}_{\eta,\text{R}}^2 + \mathbf{V}_{\eta,\text{R}}^2\mathbf{Y}_{\text{c}}^H\mathbf{Z}_{\text{A}}^H + \mathbf{Z}_{\text{A}}\mathbf{I}_{\eta,\text{R}}^2\mathbf{Z}_{\text{A}}^H]\mathbf{Q}^H \tag{2.9}$$

where B is the system noise equivalent bandwidth, $\mathbf{Z}_{\text{A}}$ is the system impedance matrix, and $\mathbf{Q}$ is defined as

$$\mathbf{Q} = \mathbf{Z}_{\text{R}}(\mathbf{Z}_{\text{R}} + \mathbf{Z}_{\text{A}})^{-1} \tag{2.10}$$

where $\mathbf{Z}_{\text{R}}$ is the load impedance matrix. $\mathbf{V}_{\eta,\text{R}}$, $\mathbf{I}_{\eta,\text{R}}$, and $\mathbf{Y}_{\text{c}}$ are diagonal matrices of noise voltage densities, noise current densities, and correlation admittances, respectively, whose values correspond to the elements in the array. This receiver noise model depends on a minimum achievable (at perfect impedance match) equivalent temperature T_{min} associated with each low-noise amplifier. In the absence of mutual coupling, $\mathbf{R}_{\text{rec}}$ is diagonal and (neglecting downstream noise in the receiver chain) the equivalent receiver noise temperature $T_{\text{rec}} = T_{\text{min}}$. When LNA noise couples back through array elements to neighboring closely

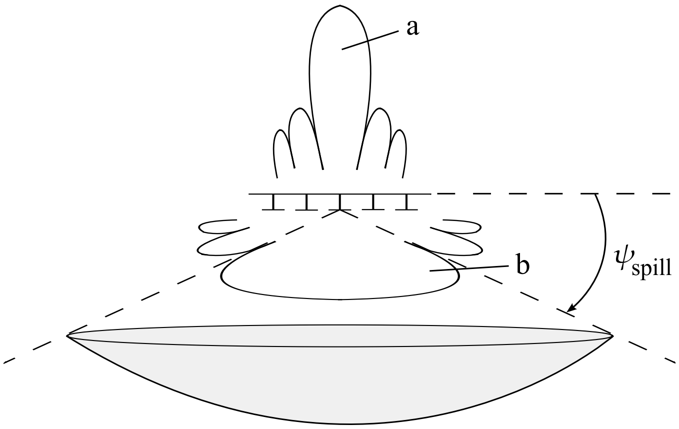

Figure 2.2: Array feed, dish, and spillover region geometry. The far field pattern illuminating the sky (a) is produced by the combined array-plus-reflector system. Dish illumination pattern (b) is determined by the PAF geometry, element response patterns and beamforming weight values $\mathbf{w}$. The spillover region spans the angle ψ_{spill} between the reflector edge and the blockage due to the array backing ground plane.

packed antennas, off diagonal terms in $\mathbf{R}_{\text{rec}}$ are non-zero, and T_{rec} is increased by a mutual coupling noise penalty.

2.2.2 Spillover Noise

The PAF beamformed illumination pattern extends beyond the edge of the reflector surface, collecting undesired signals arriving from the spillover region as shown in Figure 2.2. With a ground plane backing the array, back lobes are relatively small, so the spillover region is assumed to extend from the dish edge to the plane of the array. For a given reflector tipping angle, the spillover region includes both low temperature sky noise (~ 3 K) and relatively high temperature (280 K) ground noise.

To account for the spillover temperature distribution, we have modeled spillover noise as a dense grid of independent point sources with approximately uniform angular spacing on a spherical ring from the perspective of the PAF. For a given solid angle the intensity of the noise is modeled as constant for all look directions within the spillover region though in reality there is some variation, which we have observed in experimental data. As shown in Figure 2.3, the center is obscured by the primary reflector dish. As the dish tips, a portion of the grid rises above the extended horizon plane, corresponding to the part of the spillover illumination pattern observing cold sky rather than warm ground.

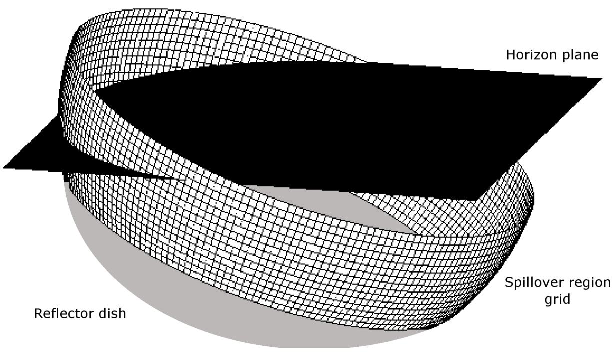

Figure 2.3: The spillover region is numerically modeled as a dense grid of statistically independent point sources, held fixed with respect to the array. Grid points cover the region spanned by ψ_{spill} in Figure 2.2. The plane represents the horizon, and for this figure the dish is tipped to 70° elevation. Grid points below the horizon plane correspond to warm ground seen by the PAF. As the dish is tipped, a portion of the spillover region passes through the sky/ground plane and is directed toward the cooler sky. The dish focal plane, containing the PAF array fround plane, also contains the upper rim of the illustrated spillover region.

Inter-element spacing between source points on the grid is uniform in angle ψ measured from the backplane, while circumferential spacing l in the azimuthal direction ϕ is

$$l = r\sin(\psi)\Delta\phi \tag{2.11}$$

where $r\sin(\psi)$ is the radius of the ring at depression angle ψ and $\Delta\phi$ is the angular separation between the grid points in azimuth. For simplicity, we assume that $r = 1$. The depression angle dependent spacing value l is used to weight the covariance matrix entries for each noise vector to ensure uniform local area average power distribution across the noise field.

Let $\mathbf{a}_i$ represent the complex voltage response across array elements to a unit amplitude source at the ith spillover noise grid position. We can approximate the spillover noise covariance matrix as

$$\mathbf{R}_{\text{sp}} = \frac{16k_bB}{\mid I_0 \mid^2}\frac{1}{2\eta_0}\sum_i T_i\mathbf{a}_i\mathbf{a}_i^H\alpha_i \tag{2.12}$$

where I_0 is the element excitation input current (for the receive array we assume $I_0 = 1$), η_0 is the intrinsic impedance of space, T_i is the noise temperature associated with the ith grid point, and α_i is the solid angle of the corresponding sky patch ($\alpha_i \propto l$).

The spillover grid remains fixed with respect to the array as the dish is tipped. For points below the horizon we use $T_i = 280$ K, representing warm ground. Points above the horizon correspond to sky noise whose temperature varies as $T_i = T_{\text{atm}}(\theta_i) + T_{\text{cmb}} + T_{\text{gb}}$ where $T_{\text{atm}}(\theta_i)$ is the elevation dependent atmospheric noise model developed below in (2.14) at zenith angle θ_i (see Figure 2.4), and T_{cmb} and T_{gb} refer to constant cosmic (CMB) and galactic background (GB) noise, respectively.

2.2.3 Main Beam Sky Noise

Atmospheric noise, CMB, and GB seen though the beamformer main lobe in the observation pointing direction all contribute to sky noise. These sources seen outside the main lobe are attenuated by the telescope's side lobe pattern and can be neglected in the total system noise model. Atmospheric noise increases as the dish is tipped toward the horizon, while CMB and GB noise are modeled as a constant at all elevations. As shown in Section 4.3.3, sky noise becomes the dominant source as the dish pointing elevation angle approaches the horizon.

Since over the span of the beam main lobe, atmospheric and sky noise appear isotropic, and since their levels seen through the beam side lobes are negligible we propose the model

$$\begin{aligned} \mathbf{R}_{\text{sky}} &= \frac{T_{\text{sky}}}{T_{\text{iso}}} \mathbf{R}_{\text{iso}}, \\ T_{\text{sky}} &= T_{\text{atm}} + T_{\text{cmb}} + T_{\text{gb}} \end{aligned} \tag{2.13}$$

where $\mathbf{R}_{\text{iso}}$ can be computed from the array pattern overlap matrix defined in [75]. Though the SOI is also seen in the main lobe, its contribution is contained in $\mathbf{R}_{\text{s}}$ while $\mathbf{R}_{\text{sky}}$ includes only noise terms.

For atmospheric noise, we use a modified plane-parallel atmosphere model. At a given elevation angle, T_{atm} is proportional to the line-of-sight thickness of the atmosphere, which for simplicity is assumed to be a solid slab of uniform thickness on a flat earth surface. As illustrated in Figure 2.4, path length through the atmosphere increases with depression angle θ according to $d(\theta) = d_0 \sec(\theta)$, where d_0 is the distance corresponding to the zenith direction. Since T_{atm} seen in the beam main lobe is approximately proportional to the corresponding

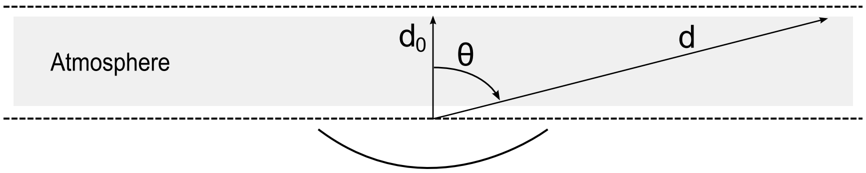

Figure 2.4: The distance through the atmosphere d is a function of the thickness of the atmosphere d_0 and the angle of depression from zenith θ. Main beam atmospheric noise T_{atm} is proportional to d along the line-of-sight.

$d(\theta)$, we have

$$T_{\mathrm{atm}}(\theta) = \begin{cases} T_{0,\mathrm{atm}} \sec(\theta) & 0 \leq \theta \leq 80^\circ \\ T_{0,\mathrm{atm}} \sec(80^\circ) + 1.3(\theta - 80^\circ) & 80^\circ < \theta \leq 90^\circ \end{cases} \tag{2.14}$$

where $T_{0,\mathrm{atm}} = 2$ K is the temperature at zenith. Note that for depression angles greater than 80° a correction is included to avoid the asymptote at the horizon [76]. We assume a constant isotropic distribution for cosmic background and galactic background noise, with $T_{\mathrm{cmb}} + T_{\mathrm{gb}} = 3$ K.

2.2.4 Loss Noise

Resistive losses in the antennas, cables, and connectors that appear before the LNA introduce a noise source that we approximate as being zero mean and independent from channel to channel, resulting in a diagonal covariance matrix model

$$\mathbf{R}_{\mathrm{loss}} = \sigma^2_{\mathrm{loss}}\, \mathbf{I}. \tag{2.15}$$

For a given T_{loss}, the appropriate value of σ^2_{loss} can be computed as

$$\sigma^2_{\mathrm{loss}} = \frac{T_{\mathrm{loss}}}{T_{\mathrm{iso}}} \frac{\mathbf{w}^H \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\mathrm{iso}} \mathbf{w}}. \tag{2.16}$$

The dominant source of loss noise in the 19 element Nagel PAF [59, 77] comes from a short length of coaxial cable which extends from the antenna to the LNA. Based on measurements of this cable we estimate T_{loss} to be 5 K [28].

2.3 Calibration Procedure

A calibration vector $\hat{\mathbf{v}}_k$ of the array voltage response to a far-field point source is required in every direction Ω_k that a beam is to be steered or where the pattern is constrained to a specified response value. Some details of the calibration procedure (reported in [28]) are repeated here since the information is crucial to understanding PAF beamformer design.

On-reflector calibration is necessary for accurate response vector estimation. Even the most detailed numerical simulations cannot predict the real physical array response with sufficient accuracy to design beamformer weights, since they must account for signal interaction with the reflector as well as gain variations between channels. Off-reflector bare array measurements are likewise unsuitable. Antenna range calibration is unrealistic since radio telescopes are physically too large and array responses drift too much over time. Additionally, due to mechanical limitations, multipath, and thermal ground noise, a reflector dish cannot be steered to sufficiently low elevations to use fixed man-made sources in the far field as calibration references. The only remaining option is to perform calibration on-reflector using the brightest available (isolated) astronomical source.

The calibration procedure can be summarized as follows. The radio telescope is steered, relative to the calibration source, in direction Ω_k for which a response vector is desired. An on-source, signal-plus-noise covariance $\hat{\mathbf{R}}_{\mathrm{x},k}$ is obtained. The instrument is then steered several degrees in azimuth away from the source at the same elevation to avoid changing the spillover ground noise pattern, and an off-source, noise-only $\hat{\mathbf{R}}_{\eta,k}$ is obtained. The calibration vector $\hat{\mathbf{v}}_k$ is computed as

$$\hat{\mathbf{v}}_k = \hat{\mathbf{R}}_{\eta,k}\mathbf{u}_k \tag{2.17}$$

where $\mathbf{u}_k$ is the principal eigenvector determined by the generalized eigenvalue problem

$$\hat{\mathbf{R}}_{\mathrm{x},k}\mathbf{u}_k = \lambda_{\max}\hat{\mathbf{R}}_{\eta,k}\mathbf{u}_k. \tag{2.18}$$

A grid of response vectors is computed in the region surrounding a calibration source. Up to 1000 distinct pointings may be required depending on the desired number of simultaneous beams in the FOV and the number of pattern constraints to be incorporated in the beamformer. This can be a time-consuming process (e.g., 4 hours), but cannot be neglected because obtaining accurate calibration vectors is fundamental to PAF beamformer design.

In general a commercial software package (such as Matlab) that is used to solve the generalized eigenvalue problem of (2.18) includes an arbitrary scaling of the eigenvectors. This scaling can be problematic since it does not retain the relative magnitude differences between the array responses for signals coming from various positions within the calibration grid. The relative magnitudes must be present in order to observe beam patterns and accurately generate power images. To maintain the relative magnitudes of the response vectors, $\mathbf{u}_k$ should first be normalized to unit length and then scaled by the square root of its corresponding eigenvalue before computing (2.17).

2.4 Beamforming Overview

One advantage of an antenna array is the ability to electronically control its radiation pattern through beamforming techniques. The appropriate choice of a set of complex weights, applied to the received signal, allows one to steer the main beam, manipulate the beam shape, and direct the placement of nulls. A beamformer may, appropriately, be considered a spatial filter, since it can block undesired signals from a given direction while giving more emphasis to those arriving from other directions. As shown in Figure 2.1 a narrowband beamformer can be represented as the inner product of a vector $\mathbf{w}$ of complex weights and the array sample vector. This dissertation considers two different categories of beamformers: statistically optimal and deterministic. A good introduction to beamforming can be found in [38].

2.4.1 Statistically Optimal Beamforming

A statistically optimal beamformer is optimal in the sense that it uses the statistics of the signal and noise environments to satisfy some specified criteria. Such a beamformer is often used in an adaptive mode since it can be updated periodically using current data samples to adapt to changes in the signal environment. Two of the most common beamformers from this category are the maximum sensitivity (max-SNR) and LCMV beamformers. A brief introduction to each is given here.

Max-Sensitivity/Max-SNR Beamformer

The optimum weight vector $\mathbf{w}_{\mathrm{m}}$ that achieves maximum sensitivity (or maximum SNR) is defined as [38, 39]

$$\begin{aligned} \mathbf{w}_{\mathrm{m}} &= \arg\max_{\mathbf{w}} \ \mathrm{SNR}, \\ \mathrm{SNR} &\triangleq \frac{\mathbf{w}^H \mathbf{R}_{\mathrm{s}} \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\eta} \mathbf{w}}. \end{aligned} \tag{2.19}$$

The maximization in (2.19) then gives the generalized eigenvalue problem

$$\hat{\mathbf{R}}_{\mathrm{s}} \mathbf{w} = \lambda_{\max} \hat{\mathbf{R}}_{\eta} \mathbf{w} \tag{2.20}$$

where $\lambda_{\max}$ is the eigenvector corresponding to the eigenvalue of maximum magnitude and the solution is the max-SNR beamformer. In practice, the estimated noise covariance matrix $\hat{\mathbf{R}}_{\eta}$ is obtained by (2.5) from noise-only samples. A distinct weight vector $\mathbf{w}_{\mathrm{m},k}$ is computed in this manner for each desired pointing direction Ω_k of the multiple simultaneously formed beams. $\mathbf{R}_{\mathrm{s}}$ is typically unknown, since this is the object we are attempting to observe with possibly unknown spatial structure. So, either a distributed source model is adopted for $\mathbf{R}_{\mathrm{s}}$, or the object is modeled as a collection of independent point source "pixels," and a separate beam is steered to each.

Assuming the point-source model, $\hat{\mathbf{R}}_{\mathrm{s}} = \sigma_{\mathrm{s}}^2 \mathbf{v}_{\mathrm{s}} \mathbf{v}_{\mathrm{s}}^H$ is rank one and the solution to (2.20) is

$$\mathbf{w}_{\mathrm{m}} = \alpha \hat{\mathbf{R}}_{\eta}^{-1} \mathbf{v}_{\mathrm{s}}^H \tag{2.21}$$

where σ_{s}^2 is the power associated with the signal of interest and α is an arbitrary scale factor that does not affect the final SNR at the beamformer output.

MVDR and LCMV Beamformers

The minimum variance distortionless response (MVDR) beamformer is designed to minimize the total output power while satisfying a single response constraint. As described in [38], the MVDR problem may be written as

$$\mathbf{w}_{\mathrm{MVDR}} = \arg\min_{\mathbf{w}} \ \mathbf{w}^H \mathbf{R}_{\mathrm{x}} \mathbf{w} \quad \text{subject to} \quad \mathbf{v}^H \mathbf{w} = f. \tag{2.22}$$

With $\mathbf{v} = \mathbf{v}_{\mathrm{s}}$ and $f = 1$, this forms a beam with a main lobe unity response in the direction of a point source corresponding to array response vector $\mathbf{v}_{\mathrm{s}}$. Solving for the beamformer weights results in

$$\mathbf{w}_{\mathrm{MVDR}} = f \frac{\mathbf{R}_{\mathrm{x}}^{-1} \mathbf{v}_{\mathrm{s}}}{\mathbf{v}_{\mathrm{s}}^H \mathbf{R}_{\mathrm{x}}^{-1} \mathbf{v}_{\mathrm{s}}}. \tag{2.23}$$

The MVDR beamformer can be generalized for multiple constraints in order to increase control over the beam pattern. For $J \leq P$ linear constraints we have

$$\mathbf{w}_{\mathrm{LCMV}} = \arg\min_{\mathbf{w}} \ \mathbf{w}^H \mathbf{R}_{\mathrm{x}} \mathbf{w} \quad \text{subject to} \quad \mathbf{C}^H \mathbf{w} = \mathbf{f} \tag{2.24}$$

which yields

$$\mathbf{w}_{\mathrm{LCMV}} = \hat{\mathbf{R}}_{\mathrm{x}}^{-1} \mathbf{C} [\mathbf{C}^H \hat{\mathbf{R}}_{\mathrm{x}}^{-1} \mathbf{C}]^{-1} \mathbf{f} \tag{2.25}$$

where the columns of $\mathbf{C} = [\mathbf{v}(\Omega_1), \cdots, \mathbf{v}(\Omega_J)]$ are calibration vectors associated with the J constraint angles and $\mathbf{f}$ is a $J \times 1$ vector of corresponding desired response values.

For radio astronomical observations it is common to obtain a separate estimate $\hat{\mathbf{R}}_\eta$ by steering the dish away from the object of interest to a relatively empty patch of sky. In this case it is often desirable to compute modified LCMV (or MVDR) beamformers using $\hat{\mathbf{R}}_\eta$ rather than $\hat{\mathbf{R}}_{\mathrm{x}}$. This reduces the possibility of SOI cancelation when observing distributed sources if there is calibration error in the estimates of the $\mathbf{v}(\Omega_J)$ steering vectors. The modified LCMV problem is

$$\min_{\mathbf{w}} \ \mathbf{w}^H \hat{\mathbf{R}}_\eta \mathbf{w} \quad \text{subject to} \quad \mathbf{C}^H \mathbf{w} = \mathbf{f}. \tag{2.26}$$

When $\mathbf{C} = \mathbf{v}_{\mathrm{s}}$ and $\mathbf{f} = 1$, this yields the max-SNR solution (2.21).

2.4.2 Deterministic Beamforming

Conventional deterministic beamforming is non-adaptive and does not rely on the statistics of the signal environment, i.e., it is data-independent. Known or estimated array response vectors $\mathbf{v}(\Omega_J)$ are used to construct beamformers with a desired response in specified directions Ω_J, or provide a desired beam pattern structure. The design objectives and methods are much like those for classical FIR filter design such as those described in [78].

A simple example is the conjugate field match (CFM) beamformer [38] which has beamformer weights defined as

$$\mathbf{w}_{\mathrm{cfm}} = \mathbf{v}_{\mathrm{s}}^{*}. \tag{2.27}$$

In spatially white noise the CFM beamformer maximizes SNR in the direction of the signal of interest specified by $\mathbf{v}_{\mathrm{s}}$. Alternatively, one can use a uniform magnitude and match only the phases $\mathbf{w}_{\mathrm{u}} = \exp(-\mathrm{j}\arg\{\mathbf{v}_{\mathrm{s}}\})$ where exp and arg operate element-wise. This works well for aperture arrays with identical elements, but due to focal effects for PAFs the element SNRs vary widely and $\mathbf{w}_{\mathrm{u}}$ performs poorly. However, by adjusting the relative gains of each weight the beam pattern can be altered to change the main beam width or side lobe levels. This is similar to incorporating a window in the design of an FIR filter.

Other approaches may require the use of iterative numerical optimizers if there is no known closed-form solution that meets the specified design objective. This is particularly true for 2-D beamformers which are considered for PAFs where closed-form analytical solutions are unknown. For example, 2-D equivalents to Parks-McClellan minimax designs are not known. Beamformers can be designed to provide a least-squares match to a desired pattern response while satisfying specified equality or inequality constraints. Numerical optimizers provide the necessary design flexibility and can be incorporated using any number of numerical software packages. However, design constraints and objectives can only be applied in directions Ω_J where calibrators $\mathbf{v}_j$ can be obtained. This is not possible beyond the main lobe and first one or two side lobes of the desired beam pattern. It is shown that fine grid numerical optimization with array models is not the best option for beamformer design.

The benefit of deterministic beamforming is the ability to design a well-prescribed and known beam pattern. Unfortunately this does not guarantee statistical optimality and achieving a high SNR without information about the statistical properties of the signal environment is a challenging problem. This is a concern for radio astronomy since signal powers are very small compared to the noise, but there is also a desire to have known beam pattern structure. Further discussion about these conflicting goals is provided in Section 4.1.

2.5 Performance Metrics

There are a number of metrics that are used to measure that performance of a PAF, including beamwidth, peak side lobe levels, sensitivity, aperture efficiency, and pattern stability. An introduction to several of these metrics is provided here for reference. Refer to [73–75] for a more thorough description of sensitivity and aperture efficiency.

2.5.1 Sensitivity

Sensitivity provides an indication of the ability of a PAF to detect weak astronomical sources. It is a scaled version of the array SNR and can be expressed as [56], [73]

$$S \equiv \frac{A_{\mathrm{eff}}}{T_{\mathrm{sys}}} = \frac{2k_{\mathrm{b}}}{10^{-26}F^{\mathrm{s}}}\,\mathrm{SNR} \tag{2.28}$$

where A_{eff} is the effective receiving area of the PAF system for the beamformer weight vector $\mathbf{w}$, k_{b} is Boltzmann's constant, and F^{s} is the flux density of the source signal of interest (Jy). Sensitivity is a very meaningful measurement since high sensitivity is required to detect the faint radio signals observed by radio telescopes.

2.5.2 Aperture Efficiency

There are a number of efficiency measurements that are used to characterize the performance of a PAF, including aperture efficiency, spillover efficiency, radiation efficiency, and noise matching efficiency [73, 74]. Aperture efficiency provides an indication of what percentage of the physical aperture is actually being used, and for a lossless antenna is defined as the ratio of the effective area to the physical area of the aperture,

$$\eta_{\mathrm{ap}} = \frac{A_{\mathrm{eff}}}{A_{\mathrm{phys}}}. \tag{2.29}$$

Aperture efficiency is computed as [56]

$$\eta_{\text{ap}} = \frac{k_b T_{\text{iso}} B}{A_{\text{phys}} F_{\text{s}}} \frac{\mathbf{w}^H \mathbf{R}_{\text{s}} \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\text{iso}} \mathbf{w}}. \tag{2.30}$$

The first term in (2.30) is independent of telescope pointing angle and is therefore constant.

Maximum antenna gain, G, relates to η_{ap} as

$$\eta_{\text{ap}} = \frac{\lambda^2}{4\pi} \frac{G}{A_{\text{phys}}} \tag{2.31}$$

where λ is the wavelength associated with the operating frequency.

From (2.28) and (2.29) we see that sensitivity is proportional to the ratio of the aperture efficiency to the system temperature,

$$S = \frac{\eta_{\text{ap}} \, A_{\text{phys}}}{T_{\text{sys}}}. \tag{2.32}$$

For additional information about aperture efficiency and the other efficiencies mentioned here refer to [73, 74].

2.5.3 Beam Pattern Stability

Measuring the beam pattern gain variations over a sufficient period of time provides an indication of the temporal stability. We compute a set of boresight beamformer weights $\mathbf{w}_{\text{bore},1}$ from an initial calibration set $\mathbf{V}_1$, where $\mathbf{V}_k = [\mathbf{v}(\Omega_1), \cdots, \mathbf{v}(\Omega_K)]$, and then apply it to array response vectors computed from subsequent calibration sets $\mathbf{V}_q$ to form a series of beam pattern vectors $\mathbf{b}_q = \mathbf{V}_q^H \mathbf{w}_{\text{bore},1}$ for Q calibration sets. The average temporal pattern variation δ in the ideally symmetric beam pattern at a given radial distance r from the main lobe peak is then determined by the following steps.

1. Normalize each of the Q beam patterns to a unity gain main beam peak.

2. Select M fixed points in the beam patterns $\mathbf{b}_q(i)$, $1 \leq i \leq M$ at which response stability will be sampled. Include several points at equal radial distance r from the main lobe

peak to sample pattern levels that should lie on an equal response contour of an ideal circularly symmetric response.

3. Identify the gain of the ith point of each beam pattern, i.e., $\mathbf{b}_q(i)$ for $q = 1, 2, \ldots, Q$, and construct the corresponding gain vector $\mathbf{z}_i = [z_{i,1} \cdots z_{i,Q}]$.

4. Divide $\mathbf{z}_i$ by its sample mean and subtract 1 to calculate the variation of each point from the mean.

5. Estimate $\delta(r)$ as

$$\delta(r) = \frac{1}{M} \sum_{m=1}^{M} \left(\frac{\mathbf{z}_m(r)}{\bar{\mathbf{z}}_m(r)} - 1 \right) \tag{2.33}$$

where the overbar indicates the sample mean of the argument.

Chapter 3

Long-term Calibration Stability

Adequate calibration is necessary in order to place multiple, simultaneous, well-formed beams on the sky and to provide an increased field of view over single feed systems. Due to the size of radio telescopes it is impractical to perform calibration measurements by any other method than using a bright astronomical source as the calibration reference. Since an array calibration vector is required for every direction in which a beam is to be steered and for every direction in which a beam pattern constraint is desired, it can take a significant amount of time to acquire an appropriate far field calibration data set. If electronic gain drift or other factors (such as mechanical variations and, for maximum sensitivity optimal beamforming, changes in spillover and sky noise) cause calibrators to vary significantly with time, then the PAF may need to be recalibrated. With regards to long-term stability, electronic gain drift is not a concern for a single horn feed telescope, but it is for a PAF where relative variations between channels can cause problems. Frequent PAF recalibration may be required to ensure that the user always has the most accurate calibrator estimates possible.

Operating under the assumption that calibrators drift significantly over time, and in order to both eliminate the need for frequent full on-sky recalibration and to always have access to accurate calibrators, a method for updating stale calibrators is desired. One proposed solution is to update the data set by using calibrator reference antennas placed at the vertex of the dish and at various locations around its rim [22, 42]. Variations in measurements of signals transmitted by these calibration sources over time can be used to account for drift in electronics or other time-dependent sources of variability. Existing calibration update systems and algorithms for PAFs have been relatively limited in complexity and capability. In [46] and [47], a single noise source was mounted at the vertex of the radio telescope. In both cases, the source was used to track short-term electronic gain variations, but in [47] it was found to provide very little benefit to updating beam patterns because the pattern variations were dominated by mechanical pointing errors.

The main purpose of this chapter is to consider the long-term stability of PAF calibration to understand the drift that occurs over periods of days and weeks. For diagnostic purposes, we have built and employed a multiple reference antenna calibration system to update stale calibration sets in response to relative receiver channel gain variations as well as more complex variations, such as fluctuations in the mechanical components of the dish and changes in the spillover and sky noise structure. Results are presented of a calibration study that shows variations in sensitivity and beam pattern stability which can be expected when using both old and current calibration data to compute beamformers. It is shown that temporal variations in the calibration vectors contribute only a small part to the total variation observed. Calibration data are relatively stable with respect to long-term variations and the use of a calibration update system is currently unnecessary.

Some results on short-term electronic stability are also presented. Data for this analysis were collected in an anechoic chamber without the reflector or any noise contributions from a realistic observation environment. This setup was used to isolate and analyze drift in the receiver electronics and data acquisition system alone (independent of dish structural stresses) as a way to separately identify these components of the overall calibration drift.

3.1 Calibration Vector Quality Metric

The quality of calibrations $\hat{\mathbf{v}}_k$ depends on SNR for the chosen bright calibrator source and integration time per pointing. SNR drops significantly for grid points outside the FOV imposed by the natural dish aperture response pattern. Thus calibrations outside the first dish side lobe are unreliable and must be discarded.

We have employed an algorithm based on the minimum description length principle (MDL) [39, 79, 80] to detect the presence of a sufficiently strong calibrator source in the received data. Eigenvalues of the noise-pre-whitened covariance $\breve{\mathbf{R}}_{\mathrm{x},k} = \hat{\mathbf{R}}_{\eta}^{1/2}\hat{\mathbf{R}}_{\mathrm{x},k}\hat{\mathbf{R}}_{\eta}^{1/2}$ are used to estimate d, the number of dominant sources present with adequate SNR. Pre-whitening is required for MDL since PAF noise is correlated. The $(P-d)$ smallest eigenvalues of (2.4) are clustered near σ_{η}^{2}, and the d remaining eigenvalues correspond to source signals. Detection of

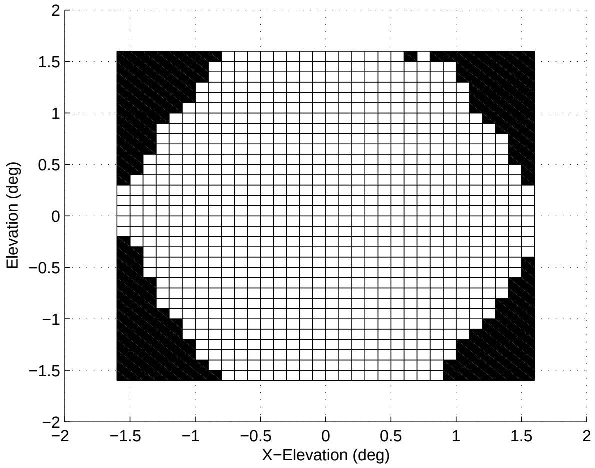

Figure 3.1: Acceptable calibration vectors are determined by the MDL quality metric, which is an objective algorithm used to identify the existence of a single source. White spaces indicate a good calibrator and black spaces indicate a poor one. The source was Cassiopeia A.

a single dominant non-noise signal is an indication that the calibration vector is acceptable. Compute

$$\hat{d} = \arg \min_{d=0,1,\cdots P} \left\{ L_d(d) + \frac{1}{2}[d(2P-d)+1]\ln(\alpha L) \right\}, \tag{3.1}$$

$$L_d(d) = \alpha L(P-d) \ln \left\{ \frac{\frac{1}{P-d}\sum_{i=d+1}^{P} \hat{\lambda}_i}{\left(\prod_{i=d+1}^{P} \hat{\lambda}_i\right)^{\frac{1}{P-d}}} \right\}$$

where $\hat{\lambda}_i$ are the eigenvalues of $\breve{\mathbf{R}}_{\mathrm{x},k}$ ordered as $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \cdots \geq \hat{\lambda}_P$, and L is from (2.5). Parameter α, $0 < \alpha \leq 1$ determines the degree of dominance (i.e., calibrator SNR) required. For $\alpha = 1$, $\hat{d}$ is the conventional MDL estimate of just detectable independent sources, while $\alpha < 1$ establishes a higher SNR threshold. If $\hat{d} \neq 1$, then SNR is assumed insufficient for calibration at that grid point.

Figure 3.1 is a real data example of a measured 33 × 33 calibration grid using a 19 element PAF on the Green Bank 20-Meter Telescope. Positions of acceptable calibration vectors using $\alpha = 1.6 \times 10^{-4}$ are marked in white, showing that there are practical limits to the range of steering vectors that can be obtained with on-reflector calibration.

An alternative method for identifying good calibrators is by manual inspection of the eigenvalues of the signal covariance matrix $\hat{\mathbf{R}}_{\mathrm{s}}$. A single dominant eigenvalue indicates the presence of the calibration source.

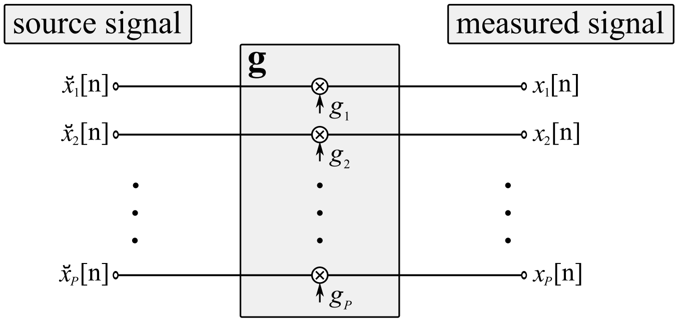

Figure 3.2: Error terms introduced in the calibration source signal before it is measured by the PAF include individual electronic gain variations ($\mathbf{g}$).

3.2 Calibration Update Strategy

The magnitude of the calibration drift can be identified using several small antennas positioned at various locations around the rim of the reflector dish to measure changes in the system over time. If necessary, this information can be used to update a set of stale PAF calibrators. At least one calibration source is needed to correct relative receiver channel gain variations. A graphical representation of these unknown per-channel error terms is given in Figure 3.2. The procedure for updating calibration data in response to electronic gain drift is described below.

At time index $q = 1$ a full grid set of calibration vectors $\mathbf{V}_1 = [\mathbf{v}_1(\Omega_1), \cdots, \mathbf{v}_1(\Omega_K)]$ is initially estimated for K calibration grid pointings using a bright astronomical calibration source. A corresponding array response to the dish-mounted single reference calibration signal $\mathbf{a}_1$ is computed according to the method described in Section 2.3, but with the calibration antenna as the source signal. As time passes electronic gain variations cause $\mathbf{V}_1$ to become outdated. Later, at time index $q > 1$, a current estimate of the calibration vector set $\mathbf{V}_q$ is desired. In practice it is inconvenient to use valuable telescope observing time to directly measure $\mathbf{V}_q$. The goal of the calibration update is to use a current rim calibration source measurement $\mathbf{a}_q$ to compute a contemporary estimate of the calibration as

$$\mathbf{V}_q = \mathbf{G}_q \mathbf{V}_1. \tag{3.2}$$

When using a single calibration source antenna to solve for independent gain correction terms per channel the qth complex per-channel update is

$$\mathbf{G}_q = \mathrm{diag}\{\mathbf{a}_1 \ ./\ \mathbf{a}_q\} \tag{3.3}$$

where the operator ./ indicates element-wise division of the vectors, and the function "diag" converts a vector into a matrix with the vector elements appearing along the diagonal.

In an effort to improve the approximation of $\mathbf{G}$ this per-channel electronic gain drift update can be extended to include contributions from multiple calibration sources by solving the least squares problem

$$\mathbf{G}_q = \arg\min_{\mathbf{G}} ||\mathbf{G}\mathbf{A}_1 - \mathbf{A}_q||_{\mathbf{F}}^2 \tag{3.4}$$

where $\mathbf{A}_q = [\mathbf{a}_q(\Omega_1), \cdots, \mathbf{a}_q(\Omega_J)]$ are the qth set of calibration vectors associated with the J calibration update system sources, $||\cdot||_{\mathbf{F}}$ is the Frobenius norm, and $\mathbf{G}$ is restricted to be diagonal. Using the property that for an arbitrary matrix $\mathbf{X}$

$$||\mathbf{X}||_{\mathbf{F}}^2 = \mathbf{tr}(\mathbf{X}^H\mathbf{X}) \tag{3.5}$$

where $\mathbf{tr}(\cdot)$ is the matrix trace, the problem can be rewritten as

$$\mathbf{G}_q = \arg\min_{\mathbf{G}} \ \mathbf{tr}\left[(\mathbf{G}\mathbf{A}_1 - \mathbf{A}_q)^H(\mathbf{G}\mathbf{A}_1 - \mathbf{A}_q)\right]. \tag{3.6}$$

Multiplying the terms and computing the derivative using the property for the derivative of diagonal matrix $\mathbf{Y}$ with respect to the trace

$$\frac{\partial \mathbf{tr}(\mathbf{XY})}{\partial \mathbf{Y}} = \mathbf{X} \odot \mathbf{I} \tag{3.7}$$

where $\odot$ is the Hadamard element-wise product and $\mathbf{I}$ is the identity matrix, the derivative is equated to zero and $\mathbf{G}$ is computed to obtain the solution to (3.4)

$$\mathbf{G}_q = \left[\left(\left(\mathbf{A}_1\mathbf{A}_1^H\right) \odot \mathbf{I}\right)^{-1} \left(\left(\mathbf{A}_1\mathbf{A}_q^H\right) \odot \mathbf{I}\right)\right]^H. \tag{3.8}$$

With multiple reference sources mounted on the dish it may also be possible to detect and correct directionally dependent variations (i.e., where the diagonal model of (3.3) does not hold) without doing a full on-sky calibration. The drift model can be extended to include a correction that accounts for additional variations that may occur in time. Define $\mathbf{G} = (\Gamma)\,(\mathrm{diag}\{\mathbf{g}\})$, where Γ is an arbitrary $P \times P$ matrix, to construct a fully filled correction matrix $\mathbf{G}$ that accounts for electronic gain drift as well as time-varying directionally dependent terms such as structural perturbations in the dish surface and feed supports due to thermal expansion and gravitational sag, mutual coupling variations possibly caused by increased LNA noise or signal reflections due to increased mismatch at the union of the antenna and LNA, variations in the focus position perpendicular to the array, or differences in mechanical response to thermal variations. Comparison of sequential rim calibration data suggests that there is a need to include Γ in the model since the differences between calibration vectors over time cannot all be accounted for with a diagonal correction. We have explored the inclusion of Γ to solve for $\mathbf{G}$ and found that it did not add any further improvement in the estimation of $\mathbf{V}_q$. As shown in Section 3.5, the development of more sophisticated algorithms that include Γ are unnecessary because the variations in the noise field (specifically spillover and main beam sky noise) are so large that even using a diagonal correction is futile.

3.3 Evaluation of Receiver Electronic Gain Drift

In October 2009, the BYU 19-element thin dipole array shown in Figure 3.4 mounted on a rotating arm, was set up in the anechoic chamber at the University of Utah and measurements were taken to measure electronic gain variations in the system. These variations are expected to provide the largest contribution to the overall drift of the PAF system.

The receiver system used for these measurements comprised 19 Ciao Wireless LNAs, receiver boxes, Filtronetics IF filters, National Instruments A/D cards, and the byu20meter Dell Server PC. A picture of the system as it was setup outside the anechoic chamber is shown in Figure 3.3. For a more detailed description of this system refer to [59, 81].

In order to quantify variations in the receiver system electronics, a collection of five-minute measurements was made every 20 to 30 minutes over an eight-hour period. The array

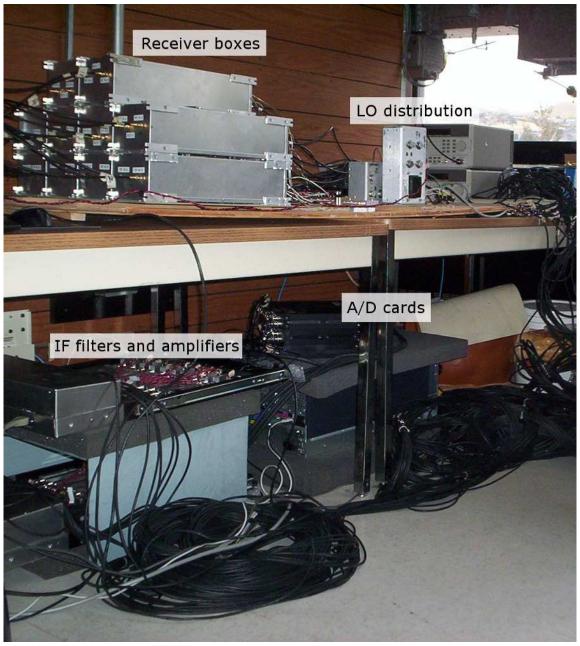

Figure 3.3: The receiving system, including receiver boxes, LO distribution network, IF filters and amplifiers, the National Instruments breakout boxes, and the ADC cards installed in the byu20meter PC. This setup was used at the anechoic chamber at the University of Utah in October 2009.

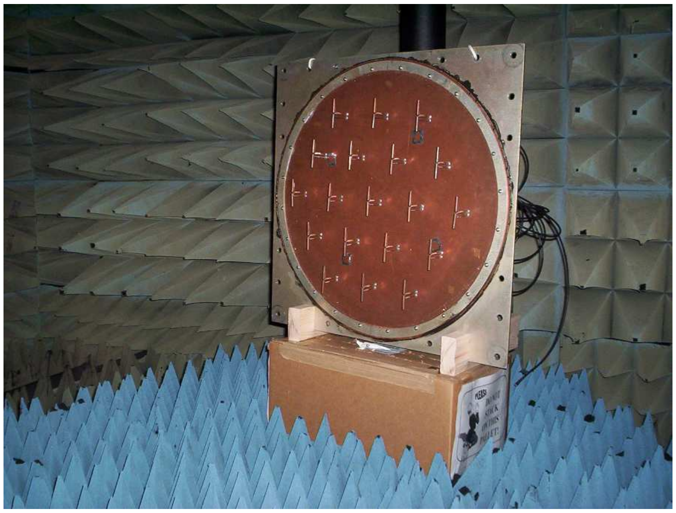

Figure 3.4: The BYU 19-element array mounted on the rotating arm in the University of Utah anechoic chamber.

and receiving system were left untouched between measurements in order to ensure that any variations from one data set to the next were due solely to drift in the electronics. A total of 17 measurements were used for this analysis.

The results presented below show the amplitudes and phases of each channel relative to a reference channel. These were obtained from correlation matrices $\mathbf{R}_{\mathrm{x},k}$ estimated for each data set.

Figure 3.5 compares the electronic drift over the full eight hours. Amplitude values are given in Figure 3.5(a) and were obtained from the diagonal elements of the covariance matrices, i.e, the values of $\sqrt{\mathrm{diag}\{\mathbf{R}_{\mathrm{x},k}\}}$ are plotted as a function of k in units of seconds. The values of each channel were normalized by the respective initial measurement. With the exception of channel 17 (the top red line), all the channel amplitudes drift in the same manner. The total range of the drift is relatively small (about 2-3% on average) and the relative gains between channels remain about the same throughout the eight-hour period. Relative phase variations are shown in Figure 3.5(b), where the phases are given in degrees and are relative to channel 1, i.e., they are drawn from the first column of each covariance matrix so that $\arg\{\mathbf{r}_{1,k}\} - \arg\{\mathbf{r}_{1,1}\}$ is plotted where $\mathbf{r}_{i,k}$ is the ith column of $\mathbf{R}_{\mathrm{x},k}$. In each case the initial phase was subtracted off so that the relative drifts could be more easily compared. The dashed line shows the zero phase of channel 1. Over the span of the eight-hour period of this experiment the largest phase drift between any two channels is between 0.5-0.6°. This is relatively small, especially considering that in general there was no special attention given to minimize noise and drift when selecting the components of the receiver system.

Relative phase drifts between channels for a single five-minute data set are shown in Figure 3.6. The dashed line shows the channel 1 gain. This result is from the first five-minutes but similar results can be produced by using any of the other 16 data sets. The total phase drift after five minutes is less than 0.1°, but a significant portion of that is due to the initial jump that is seen to occur within the first 30 seconds of the data. Interestingly, there are three channels (channels 2, 3, and 4) that don't experience this jump and seem to drift much less relative to channel 1. Channels 1-4 are on the same data acquisition cards as the channel 1 reference channel. To verify that this initial jump was a phenomenon of

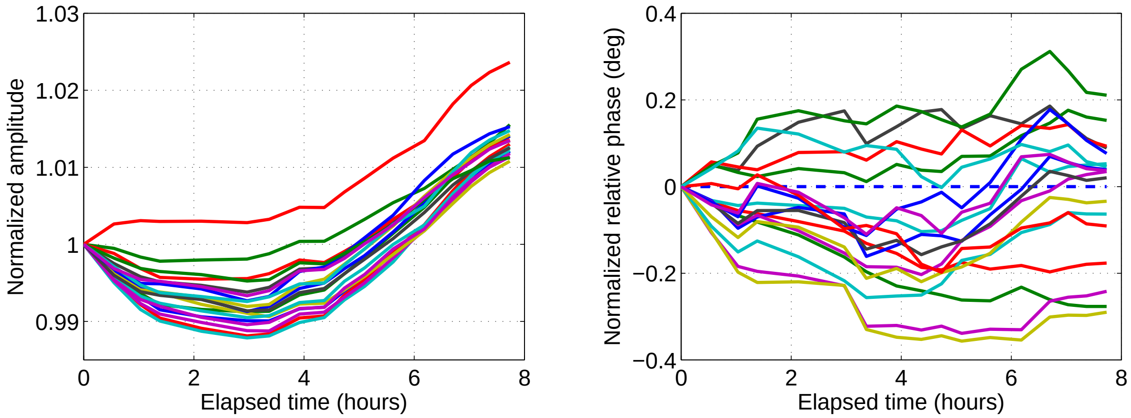

(a) Amplitude drift over 8 hours.

(b) Phase drift over 8 hours.

Figure 3.5: Receiver system electronic drift for an eight-hour period. (a) shows the amplitude drift of each channel, normalized by the first measurement. (b) is the phase drift of each channel relative to the first channel with the initial measurement subtracted off to allow for side-by-side comparisons.

the data acquisition cards, additional measurements were collected by bypassing the receiver electronics and inputting a signal directly into the cards. As seen in Figure 3.7 the jump between ADC cards at the beginning of the file is still present with this setup. There is much less total measured drift without this jump, which appears to just be a problem associated with the first data acquisition card since all channels not on card 1 seem to drift in the same manner. Still, the peak relative drift here of 0.1° is considered to be an insignificant source of calibration error.

3.4 Long-term Calibration Stability Tests

Far-field calibration data were collected in 2011 at the NRAO facility in Green Bank, WV for analysis of the calibration drift that occurs in a real observation environment over long periods of time. Principal goals for these tests were to understand how frequently calibration grids need to be measured, and to evaluate effectiveness of using on-dish calibration reference sources to update far-field calibration vectors. Experiments were conducted on the 20-Meter dish shown in Figure 3.8, which is a 20-meter diameter parabolic reflector with a focal length over diameter (f/D) ratio of 0.43. Four arms extend from the dish surface to support the feed at the focal point. The array, known as the Carter array, appears in Figure

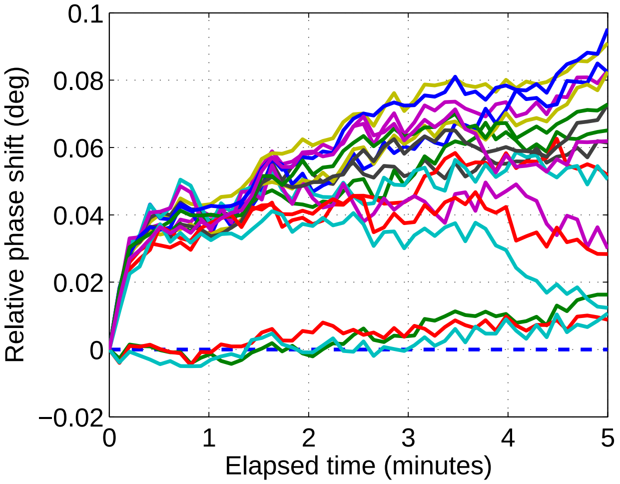

Figure 3.6: Receiver system electronic phase drift for a five-minute period. The phase drift of each channel is relative to the first channel with the initial measurement subtracted off to allow for side-by-side comparisons.

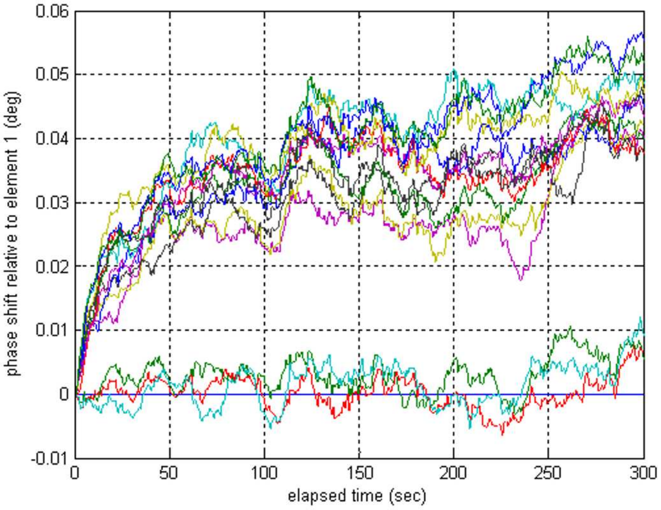

Figure 3.7: Data acquisition ADC card phase drift for a five-minute period. The phase drift of each channel is relative to the first channel with the initial measurement subtracted off as in the previous figures.

3.9, which consists of 19 hexagonally spaced half-wave thickened dipoles, separated by 0.6λ on a ground plane [82]. Each element was designed to provide an active match to account for and reduce the effects of mutual coupling [58, 82]. The PAF is mounted to the exterior of the "front-end" box at the focal point of the reflector. The data acquisition system was

Figure 3.8: Green Bank 20-Meter Telescope.

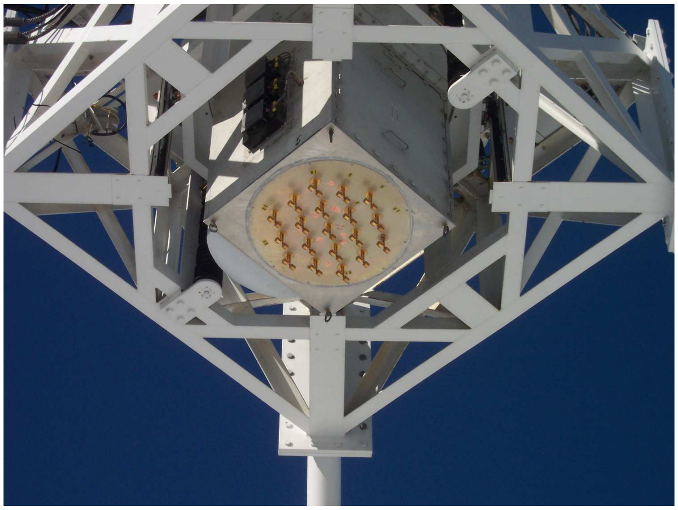

Figure 3.9: The BYU active-impedance-matched 19-element Carter PAF: hexagonally spaced half-wave thickened dipoles, separated by 0.6λ on a ground plane. The PAF is attached to the end of the front-end box which is mounted to the 20-Meter Telescope. The front end box contains the LNAs, receiver cards, power supplies, and LO distribution hardware.

built by NRAO Green Bank using a single Adlink Technology Inc. ADC card in each of 10 PCs. This is similar to the BYU system described in Section B.1.1.

A high-level description of the design and construction practice for the PAF receiver system is provided below to establish that the calibration stability performance reported below is achievable with very ordinary means, and with no special attention to drift. A calibration update system is also introduced along with a description of the experiments performed.

3.4.1 Data Acquisition Hardware

The front-end box contains the low noise amplifiers (LNAs), down-converter receiver cards, power supplies, and local oscillator (LO) distribution hardware as explained in [81]. All of the electronics used in this system are built from commercially available off-the-shelf components.

As shown in Figure 3.10 the LNAs are connected directly to the output of the antenna elements (on the back side of the ground plane) in order to minimize loss and noise. An 18-inch piece of low-loss coaxial cable connects each LNA output to a receiver card input. The receiver card, shown in Figure 3.11 was jointly designed by Vikas Asthana and myself. It uses a two-stage super-heterodyne design that mixes the input L-band signal down to a final IF frequency of 2.8125 MHz. Each of the 10 cards contains four receiver channels which achieve cross channel isolation of at least 40 dB. The RF mixers, LO splitters, amplifiers, and most of the filters are standard components available from Mini-Circuits and Analog Devices, Inc. The LO signals are provided by Agilent/Hewlett Packard HP8648D signal generators, which are located in the antenna mount base pintel room and are remotely controlled by the data acquisition system through a General Purpose Interface Bus (GPIB) interface. Aside from the signal generators, all other components are standard, non-military, non-instrumentation grade components. A VanLong SAW filter is used at the first IF stage and a custom LC bandpass filter from KR Electronics is used at the second IF.

Separate 170-foot long coaxial cables for each channel carry the receiver final IF signal from the front end box down to the data acquisition system located in the pintel room. This long exposed cable run is another potential source of calibration drift. The down-conversion is deliberately done in the front-end box so that phase variations between receiver channels are relatively unaffected by long cable lengths. The front-end box is temperature controlled but there is no temperature compensation incorporated in the receiver system design.

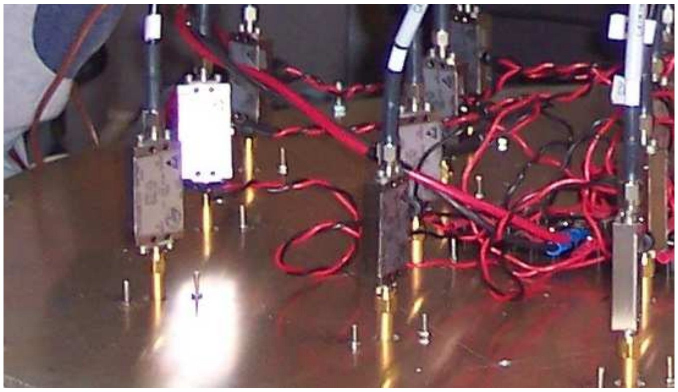

Figure 3.10: The array dipole antennas were designed to connect directly to the LNAs, on the back side of the ground plane, to minimize loss.

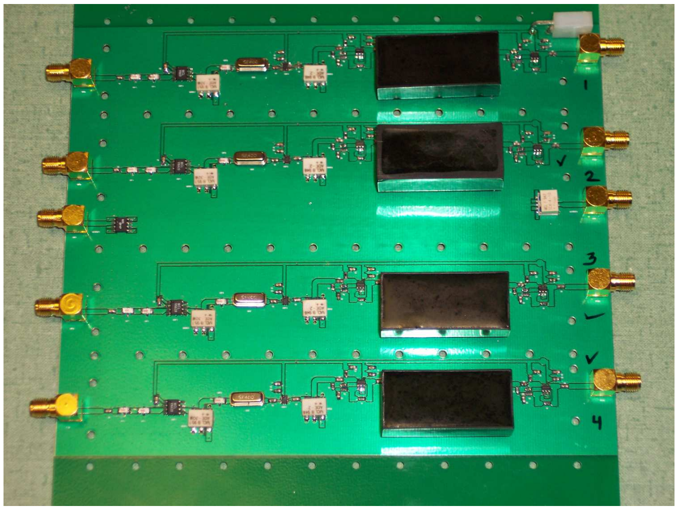

Figure 3.11: The four-channel down-converter/receiver board (one of five).

3.4.2 Calibration Update System

The calibration update reference source system employed on the telescope consisted of four cylindrical L-band horns, each with the purpose of periodically transmitting reference signals to the array which could be used to quickly update far field calibration vector estimates. One horn was mounted at the vertex of the dish and the remaining three at the base of three of the four dish feed support arms, as shown in Figure 3.12. Each calibrator antenna was pointed directly at the bare array and aligned with the polarization direction of the array elements for maximum efficiency. Originally, a fifth horn was also mounted next

Figure 3.12: Calibration horns mounted on the Green Bank 20-Meter Telescope. Inset is a closer look at one of the horns attached to a feed arm.

to the PAF, directed at the dish surface, but it provided no useful information so it was excluded from the analysis.

To avoid having a separate signal generator for each calibrator antenna, a coaxial switch was used to route a single signal to any selected antenna. The Dynatech single pole, six throw (SP6T) 28-volt coaxial switch is rated for 18 GHz. A Hewlett-Packard 11713A switch driver was used to control the coaxial switch (see Figure 3.13). It is capable of manual control or remote computer control using a GPIB interface. The GPIB feature was important because it allowed the calibration system to be incorporated into the existing automated data acquisition code. Since the coaxial switch was to be located in the open air at the vertex of the dish it was mounted in a metal box to protect it from the weather (shown in Figure 3.14).

3.4.3 Experimental Procedure

Far-field calibration grids were measured about every 12 hours over a five day period during February and March of 2011. In all, the main body of data collected for this analysis comprises a total of 10 calibration grids. Each of these data sets, measured on the astronomical source Cassiopeia A (Cas A), is a 23×23 grid of five-second pointings spaced 0.15° apart and ranging from -1.65° to $+1.65^\circ$ in both elevation and cross-elevation (i.e.,

Figure 3.13: Single pole, six throw switch driver for calibration update system.

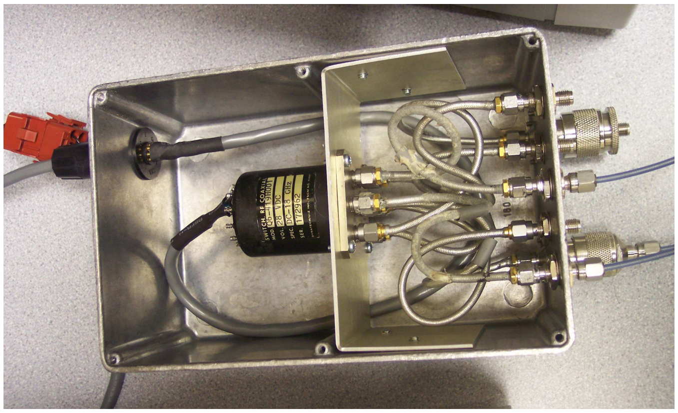

Figure 3.14: Weather resistant metal box containing the coaxial switch and associated cabling for the calibration update system.

the arc direction perpendicular to elevation) with respect to boresight. Two off-source measurements were taken at each elevation, ±5° from boresight, as estimates of the noise field corresponding to all calibrators at that elevation.

In addition to these 10 calibration grids, data from other smaller grids that had originally been obtained for other purposes was used where available. These grids were also measured using Cas A and provide a larger sample size for a more thorough study of PAF calibration. In some cases there were smaller grids of 11×11 pointings and in other cases only a single on and off pair were needed to compute a max-SNR beamformer.

When using multiple calibration grids to compare the responses to a given direction it is important to ensure that response vectors from each are actually pointing in the same relative direction. All the calibration grids collected during the 2011 experiments were properly aligned, i.e., they had the same initial elevation and cross-elevation offsets relative to the center of the calibrator source. To confirm that this was indeed the case, each calibration grid was checked to verify that the power measured on element one (the center element) for the center grid point had the highest measured power of all the grid points. This check helped to ensure that there is not pointing mismatch errors in the analysis. For grids of different sizes, only those points that corresponded to the same relative direction could be used for comparison. A method for aligning multiple grids with different initial offsets (such as those obtained during the experiments of 2008) is given in Appendix C.

A signal transmitted from each of the calibrator horns was measured before and after each far-field calibration grid. This involved sending a series of sine-wave calibration signals from the single function generator, through the remotely-controlled coaxial switch, to each of the calibration horns in sequence. Simply using a splitter to transmit a single calibration signal from each calibrator was inadequate since this does not produce independent steering vectors from each source. Each rim calibrator transmitted a signal at a different frequency within the passband of the system as a failsafe way to distinguish the different signal during later processing. Function generator output power was −50 dBm, which is 30 to 40 dBm above the noise floor.

3.5 Experimental Results

In this section the metrics of sensitivity and beam pattern stability are used to determine the extent of calibration drift over time. By quantifying the drift we can determine whether there is a need for frequent recalibration of a phased array feed.

3.5.1 Calibration System Integrity

Before drawing conclusions based on the rim calibrator data, it was important to validate the system. Each of the four dish-mounted calibrator sources should be able to measure electronic drift. Since they are located at different positions relative to the array, the associated response vectors will vary. What should be common is the change in the response over time, i.e., the diagonal correction matrices of each individual calibrator source. Figure 3.15 shows a comparison of the electronic updates computed from each individual calibrator

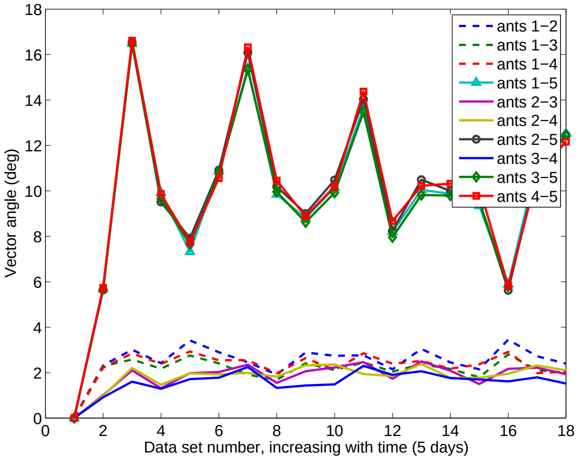

Figure 3.15: The electronic updates computed from each calibrator horn are compared with a vector angle measurement. Antenna 1 is at the vertex, antennas 2-4 are on the rim, and antenna 5 is at the apex. In the legend, *ants i-j* indicates a plot of the vector angle between array responses from calibrator antennas i and j. There is a much larger angle between the apex horn and the four others, which appear to vary on the same order.

source over the five-day period. Each line is a comparison of vector angle between a different pair of antennas, computed as

$$\theta_{\text{vec}} = \cos^{-1} \frac{|v_1^H \ v_2|}{||v_1||_2 \ ||v_2||_2}. \tag{3.9}$$

We expect a difference from one calibrator antenna to the next but that difference should remain somewhat constant over time. All of the calibrators seem to agree with each other to the same degree, with the exception of those that involve a comparison to the apex antenna (the fifth antenna) which was located near the array (this result supports the decision to exclude the fifth antenna as mentioned in Section 3.4.2). However, from this analysis it is unclear if a 2-3° difference between different calibration source measurements means that each is giving the same information. Simply comparing the vector angles between the diagonal update matrices is not conclusive since this does not provide any sense for what is a reasonable difference.

More information can be obtained by computing a series of diagonal update matrices from one of the calibrators and then using that information to estimate the steering vectors

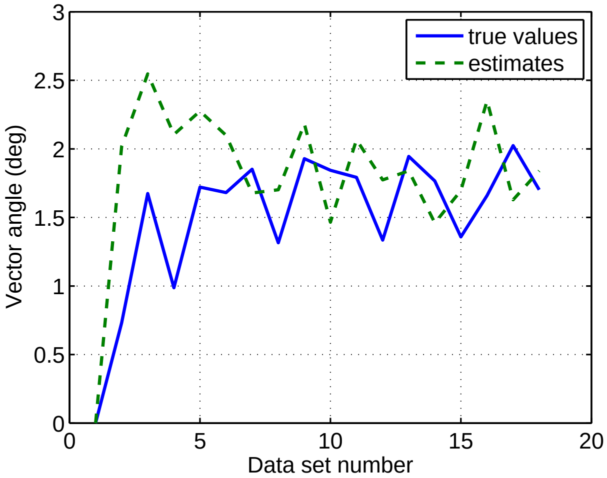

Figure 3.16: The variation between the measured response vectors of a single calibrator horn is on the same order as the variation between the estimated response vectors of that same calibrator. The estimates were computed from diagonal updates derived from another of the calibrator horns. In this case the updates were taken from the vertex horn data and then applied to the data from one of the rim calibrators.

of another calibrator source. Vector angles between the true response vectors of this second source and those of the estimated response vectors can then be compared. If the estimated angles match up with the true angles then we can conclude that the 2-3° difference identified previously is acceptable, and that both calibrator sources measured the same electronic gain drift over time. Figure 3.16 presents this analysis. The vertex calibrator was used to compute a series of diagonal update matrices that correspond to the changes in electronic gain over the five-day experiment period. Each of the updates was then applied to the initial steering vector of one of the rim-mounted calibrators in order to compute an estimate of the true steering vector over the five days. The vector angle between the true steering vectors at each time and the initial steering vector is then compared to the vector angle between the estimated steering vectors at each time and the initial steering vector. The result shows the variation between the estimated steering vectors is on the same order as that of the true steering vector, implying that the same electronic gain drift is measured from both of the different calibrator sources. Similar comparisons between other pairs of calibrators are consistent with this result.

3.5.2 Sensitivity Variations

Sensitivity is one figure of merit used to characterize the quality of a beamformer. The useful life of a set of calibrators is determined by examining the sensitivity degradation over time computed using max-SNR beamformers. Besides the 10 main calibration grids there were eight additional data sets from which max-SNR beamformers were computed. The total span of time covered by this collection of calibrators was 70 days, between January 4 and March 14. For each observation the telescope was moved as needed to point to the astronomical calibration source.

Sensitivity variations for three different cases of the boresight max-SNR beamformer is shown in Figure 3.17. At each point the current signal and noise covariance estimates are used to compute the sensitivity. The differences between the plotted lines arise from the use of different approaches to computing the beamformer of (2.19). The three approaches shown here are: 1) the fresh beamformer is computed using $\hat{\mathbf{R}}_{\mathrm{s}}$ and $\hat{\mathbf{R}}_{\eta}$ from the current calibration set which represents the best sensitivity that can be achieved at the given moment in time, 2) the stale beamformer which is computed using $\hat{\mathbf{R}}_{\mathrm{s}}$ and $\hat{\mathbf{R}}_{\eta}$ from the first calibration set, and 3) the noise-updated beamformer which is computed using $\hat{\mathbf{R}}_{\mathrm{s}}$ from the first calibration set and $\hat{\mathbf{R}}_{\eta}$ from the current calibration set.

This figure shows significant variation (up to 7%, or 0.26 dB) in the attainable sensitivity (achieved with the fresh beamformer) from one calibration grid to the next. Variations are primarily due to changes in the noise field as the dish is steered to different parts of the sky. Specifically, the spillover and main beam noise seen by the array are dependent on elevation angle. The best example of this was from the five-day data where diurnal variations were observed between the calibration grids that were measured in the morning and those from the evening. The physical temperature of the operating environment may also affect the electronics and mechanical components of the system, which could appear as noise fluctuations.

If the stale beamformers are used instead of obtaining a fresh calibration grid each time beamformers are needed then a slight penalty is incurred. The average loss in sensitivity is about 0.9%, or 0.015 dB, throughout the 70 day period. It is remarkable that there is relatively little loss in sensitivity for using a stale beamformer. In fact, the loss due to using

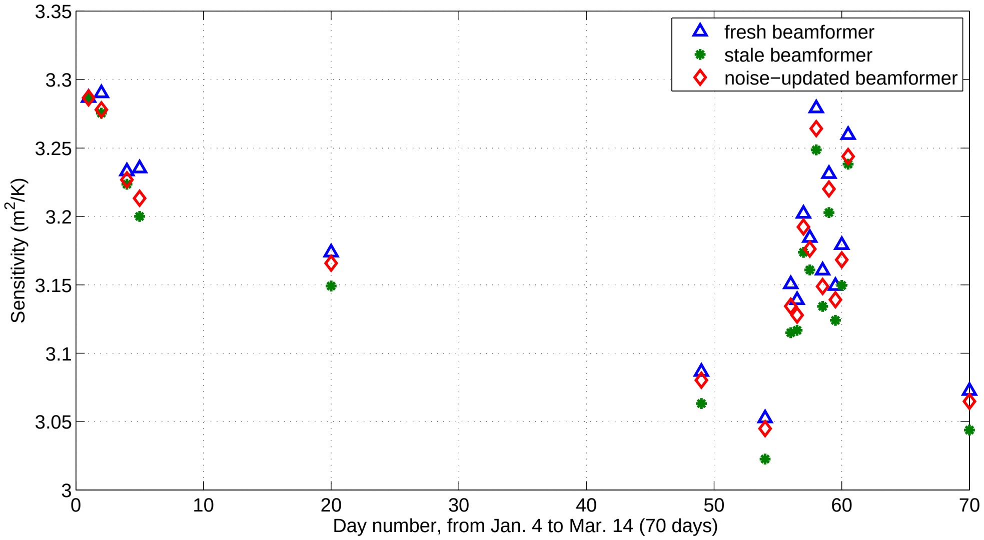

(a) Sensitivity variations, 70 days.

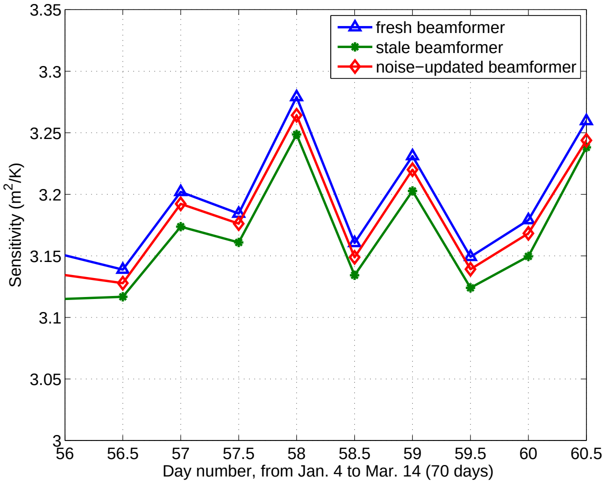

(b) Sensitivity variations, 5 days.

Figure 3.17: Comparison of the array sensitivity when using fresh and stale max-SNR boresight beamformers. If stale beamformers are recomputed using a fresh noise field estimate then the loss in sensitivity is reduced by half. (a) shows the full 70 days. (b) shows days 56-60, for which a calibration grid was obtained each morning and evening (note the diurnal pattern).

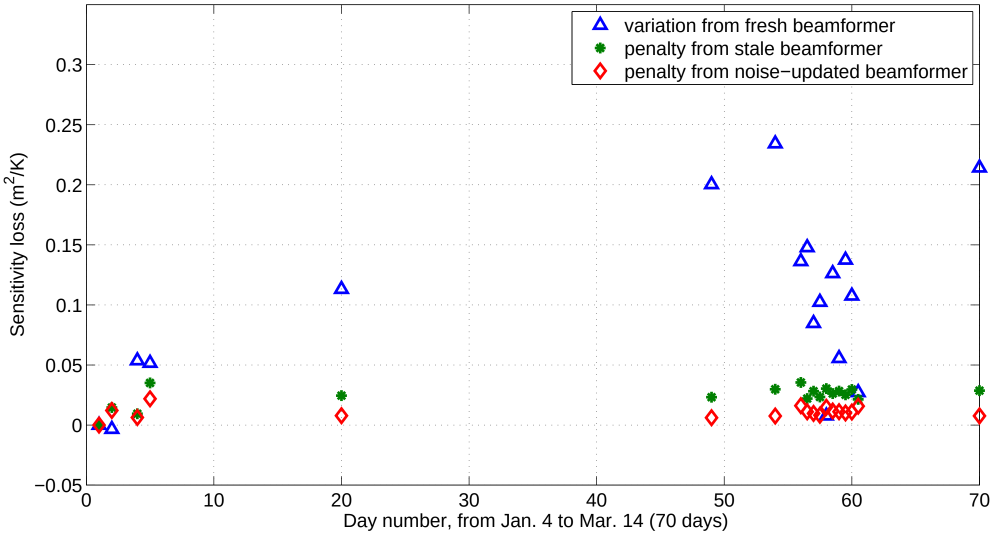

(a) Sensitivity penalty, 70 days.

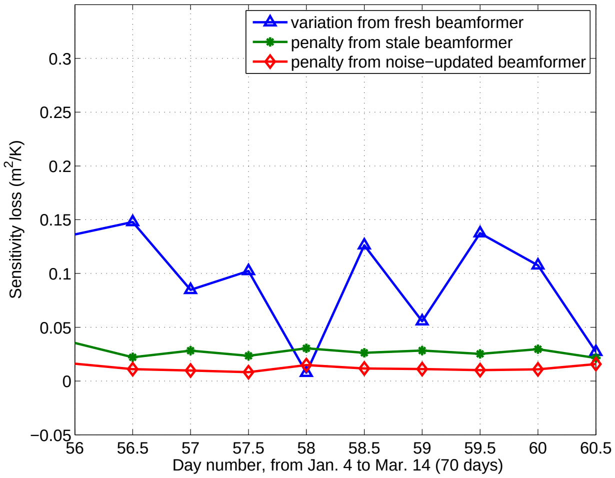

(b) Sensitivity penalty, 5 days.

Figure 3.18: The penalty for using a stale boresight beamformer is much less than the variation due to other factors. The penalty is reduced by updating the max-SNR beamformer using a fresh estimate of the noise field. (a) shows the full 70 days. (b) shows days 56-60, for which a calibration grid was obtained each morning and evening.

a stale beamformer is much less than the natural change that can be expected by pointing to a different part of the sky. This penalty is more clearly seen in Figure 3.18. Interestingly though there is an initial increasing penalty, but after a few days a ceiling is reached above which the penalty does not increase.

Sensitivity can be further reduced by simply using the noise-updated beamformer. A current noise estimate can be easily obtained without significantly delaying an observation. The consequent sensitivity loss relative to the fresh beamformer is shown in Figs 3.17 and 3.18. Average loss is only 0.35% over 70 days. This result shows that calibrators remain fairly stable over time and that a large portion of the degradation is due to a changing noise field. For a few days the stale beamformer works almost as well as the noise-updated beamformer but even after that point there is only a slight difference between the two.

Similar results are obtained when evaluating an off-boresight pointing, where the distribution across the array of the larger magnitude beamformer weights in $\mathbf{w}_k$ is not so heavily concentrated on the center element, but spread out among several elements.

To understand if the inherent variation in sensitivity is acceptable it is helpful to compare to the performance of a conventional single horn feed. Lacking single horn data the PAF center-element beamformer (coefficients are all zeros except for the one corresponding to the center element) is used to give an general idea of variation. Obviously the center element of the array does not represent the best performance that can be achieved with a horn feed due to a larger spillover component in its dish illumination pattern, but variations in sensitivity should provide some insight. Figure 3.19 shows that variations of the PAF system are on the same order as the fixed, calibration free, center element beam.

3.5.3 Beam Pattern Variations

To study the effect of using stale beamformers on pattern stability the equiripple deterministic beamformer, which is presented in Section 4.1.2, is implemented. An equiripple beamformer was computed using the first calibration data set. The main beam radius was 0.8° and four constraint points were equally spaced near the edge of the main beam, 35 dB down from the main beam peak. Fresh beamformer weights for each subsequent calibration set $\mathbf{V}_k$ were then obtained by solving the least squares minimization problem

$$\mathbf{w}_{\mathrm{e},k} = \arg\min_{\mathbf{w}} ||\mathbf{w}_{\mathrm{e},1}\mathbf{V}_1 - \mathbf{w}\mathbf{V}_k||_2^2 , \; 2 < k < K \tag{3.10}$$

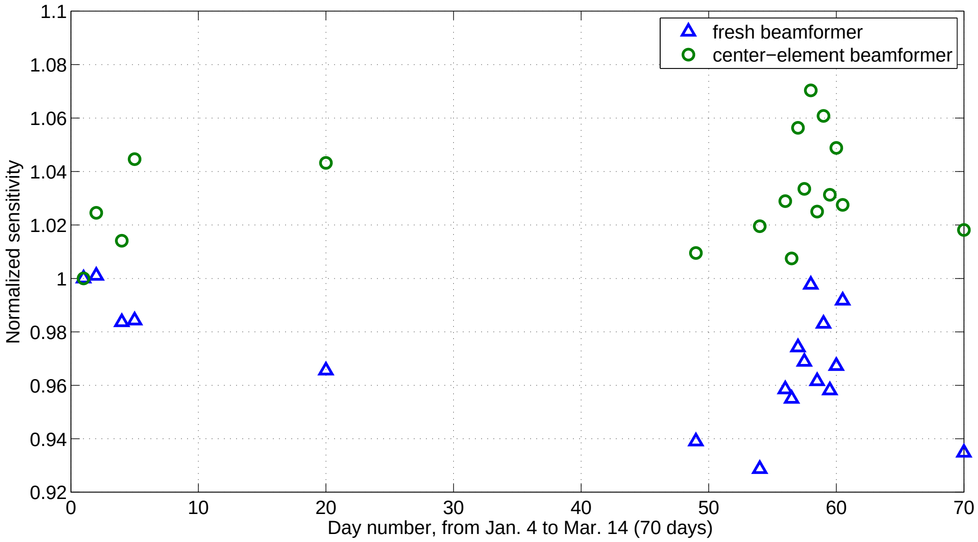

(a) Sensitivity comparison, 70 days.

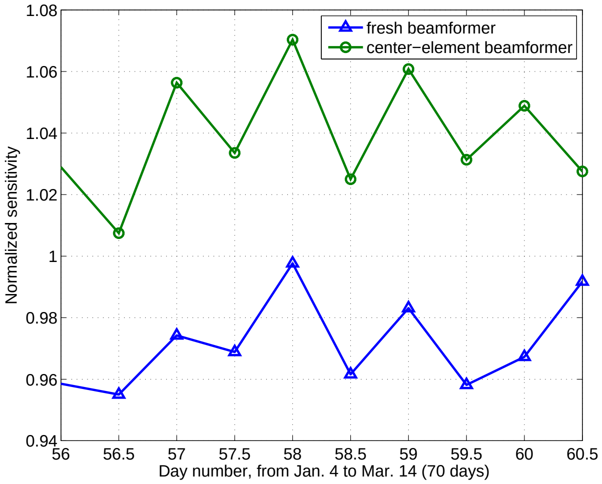

(b) Sensitivity comparison, 5 days.

Figure 3.19: The boresight beam sensitivity variation of the phased array feed is consistent with the variation of a single feed system. (a) shows the full 70 days. (b) shows days 56-60, for which a calibration grid was obtained each morning and evening (note the diurnal pattern).

where K is the number of calibration sets used in the analysis. The solution is the beamformer $\mathbf{w}_{\mathrm{e},k}$ that minimizes pattern variation relative to the initial beam as sampled at calibration grid points. The fresh beamformers obtained in this manner represent the best stability that can be expected.

If PAF pattern stability is computed according to (2.33), the expected degradation from using an old calibration set can be determined. The following results compare the performance of a fresh beamformer, an old beamformer, and a gain-corrected old beamformer. Four points within the main beam (at about 0.3° from boresight) and four points in the first side lobe (at about 1.05° from boresight) were used to compute the main beam variation and the side lobe variation, respectively. The gain correction was computed with all four of the calibration horns using (3.8). While a single calibration horn is adequate to correct for electronic gain drift, it was found that using it alone introduced additional variation in the beam pattern. Using the four calibration horns provides the same level of gain correction without sacrificing pattern stability. Remarkably, the long-term stability results presented in this section are consistent with the short-term measurements presented in [43, 47].

Figure 3.20(a) shows that the main beam stability of the boresight beam pattern is only slightly degraded when using an old beamformer. Additionally, the per-channel gain correction tracks the old beamformer plot and does not offer any improvement in stability. This behavior is repeated in the side lobe variations of Figure 3.20(b). For the five days of calibration data, there is a 2-3% variation in the main beam and 3-4% in the side lobes. Similar results are obtained with off-boresight steered beams, as shown in Figs 3.21(a) and 3.21(b), where the beam is steered to -0.45° in cross-elevation relative to boresight.

A visual inspection of the beam patterns helps to give physical meaning to the quantitative values given above. Figure 3.22(a) is the boresight beam pattern for calibration set 1 (day 1) using the beamformer from calibration set 1. Figure 3.22(b) is the boresight beam pattern for calibration set 9 (day 5) using the beamformer from calibration set 9. Figure 3.22(c) is the boresight beam pattern for calibration set 9 using the old beamformer from calibration set 1. Figure 3.22(d) is the boresight beam pattern for calibration set 9 using the gain-corrected beamformer from calibration set 1. Calibration set 9 was chosen arbitrarily

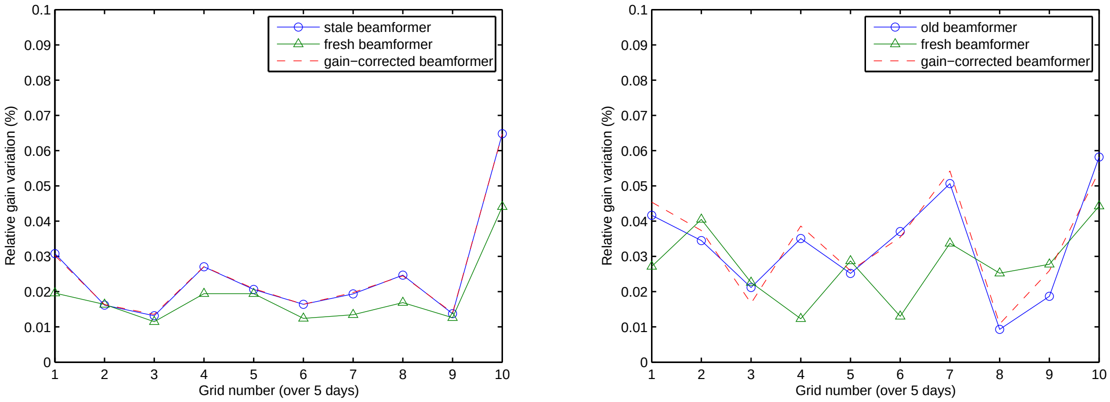

(a) Main beam variation.

(b) Side lobe variation.

Figure 3.20: Beam pattern variations in the main lobe (a) and side lobe (b) of a boresight directed beam. Using an old beamformer results in very little degradation. Applying the electronic gain drift correction does not improve the performance.

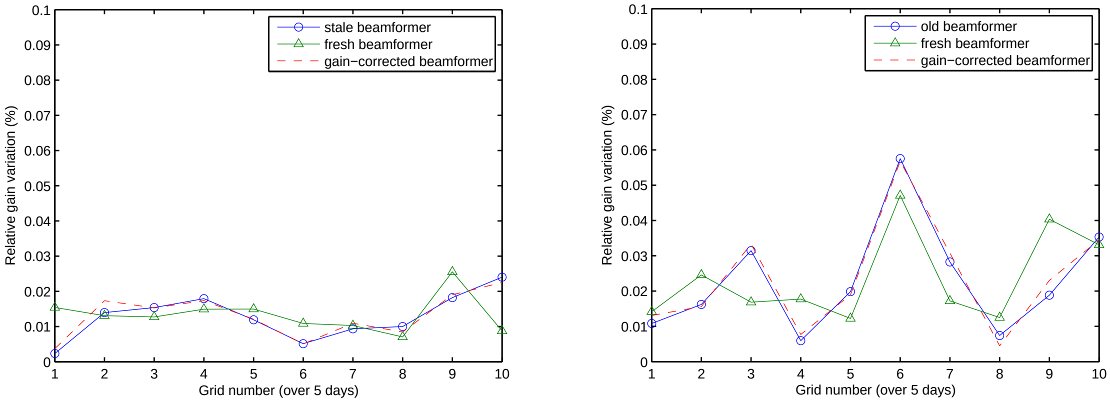

(a) Main beam variation.

(b) Side lobe variation.

Figure 3.21: Beam pattern variations in the main lobe (a) and side lobe (b) of an off-boresight directed beam. Using an old beamformer results in very little degradation. Applying the electronic gain drift correction does not improve the performance.

for use in this comparison. The beam patterns from each of these cases are all very similar. From visual inspection alone it is difficult to see the differences.

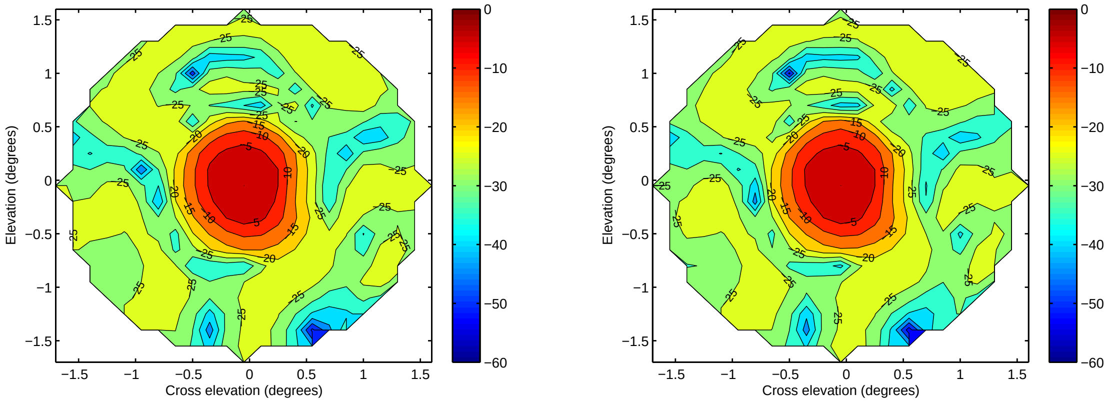

(a) Day 1 pattern with fresh beamformer.

(b) Day 5 pattern with fresh beamformer.

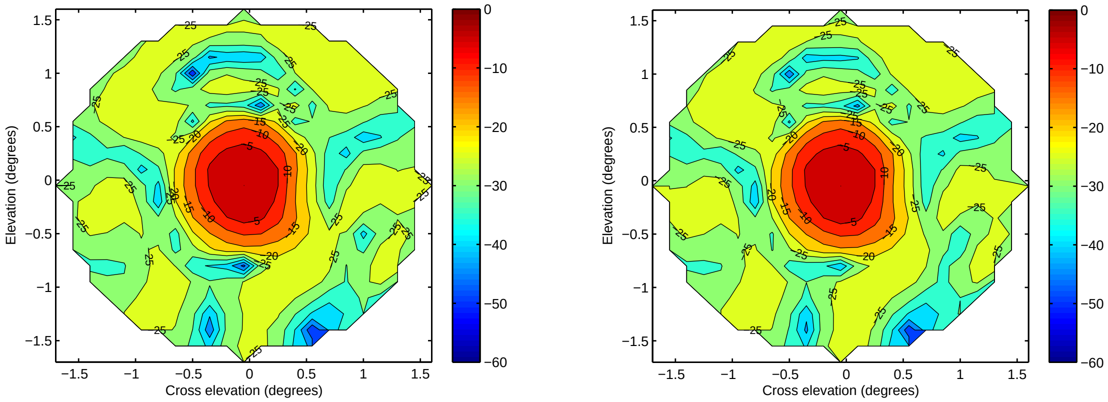

(c) Day 5 pattern with old beamformer.

(d) Day 5 pattern with gain-corrected beamformer.

Figure 3.22: Boresight beam patterns from days 1 and 5 using different beamformers: fresh from the current day, old from day1, and gain corrected from day 1. There are only slight differences between each of these beam patterns.

3.6 Conclusion

Accurate PAF calibration data is required to compute desired beamformers for radio astronomy applications. PAF calibration varies with time depending on a number of factors including gain drift in the receiver electronics and array-dish geometry drift. For max-SNR beams, changes in spillover noise as the telescope is pointed in different directions requires current noise field estimates to achieve maximum sensitivity. Frequent acquisition of a full set of calibrators is a time-consuming process that could interrupt valuable observations. In this chapter contributions from the drift of receiver electronics have been quantified and analysis of long-term calibration stability presented. We conclude that calibration measurements are stable for at least five days, and possibly much longer.

The variation in sensitivity was found to be affected more by other factors than by the age of the calibrators from which the beamformer was derived. For a period of time covering 70 days there were total variations of up to 7% measured, while the average loss in sensitivity due to the use of an old beamformer was only about 0.9%. This penalty was shown to be reduced even further, to about 0.35% by simply using a fresh noise field estimate in the computation of the max-SNR beamformer. In either case, though, the loss is negligible because of the large variations due to changes in spillover noise and mechanical pointing error.

The stability of beam patterns was also used to study the extent of calibration drift. For the five-day period for which we collected calibration data, the measured beam pattern variations were about the same for both the fresh and old beamformers. The variations were no greater on average than about 2-3%. This was true for main lobe variations, side lobe variations, boresight directed beams, and off-boresight beams. Additionally, correcting for electronic gain drift was shown to be unnecessary because any improvement that this may provide is relatively small compared to the total variations due to the other factors mentioned previously.

These results show that radio astronomical PAF calibration data is sufficiently stable over a five-day period (and possibly for up to 70 days or more). This is a welcome finding because it means that there is not a need for frequent recalibration or a sophisticated calibration update system. These conclusions may depend on specific electronics and rigidity of the reflector dish construction. Our receivers used conventional electronic components and design methods, and the Green Bank 20-Meter Telescope was designed for accurate steering and may be more rigid than some larger dishes [83]. Further research is needed to determine if calibrator corrections can be made to account for the effects of mechanical pointing error and spillover noise changes. However, at this point, the long-term stability of the PAF is acceptable since it appears to be consistent with that of a traditional horn feed radio telescope.

Chapter 4

PAF Beamforming

PAF beamformer design methods must balance the demand for high sensitivity with a known, stable beam pattern spatial structure. High sensitivity is required to detect the faint radio signals emitted by distant celestial bodies, which are typically tens of dB below the noise floor [84], while gain and noise stability is needed for radiometry, high dynamic range, and weak source detection. This tradeoff occurs because the maximum sensitivity beamformer now used by all PAF development groups achieves its optimal SNR performance by adapting the beam pattern to the noise field and array response parameters. This causes some variation across beams in the far-field beam pattern. It is shown that the inherent tradeoff between these two parameters is a motivating force in PAF beamformer design. This chapter compares several beamforming methods and discusses the implications of using each. Additionally, an example demonstration shows the benefits of a PAF for adapting to changes in the sky and spillover noise environments.

4.1 PAF Beamformer Design

A variety of beamformers have been suggested for PAF use [56, 85], but the data-dependent max-SNR (max-sensitivity) beamformer [49], simultaneously introduced for astronomical PAF use by the ASTRON and BYU-NRAO teams, has been the only one successfully applied to create images of experimental PAF data [28,50–55]. Other PAF beamforming work has been limited to simulation models but has intriguing potential. This includes the use of eigenbeams to reduce data transfer and storage requirements [61] and numerically optimized Gaussian beams steered without distortion, while accounting for polarization effects [62]. The linearly constrained minimum variance beamformer (LCMV) is appealing because of its ability to provide beam pattern constraints while minimizing the overall noise response [38]. Conjugate field match beamforming has also been attempted, but has been found to be unsuitable due to the inability to control beam pattern shape when there is significant gain variation across the sensor array for far-field sources [56–59].

Each of these design methodologies has drawbacks. The max-SNR beamformer does not guarantee pattern stability since it can be recomputed to optimize over variations in the noise field. The closely related LCMV approach is promising, though in some cases there are insufficient degrees of freedom in the beamformer for a typical PAF to constrain the side lobe pattern to be uniformly low while controlling the main beam shape. Deterministically derived beamformers offer complete control of the beam pattern shape (within the calibration field of view) but lack the necessary sensitivity for useful observations. Accurate array calibrations are required in any case, and this is a significant technical challenge. Desires for maximum sensitivity and complete control of the beam shape cannot be mutually satisfied.

This section compares options that are available to address the new challenges of practical PAF beamformer design. Of the several candidate beamformer methods mentioned above, we will consider the two that most obviously represent the extremes in trading off high sensitivity with direct pattern control for stable on-sky responses: the max-SNR and numerically optimized approaches. A hybrid beamformer is introduced, which combines the positive aspects of these approaches to balance the tradeoff between sensitivity and beam pattern shape control. A review of the practical implementation challenges for each given realistic array calibration data, and a performance analysis and comparison is presented. While it is possible to design a deterministic beamformer using simulation models, a transformation step is required before it can be satisfactorily applied to measured data. The hybrid beamformer is shown to provide a solution suitable for PAF operation.

4.1.1 Max-SNR Beamformer

The max-SNR beamformer maximizes PAF sensitivity in a desired direction. This beamformer is easily implemented since steering a beam in direction Ω_k only requires a single pair of on and off-source measurements, and a single calibration vector $\mathbf{v}(\Omega_k)$. When computed with previously acquired calibration data the beamformer is said to operate in “fixed-adaptive” mode, meaning that it is optimal for the calibration set but does not rely on current observation array covariance estimates. When noise spatial distribution across the array and electronic and mechanical instrumentation responses are stable (a challenge for this multichannel system) then near optimal high sensitivity performance is achieved for days or weeks without recalibration (see Chapter 3). Fully adaptive operation requires frequent updates of $\hat{\mathbf{R}}_{\mathrm{n}}$ throughout the observation time in order to cancel interfering signals or account for changes in the noise spatial structure such as spillover noise variation as a function of pointing elevation (see Section 4.3) [28].

The max-SNR beamformer offers little direct control of the beam pattern shape, since it naturally responds to the noise covariance structure. Variations in the observation environment due to changes in spillover noise structure and electronic drift in the receiver system change the underlying array response vectors. This presents a dynamic range challenge since unpredictable beam pattern variations may introduce undesired signals from a nearby bright source. A natural extension of the max-SNR beamformer that provides constraints on the beam pattern structure is the LCMV beamformer, which is discussed in Section 4.2.2. We will see that it has limitations similar to those of the max-SNR beamformer.

4.1.2 Equiripple Beamformer

Deterministic beamformers are designed to achieve a specified beam pattern structure and can be derived using a number of methods. This work utilizes an iterative optimization routine based on the minimax principle to construct an equiripple beamformer using calibration data (modeled or measured). The equiripple method was chosen for convenience and is simply a representative of a large class of deterministic beamformers. Complex beamformer weights are found by iteratively evaluating the far-field response at available calibration points throughout the beam pattern to achieve minimum gain equal ripple side lobes within the FOV while meeting designated main lobe shape constraints. Since this approach optimizes the beamformer weights only with respect to specified response pattern goals, the correlated spillover and mutually coupled array noise response is not considered, so true maximum sensitivity in not achievable. However, when strict control providing known pattern shapes across all formed beams is required, this approach may be desirable.

A similar numerical optimization with respect to the dish illumination pattern could be used to control illumination spillover and thus directly reduce noise levels, but obtaining a sample calibration of points on the dish rather than in the on-sky far-field beam is difficult or impossible using available calibrator sources. A system of calibration reference sources placed near the dish rim is described in Section 3.4.2, and used for improved beamformer stability in Section 5.3.2. This does provide limited control in the spillover region but cannot support full illumination pattern calibration. Thus only far-field pattern optimization has been considered in this analysis.

A pattern magnitude response equality constraint in direction Ω_k is specified as

$$|\mathbf{v}^H(\Omega_k)\mathbf{w}| = c_k \tag{4.1}$$

where $\mathbf{v}(\Omega_k)$ is the array response vector for direction Ω_k, H indicates the Hermitian transpose operation, and c_k is the specified gain referenced to a unity gain main lobe peak. Additionally, the phase of the first element in $\mathbf{w}$ is set to zero, eliminating an ambiguous degree of freedom. The peak response across all side lobes for a given weight vector is given by

$$F(\mathbf{w}) = \left|\mathbf{V}_{\text{side}}^H \mathbf{w}\right|_\infty \tag{4.2}$$

where $|\cdot|_\infty$ indicates the l_∞ norm (i.e., $\max(|\cdot|)$) and the columns of $\mathbf{V}_{\text{side}} = [\mathbf{v}(\Omega_1), \cdots, \mathbf{v}(\Omega_{K'})]$ are the calibration vectors associated with the K' side lobe angles where the beam pattern is to be minimized. The constrained minimax side lobe response beamformer is found by applying a commercial numerical optimization code (such as fmincon in the MATLAB Optimization Toolbox) to the objective function

$$\begin{gathered} \mathbf{w}_{\text{e}} = \arg\min_{\mathbf{w}} F(\mathbf{w}) \text{ subject to} \\ |\mathbf{v}^H(\Omega_k)\mathbf{w}| = c_k \;\; \{\forall\, k, 1 \le k \le J\} \text{ and } \angle[\mathbf{w}]_1 = 0 \end{gathered} \tag{4.3}$$

where J is the number of point response constraints applied to the beam.

This design approach appears best suited for simulation, where a dense grid of modeled calibrators over a wide field of view can be quickly and easily obtained. With a limited

FOV for valid calibrators in a real-world scenario, the equiripple constraint can only be applied to the first two far-field side lobes. However, it is shown that beamformer weights designed from modeled calibrations, even with a detailed electromagnetic model, cannot be directly applied to PAF data without significant distortions, and that even after applying a corrective transformation, better performance is achieved with measured calibrators. In either case, while the resulting beamformer provides a closely controlled beam pattern shape, the associated sensitivity penalty makes this method undesirable for typical PAF use.

4.1.3 Hybrid Beamformer

We propose a hybrid beamformer that combines the benefits of both the max-SNR and numerically optimized equiripple beamformers, providing maximum sensitivity for a given amount of beam shape control. The beamformer is obtained through the numerical optimizer described above by solving

$$\begin{aligned} \mathbf{w}_{\mathrm{h}} = \arg\min_{\mathbf{w}} U(\mathbf{w}) \text{ subject to}& \\ ||\mathbf{v}^H(\Omega_k)\mathbf{w}| - c_k| \le 2\gamma|a_k - c_k| \;\; \{\forall\, k, 1 \le k \le J\}& \\ \text{and } \angle[\mathbf{w}]_1 = 0& \end{aligned} \tag{4.4}$$

where $a_k = |\mathbf{v}^H(\Omega_k)\mathbf{w}_{\mathrm{m}}|$ are the magnitude of the max-SNR beam response in direction Ω_k, referenced to a unity gain main lobe peak and

$$\begin{aligned} U(\mathbf{w}) &= \gamma\frac{1}{\mathrm{SNR}(\mathbf{w})} + (1-\gamma)\left|\mathbf{V}_{\mathrm{side}}^H\mathbf{w}\right|_\infty, \\ \mathrm{SNR}(\mathbf{w}) &= \frac{\mathbf{w}^H\mathbf{R}_{\mathrm{s}}\mathbf{w}}{\mathbf{w}^H\mathbf{R}_{\mathrm{n}}\mathbf{w}}. \end{aligned} \tag{4.5}$$

The user specified weighting parameter $0 < \gamma < 1$ controls whether emphasis is placed on maximizing the array SNR [73], or on minimizing the equiripple side lobe levels. In order for $\gamma = 1$ and $\gamma = 0$ to produce the max-SNR and equiripple beamformers respectively, the equality pattern constraints of (4.3) were changed to γ-dependent inequality constraints. In (4.4), when $\gamma \to 0$ the right hand side inequality constraint $2\gamma|a_k - c_k|$ is forced to zero and acts just like the constraint in (4.3). As $\gamma \to 1$ the constraint region grows large enough around c_k to include the max-SNR solution with some excess "elbow room" to allow the optimizer latitude in its search path. The hybrid beamformer can be designed using either modeled or measured calibration data, however, as mentioned previously, it is shown that measured calibrators are the best choice for this method. This hybrid beamformer will also be shown to provide both beam pattern shape control and sufficient sensitivity for PAF operation.

4.1.4 Transforming Modeled Beamformers

Using modeled calibration vectors, deterministic beamformer design can be done relatively quickly, with constraints over an arbitrarily large field of view with a dense grid of points. However, these beamformer weights cannot be directly applied to real-world data because modeling inaccuracies lead to distortions in the final pattern. In situ calibration data is required to account for the variability of the operating environment.

Figure 4.1 verifies this conclusion. One half of both a modeled and measured equiripple beam pattern are shown side-by-side for comparison. The left half shows the modeled beam pattern resulting from an equiripple beamformer $\grave{\mathbf{w}}_{\mathrm{e}}$, designed using a very detailed electromagnetic model closely matching the real instrument (see the element pattern comparisons of Figure 4.2). Response constraints of unity gain at the main lobe peak and a gain of -36.5 dB at four additional points spaced 0.8° from boresight provide a null-to-null width of 1.6° (constraints placed on the side lobe or at the 3 dB point are just as reasonable). Note that the main beam is of uniform radius and the first side lobe offers 27 dB of attenuation and is of constant height. The right half figure results when measured calibration vectors $\hat{\mathbf{v}}(\Omega_k)$ are used to calculate a beam response for $\grave{\mathbf{w}}_{\mathrm{e}}$. The result is noticeable and undesirable distortion in the main beam and throughout the side lobes.

Such unacceptable results can be significantly improved by applying a transformation to map the modeled beamformer onto observed calibration data to correct for the inherent modeling inaccuracies. Even with a transformation, though, the modeled calibrators do not provide the best option for deterministic beamformer design. However, since the possibility of using modeled beamformers is reasonable and appealing, we present the background and

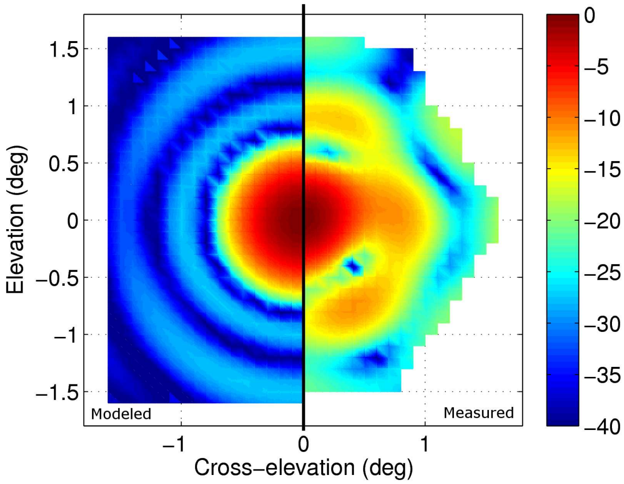

Figure 4.1: The effect of not performing a transformation before applying modeled beamformer weights to measured data. The left half is the modeled beam pattern. The right half is the measured beam pattern using the same modeled beamformer weights. There is noticeable distortion throughout the measured pattern.

results of this approach for completeness, even though it is shown to be impractical in the field.

The goal of a transformation is to find a mapping matrix $\mathbf{T}$ such that real data beamformer $\hat{\mathbf{w}}' = \mathbf{T}\grave{\mathbf{w}}$ produces a measured beam pattern closely matching the desired modeled designed beam pattern. A transformation of this kind could allow the advantages of designing modeled beamformers to be successfully exploited.

Let $\mathbf{b}^T = \mathbf{w}^H\mathbf{V}$ be the beam pattern vector, where columns of $\mathbf{V} = [\mathbf{v}(\Omega_1), \cdots, \mathbf{v}(\Omega_Z)]$ are steering vectors associated with the Z angles where pattern matching is to be enforced. Modeled and measured beam patterns are given by $\grave{\mathbf{b}}^T = \grave{\mathbf{w}}^H\grave{\mathbf{V}}$ and $\hat{\mathbf{b}}^T = \hat{\mathbf{w}}^H\hat{\mathbf{V}}$ respectively, with $\hat{\mathbf{V}}$ consisting of measured calibration vectors, and $\grave{\mathbf{V}}$ those of the model. The $\hat{\mathbf{w}}'$ that minimize the squared difference between the beam patterns are computed as

$$\begin{aligned}\hat{\mathbf{w}}' &= \arg\min_{\hat{\mathbf{w}}} ||\hat{\mathbf{w}}^H\hat{\mathbf{V}} - \grave{\mathbf{w}}^H\grave{\mathbf{V}}||_2^2 \\ &= (\hat{\mathbf{V}}^H)^\dagger\grave{\mathbf{V}}^H\grave{\mathbf{w}} \qquad (4.6)\end{aligned}$$

where $(\hat{\mathbf{V}}^H)^\dagger = (\hat{\mathbf{V}}\hat{\mathbf{V}}^H)^{-1}\hat{\mathbf{V}}$ is the left matrix pseudoinverse of $\hat{\mathbf{V}}^H$.

It is more important that the transformation preserves the shape of the higher gain main beam than the side lobe structure, but the least squares solution fails to do this since it gives equal weighting to all points throughout the beam pattern. The obvious solution is a weighted least squares treatment, incorporating diagonal weighting matrix **G** which determines relative priority for fitting points in the beam pattern. The weighted least squares version of (4.6), with diagonal weighting matrix **G**, is written as

$$\begin{aligned} \hat{\mathbf{w}}' &= \mathbf{T}\grave{\mathbf{w}}, \\ \mathbf{T} &= (\hat{\mathbf{V}}\mathbf{G}\hat{\mathbf{V}}^H)^{-1}\hat{\mathbf{V}}\mathbf{G}\grave{\mathbf{V}}^H. \end{aligned} \tag{4.7}$$

This result provides more control over the transformation with improved matching at the desired points and a least squares fit over the remainder of the pattern. The main beam structure can be satisfactorily preserved by increasing the weighting of just a few points within the main beam. For closest matching at the desired points, care should be taken to limit the number of weighted points in the beam pattern since there are a limited number of degrees of freedom provided by the beamformer weight vector.

4.2 Beamforming Results

Calibration data for the beamforming analysis was collected in 2008 on the Green Bank 20-Meter Telescope.

Great effort has been made to design a simulation model that closely mirrors the true PAF and operating environment for the experimental array and 20m dish [28]. Its accuracy is verified in Figure 4.2, where modeled and measured individual element power patterns representing an azimuth slice through the center of the astronomical source Cassiopeia A (Cas A) are shown to be in close agreement. However, as shown in Figure 4.1, this agreement is inadequate for the precision required for designing PAF beamformers in simulation.

The calibration data set, measured with Cas A, is a 33×33 grid of 10-second pointings spaced 0.1° apart and ranging from $-1.6°$ to $+1.6°$ in both elevation and cross-elevation (i.e., the arc direction perpendicular to elevation) with respect to boresight. An off-source measurement was taken at each elevation, 8° from boresight, as an estimate of the noise field

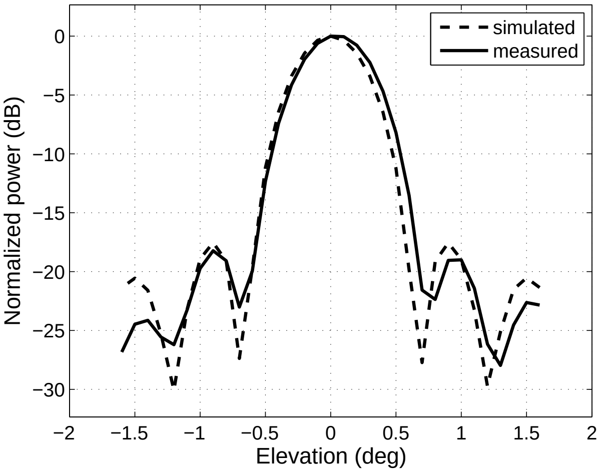

Figure 4.2: Modeled and measured power patterns of array element 1 (center element). The close agreement between the modeled and measured patterns is an indication of the accuracy of the simulation model.

corresponding to all calibrators at that elevation. The MDL algorithm described in Section 3.1 was used to identify useful calibrators (Figure 3.1).

4.2.1 Comparison of Beamformer Methods

To compare performance of these beamformers we first look at the differences in the structure of each corresponding beam pattern.

The max-SNR beam pattern is shown in Figure 4.3 and demonstrates the concerns discussed previously: nonuniform side lobes with a peak only 13 dB below the main lobe. Nevertheless, the structure is optimal for obtaining maximum sensitivity in the calibration noise field.

The modeled equiripple beamformer of Figure 4.1 was constrained and transformed as described in Section 4.1.4. The terms of the weighting matrix **G** corresponding to the five optimization constraint points were set to a value of 10^8 while the remaining diagonal terms were initialized to 1. This places a heavy penalty on deviations from the design equality constraint points. Figure 4.4 compares the measured beam pattern after transformation to the original modeled pattern. In the measured right half image, the shape of the beam main lobe is slightly distorted, there is significant variation in the side lobe structure and an increase

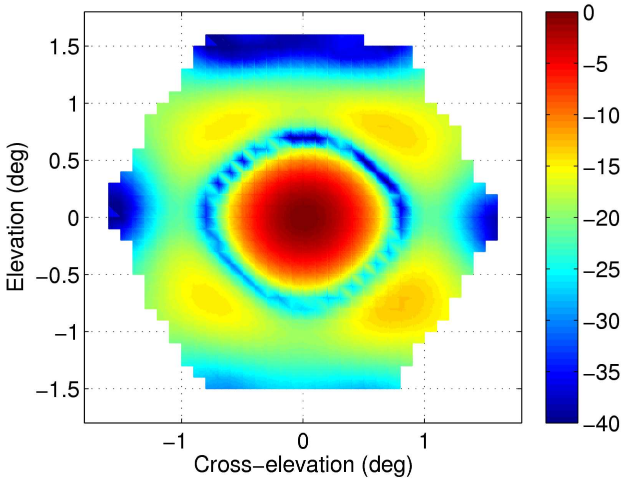

Figure 4.3: Measured far-field pattern of the max-SNR beamformer. The pattern exhibits high, uncontrolled side lobes.

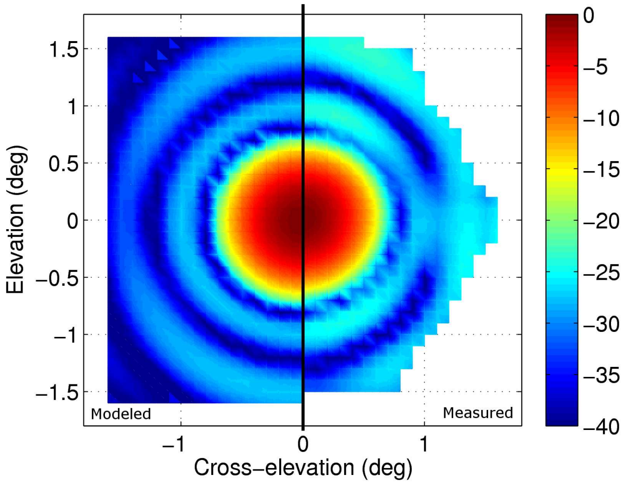

Figure 4.4: The effect of performing a weighted least squares transformation on the modeled beamformer weights before applying them to measured data. The left half of this figure is the modeled beam pattern. The right half is the measured pattern using the transformed modeled beamformer weights. The transformation causes some distortion to the pattern but makes modeled deterministic beamformers possible.

in the maximum side lobe level of 3-4 dB. Still, the shape much more closely resembles the desired pattern than when the modeled weights were applied without transformation.

The need for a transformation can be entirely eliminated by applying the measured calibrators directly in the optimization of (4.3). In Figure 4.5 the measured beam pattern

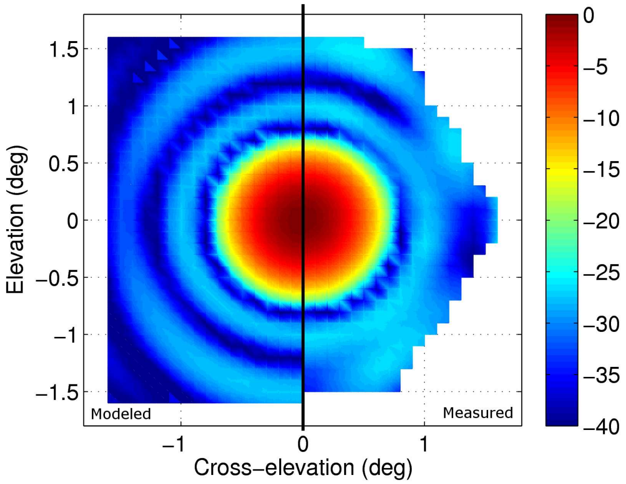

Figure 4.5: Deterministic beamformers designed with measured calibration vectors closely resemble the modeled version. The left half of this figure is the modeled beam pattern. The right half is the measured beam pattern using the measured beamformer weights from a numerical optimizer. There are only slight differences noted between the two patterns.

of the right half image conserves the main and side lobe shapes of the modeled pattern. It still does not perfectly match the modeled image on the left (especially at the second null), but this approach does appear to provide the best overall option for deterministic design. This raises the question of whether there is a need for modeled beamformer design, which is addressed in Section 4.2.3.

The hybrid beamformer, whose beam pattern with $\gamma = 0.25$ is shown in Figure 4.6, offers characteristics of both the max-SNR and equiripple beamformers. The influence of the equiripple beamformer is seen in the uniformity across the pattern, while the increased width of the side lobe is more characteristic of the max-SNR beamformer. The 21 dB side lobes are several dB higher than those of the equiripple approach, but the tradeoff for increased side lobe levels is a desirable increase in sensitivity.

A numerical comparison of the beamforming methods is given in Table 4.1. Results of the hybrid beamformer with $\gamma = 0.5$ and $\gamma = 0.25$ show that low side lobes can be achieved without fully sacrificing sensitivity. The hybrid beamformer with $\gamma = 0.5$ results in a 4.7 dB increase in side lobe suppression compared to the max-SNR case, with only a 4% reduction in sensitivity. Decreasing γ to 0.25 suppresses the side lobes an additional 3.3 dB and reduces the total sensitivity by 15%. When compared to the equiripple beamformer

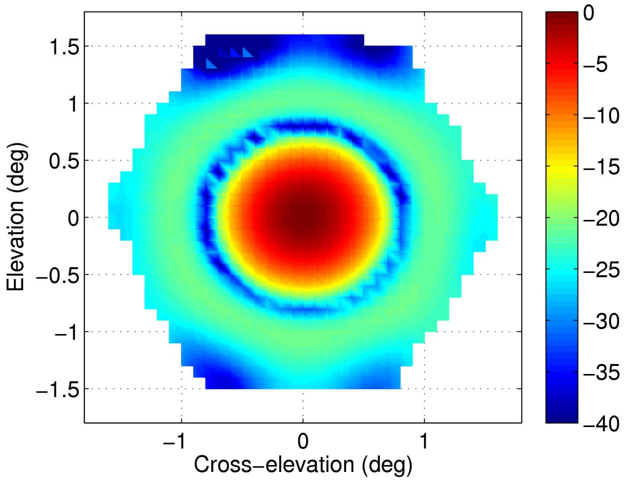

Figure 4.6: Measured far-field pattern of the hybrid beamformer with $\gamma = 0.25$. The circular shape of the main beam and the constant side lobe level are characteristics of the equiripple beamformer. The peak side lobe level is slightly greater than with the equiripple beamformer, but there is also a significant increase in sensitivity.

Table 4.1: Comparison of beamformer techniques.

Beamformer	Beamwidth	Peak side lobes	Sensitivity
max-SNR	1.6°	−13.03 dB	2.97 m^2/K
equiripple	1.6°	−26.40 dB	1.84 m^2/K
hybrid ($\gamma = 0.5$)	1.6°	−17.70 dB	2.86 m^2/K
hybrid ($\gamma = 0.25$)	1.6°	−21.02 dB	2.54 m^2/K

the hybrid approach represents a significant improvement in sensitivity, yet it maintains the desired uniform structure not provided by the max-SNR beamformer. Certainly, based on the information in Table 4.1, the hybrid beamformer is an interesting alternative for PAF operation.

Figure 4.7 shows results of the hybrid beamformer for all values of γ and the described set of constraints. As expected, as the value of γ approaches 1 the side lobes and the sensitivity both increase. The results of these figures can be used to determine the proper value of γ that should be used for designing a hybrid beamformer for a specific application.

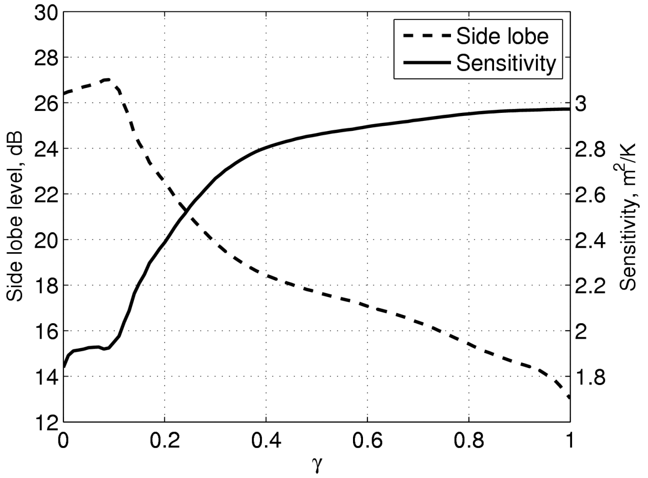

Figure 4.7: The ability to suppress the beam pattern side lobes is greatest for the equiripple beamformer corresponding to $\gamma = 0$, and gradually decreases as the hybrid beamformer approaches the max-SNR result. Similarly, the sensitivity of the hybrid beamformer is greatly improved as the beamformer becomes less deterministic.

4.2.2 LCMV Beamformer

With the ability to minimize output variance while meeting desired constraints, the LCMV beamformer appears to be a strong candidate for PAF beamforming.

This beamformer is easy to implement and only requires calibration vectors at the desired constraint points. However, due to the limited degrees of freedom available to the beamformer it is not possible to control the entire FOV and avoid undesired beam pattern structure. Introducing additional constraints to obtain more control uses degrees of freedom that are needed to minimize the variance, causing a decrease in sensitivity and distortions in the beam pattern.

Example beam patterns using the LCMV beamformer are shown in Figure 4.8(a) and 4.8(b). Figure 4.8(a) was constructed with five constraint points matching those described in Section 4.1.4. The constraints are met, but the remainder of the pattern is unpredictable. In order to mitigate the distortion four additional equally-spaced constraints were added to the first null. The result given in Figure 4.8(b) compares well with that of the hybrid beamformer with $\gamma = 0.5$, achieving a sensitivity of 2.86 m^2/K, but due to the limited constraints in the

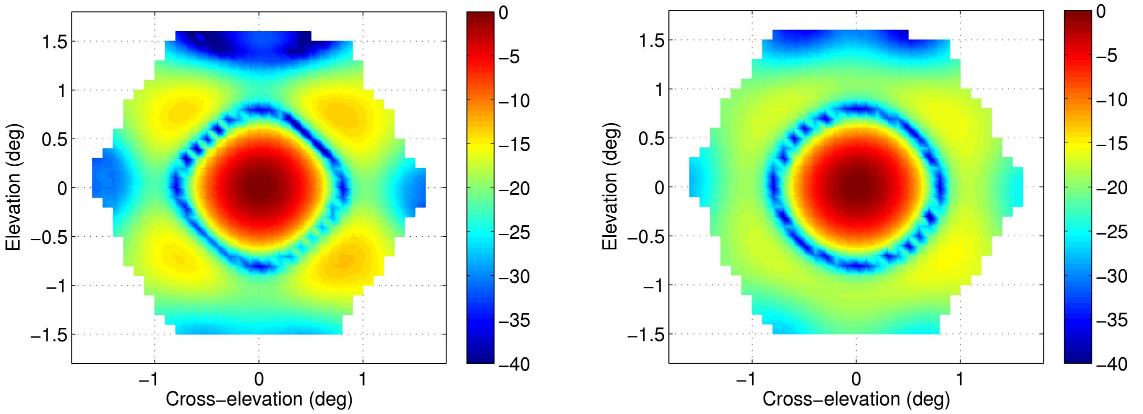

(a) LCMV with five constraints.

(b) LCMV with nine constraints.

Figure 4.8: Measured beam patterns of the LCMV beamformer with a main beam constraint plus additional constraints evenly spaced within the first null. (a) Four additional constraints are not adequate to provide the uniformly low side lobe structure given by the hybrid beamformer. (b) Eight additional constraints produces a pattern that closely matches that of the hybrid beamformer with $\gamma = 0.5$ but has higher side lobes and decreased sensitivity.

LCMV, there is noticeable variation in the side lobe structure and the peak side lobe level of 15.24 dB is nearly 2 dB higher than in the hybrid result.

4.2.3 Value of Modeled Beamformers

The transformation process described in Section 4.1.4 requires a set of measured calibrators to best preserve the desired beam pattern structure, but this process can be bypassed if we use the measured calibrators directly to design the beamformer. Still, there are some potential benefits of modeled design that make it an appealing consideration. Modeled design is useful if there is an advantage to (1) optimizing over a more dense grid of calibrators than is required for the transformation, or to (2) having calibration points on an increased span of angles. We explore these possibilities below.

1. *Sparse Transformation*: Using the MDL algorithm to remove poor calibrators from our highly oversampled 33×33 calibration grid, we are left with 817 of the available 1089 pointings. Thinning this grid to just 44 points reduces calibration time by about 3 hours, but a modeled beamformer transformation based on this sparse calibration set results in beam pattern distortion as seen in Figure 4.9(a). The left half pattern

is the result of a full 817 point transformation and the right half image is that of the more sparse. A significant difference in the sensitivity is also obtained: 1.74 m^2/K for the full transformation and 1.66 m^2/K for the sparse. The pattern distortion and sensitivity loss may be tolerable for the given reduction in calibration measurement time, but the result is still less desirable than when the sparse calibration set is directly used itself to design the beamformer. The equiripple beamformer computed directly with the sparse set of measured calibrators still exhibits distortion, but offers much better sensitivity, achieving 1.78 m^2/K. Because of the great dependence on the details of the transformation data set, any potential benefits of modeled dense calibration beamformer design are lost in the transformation process.

2. *Increased Angular Span*: Designing modeled beamformers over a larger angular region than can be covered with measured calibrators is only beneficial if the beam shape control in the extended region is not forfeited during the transformation procedure. The left half image of Figure 4.9(b) was again obtained after a transformation with a full calibration set, while the right half image was transformed with a reduced size grid of 19×19 points, bounded by the dashed line. The result is a loss of control of the side lobes outside the calibration region once a transformation is applied. Again we see that the benefits of modeled design are restricted by the need for, and the limits of the calibration set used in the transformation process.

Based on these results, we conclude that it is impractical to design measured deterministic PAF beamformers using simulation models. Deterministic PAF beamformers are best designed using measured calibration data directly, avoiding pattern distortion and sensitivity loss that accompany a transformation from the model.

4.2.4 Angular Limits of Pattern Control

As noted in Section 4.2.3 pattern shape control with deterministic beamformers is limited by the angular range of the calibration vectors used in the design process. This introduces concerns about the behavior of the beam pattern beyond the reach of good calibration vectors. The inherent dish directivity pattern, governed by the properties of the

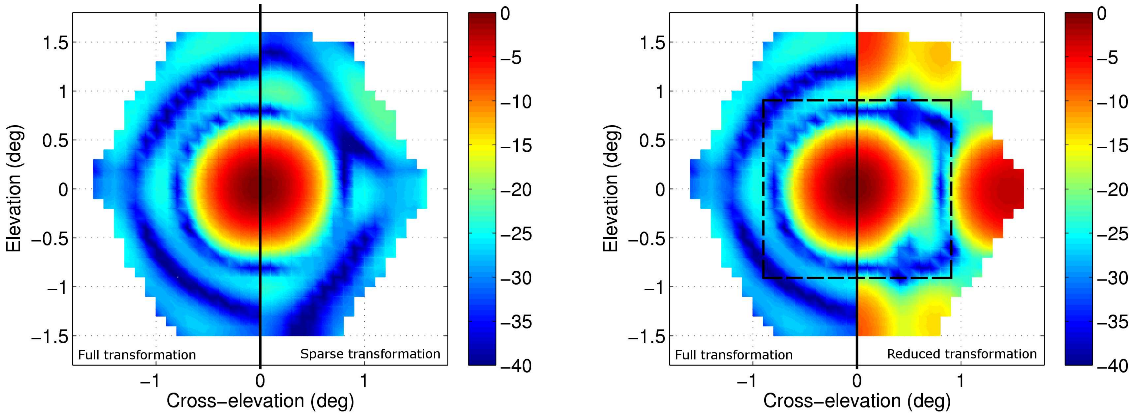

(a) Sparse transformation.

(b) Increased angular span.

Figure 4.9: (a) The original measured calibration set used in this analysis was highly oversampled. There is only a slight difference between the measured beam patterns after full (left half) and sparse (right half) transformations of the modeled beamformer, but the calibration time difference is about 3 hours. The sparse transformation introduces distortions and a decrease in sensitivity. (b) The beam pattern that extends beyond the calibration region cannot be controlled. A full transformation from the model gives the left half pattern and a reduced 19×19 transformation region (represented by the dash-lined box) gives the right half pattern. The modeled weights were computed based on a larger 33×33 grid of calibrators. The patterns are in relative agreement within the box, but not at all outside the box.

reflector dish, begins to dominate the combined array-dish pattern at some angle, after which we lose most control of the beam pattern shape. It is the region between the edge of the calibration set and the start of the dish aperture dominance that is of concern. Since we cannot measure good calibration vectors in this region of interest, we must draw conclusions through analysis of modeled results.

We are interested in knowing whether the range over which good calibrators can be obtained extends to the angle at which the dish pattern begins to the dominate the PAF far-field pattern. If it does not, as seen in Figure 4.9(b), there is a region of the beam pattern that is uncontrollable. To determine if the measured calibration range is large enough to fill this gap, we have computed modeled equiripple beamformers using both a 33×33 grid (matching that obtained in practice) and a larger 101×101 grid of calibration points (same grid point density in each case). Figure 4.10 shows a comparison of the resulting beam patterns plotted over $\pm 5°$ in both elevation and cross-elevation directions, with the 33×33 grid beamformer on the left and the 101×101 grid beamformer on the right. There is

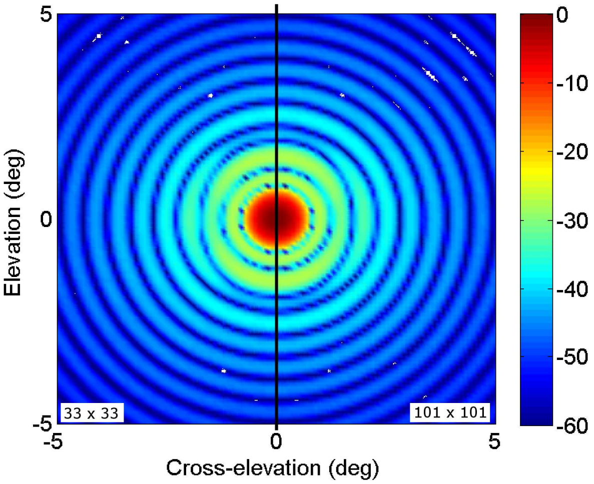

Figure 4.10: Comparison of modeled beam patterns of beamformers produced from different size calibration grids. Grids of size 33×33 and 101×101 were used to compute beamformers that generated the left and right half images, respectively. There is no visible difference between the patterns, indicating that there is no unreachable and uncontrollable region between the angular extent of a measured PAF calibration grid and the angle at which the dish aperture patterns begins to dominate.

no noticeable difference between the two images. Although in practice the calibration set is limited by the SNR of the calibration source, we conclude that it is adequate to provide control of all parts of the beam pattern that are not dominated by the dish aperture pattern.

4.3 PAF Beamforming Demonstration: Elevation Dependent Noise

Adaptive beamformers can be used to suppress undesired signals and noise and are updated in response to changes in the operating environment. There have been a number of publications discussing approaches to mitigating interfering signals using PAFs [72, 86–88]. This section explores the effects and suppression of the noise component, specifically elevation-dependent noise. Additive noise from the receiving system, the surrounding environment and atmosphere are collected by a radio telescope and make it difficult to detect weak astronomical sources. Reducing the noise provides an increase in sensitivity to signals of interest. Ideally, the noise field is uniform with respect to the steering angle of the dish. However, variations in the characteristics of the received spillover and atmospheric noise signals occur as a radio telescope moves position. These variations are most strongly noted

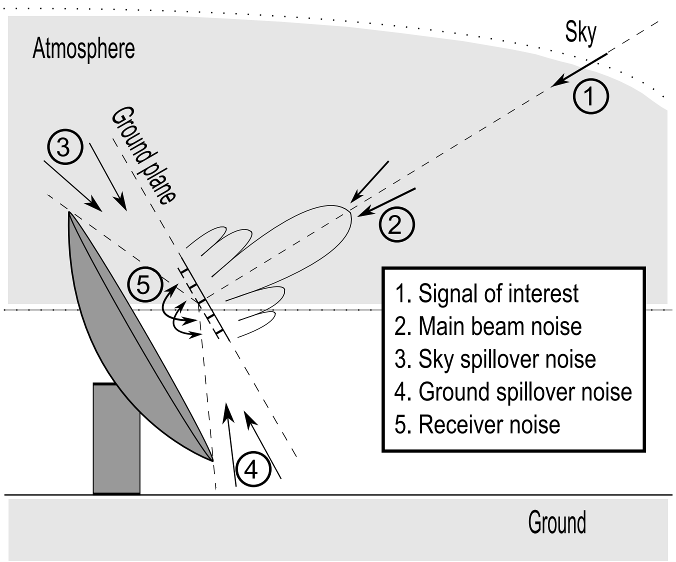

Figure 4.11: The main sources of noise that contribute to the signal received by a radio telescope include mutually-coupled receiver noise between antenna array elements, spillover noise originating from outside the edge of the reflector dish, and atmospheric noise collected by the main beam of the antenna radiation pattern.

with elevation changes while tracking a celestial source moving across the sky. Elevation-dependence of the noise can be mitigated through the use of noise-minimization adaptive beamforming. An example is provided in this section that demonstrates this capability as well as a discussion about how such an adaptive beamformer balances the contributions of each noise source to meet its objective. Ultimately it is shown that, although minimal, there is some benefit in response to elevation dependent noise that is gained by using a fully adaptive beamformer as opposed to a fixed-adaptive beamformer.

4.3.1 Components of the Noise Field

As depicted in Figure 4.11, the noise acquired by a radio telescope system is composed of three main sources: receiver noise, spillover noise, and main beam sky noise. (Additional loss comes from antenna elements but this is assumed negligible for the purposes of this demonstration.) While each of these makes a significant contribution to the overall system noise, the source that dominates and therefore drives the beamformer response depends greatly on the reflector dish elevation angle. These different noise components are discussed in Chapter 2.

4.3.2 Experimental Results

Initial results showing the elevation-dependence of PAF noise were based on data collected in November 2007 on the Green Bank 20-Meter Telescope [81]. Sixteen measurements were collected in five-degree increments for elevation angles between 85° down to 10° above the horizon. Care was taken to direct the dish toward a portion of the sky that was not known to contain any bright sources that could interfere with the experiment. It was initially anticipated that there would be a decrease in the system noise as the dish was tipped toward the horizon due to the changing background in the spillover region. In reality, the contribution of the main beam noise to the total system noise was greater than anticipated and resulted in a net increase as the dish approached the horizon. Results of this initial experiment are shown in Figure 4.12(a), where the beamformed noise power is shown with respect to elevation angle. The two plots compare the performance of the fixed-adaptive and adaptive MVDR beamformers discussed in Section 2.4. The temporary increase in measured power seen at 40° elevation may be due to an unexpected astronomical source in the field of view. It is clear that the adaptive beamformer more effectively suppresses noise as it adapts to changes in the noise environment. The fixed-adaptive beamformer is expected to provide the closest possible match to the adaptive beamformer of any fixed beamformer. If a truly deterministic beamformer, i.e., one that has no knowledge of current noise covariance structure, had been used the anticipated improvement obtained by the adaptive beamformer could be much greater in contrast. The value of $T_{\min}$ for the Mini-Circuits LNAs used in the receivers for this experiment was originally estimated to be between 105-120 K based on the specifications provided by the manufacturer and our own measurements [77]. More sophisticated measurements done by Dr. Leonid Belostotski at the University of Calgary found the true value of $T_{\min}$ to be about 85 K [74]. By comparison, Figure 4.12(b) shows the modeled results under similar conditions, using the full noise model described in Section 2.2 with $T_{\min} = 85$ K, $T_0 = 2$ K, and the ground temperature T_{gnd} set to 280 K.

In anticipation of more detailed analysis, the Green Bank 20-Meter Telescope was used again in 2008 to repeat the experiment. A total of 33 data points were collected with increments of 2.5 degrees from zenith down to 10° above the horizon. Receiver system upgrades prior to the 2008 experiments included new LNAs with a much lower noise figure.

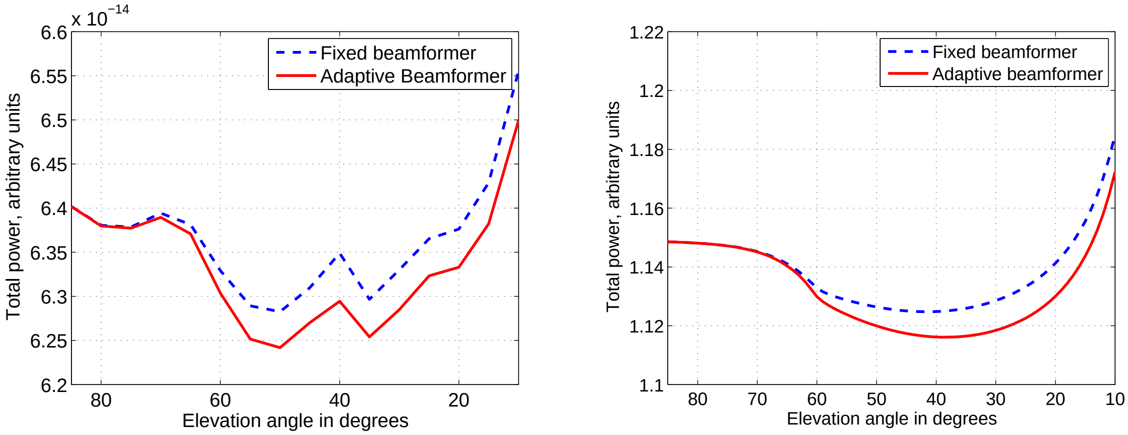

(a) Experimental results.

(b) Modeled results.

Figure 4.12: A comparison of experimental and modeled results showing the effect on the system noise levels by tipping a radio telescope toward the horizon. The model parameters were set to match those of the 2007 experimental system, including the approximated LNA noise value of 85 K. The decrease and subsequent increase in system noise as the reflector dish is tipped from zenith toward the horizon is consistent with the expected behavior of the spillover noise and main beam noise, respectively. The adaptive beamformer suppressed the noise better than the fixed-adaptive beamformer.

The T_{min} associated with these new Ciao Wireless LNAs has a measured value of 33 K [82], which was incorporated into the simulation model. Results of this more recent experiment are shown in Figs 4.13 and 4.14 along with modeled results for comparison. The results were obtained using the adaptive and fixed-adaptive MVDR beamformers. Figure 4.13(a) is in arbitrary units of total power while the remaining plots show equivalent temperature values in units of Kelvins. The conversion from power to temperature is computed by dividing by the isotropic noise power as in (2.8). For the case of a fixed beamformer, the isotropic power is constant. The noise temperature for the fixed beamformer is then just a scaled version of the fixed beamformer noise power.

It is clear from Figure 4.13(a) that the adaptive beamformer is again better able to suppress the noise at the different elevation angles as we saw in the previous result. Figure 4.13(b) shows modeled and measured T_{sys} values which help to provide a more complete understanding of the situation. Interestingly, the adaptive beamformer measures a higher T_{sys} relative to the fixed-adaptive beamformer as the elevation angle decreases. This counter-intuitive result can be explained by looking at the aperture efficiency and sensitiv-

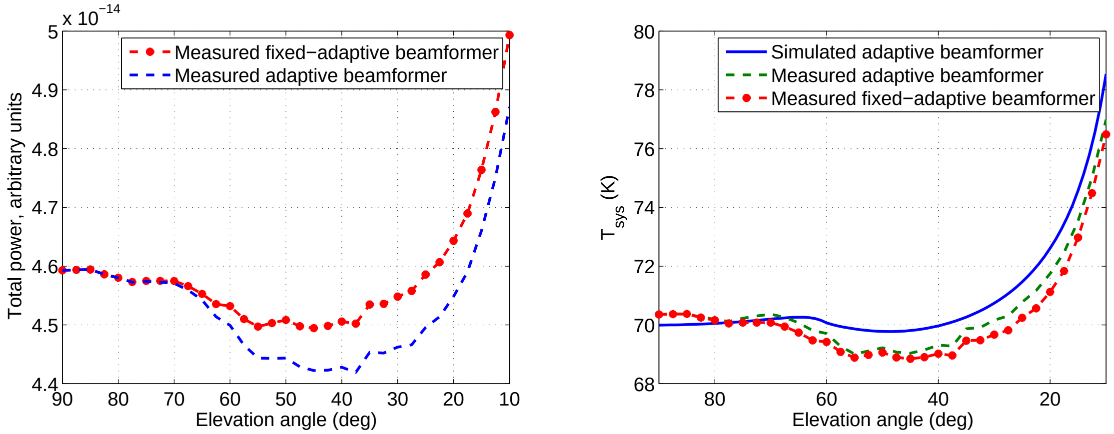

(a) Power.

(b) Temperature.

Figure 4.13: A comparison of the noise power and T_{sys} values for measured and modeled results of the 2008 system. For T_{sys} there is close agreement with the model. The model result was obtained using an adaptive MVDR beamformer, while both adaptive and fixed-adaptive MVDR beamformers were applied to the measured data.

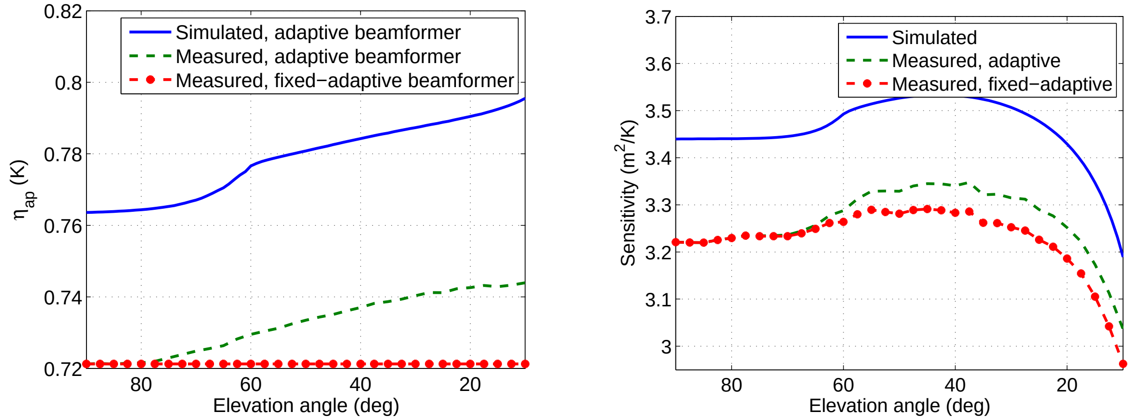

(a) Aperture efficiency.

(b) Sensitivity.

Figure 4.14: A comparison of the aperture efficiency and sensitivity for both measured and modeled results of the 2008 system. The model uses an adaptive MVDR beamformer, while both adaptive and fixed-adaptive MVDR beamformers were applied to the measured data. The measured and modeled results exhibit the same behavior and vary only by a scale factor.

ity variations as a function of elevation angle (see Figure 4.14). Since the beamformer is designed to maximize sensitivity, which is proportional to $\eta_{\mathrm{ap}}/T_{\mathrm{sys}}$ (see Section 2.5.2), the increase in T_{sys} indicates a corresponding increase in aperture efficiency. The aperture efficiency associated with the fixed-adaptive beamformer is included in the figure for reference, but remains constant since the weights are held fixed throughout the data set.

From these results it is clear that there is some improvement in the ability to respond to elevation dependent noise variations by using an adaptive beamformer. However, since the improvement is relatively small it can be concluded that a fixed-adaptive beamformer of this kind works well for PAF radio telescopes.

The significant agreement between the modeled and measured results shown here provides a confirmation of the accuracy of the simulation models. This allows for modeling of a variety of different scenarios that will further increase our understanding of how the adaptive beamformer acts to minimize noise power in response to changes in the structure of the noise fields.

4.3.3 Modeled Results

As shown in (2.7) the total system noise covariance matrix $\mathbf{R}_\eta$ is the sum of the spillover, receiver, and main beam noise covariance matrices, which are designated as $\mathbf{R}_{\text{rec}}$, $\mathbf{R}_{\text{sp}}$, and $\mathbf{R}_{\text{sky}}$ and which are defined in (2.9), (2.12), and (2.14) respectively. Since this is the case we can perform analysis on each one to understand how each component affects the total system noise temperature T_{sys}. Additionally, the simulation model provides a clearer understanding of the benefits of implementing an adaptive beamformer to efficiently respond to noise field variations. As with the analysis of the experimental data, the modeled results shown here offer a comparison between the adaptive and fixed-adaptive MVDR beamformers. The values of the model parameters are the same as those described in sec 4.3.2 with $T_{\text{min}} =$ 33 K. One difference however, is that the model incorporates active impedance matching [58] to minimize the effects of cross-coupling between the array elements which reduces the overall system noise.

Figure 4.15 provides a graphical comparison of adaptive and fixed-adaptive beamformers for each of the sources of noise mentioned previously and the total noise. The output of the fixed beamformer offers some intuitively pleasing results. There is no change in the receiver noise (Figure 4.15(a)) since both the beamformer weights and the input receiver noise power are independent of the direction that the dish is pointed. A decrease in

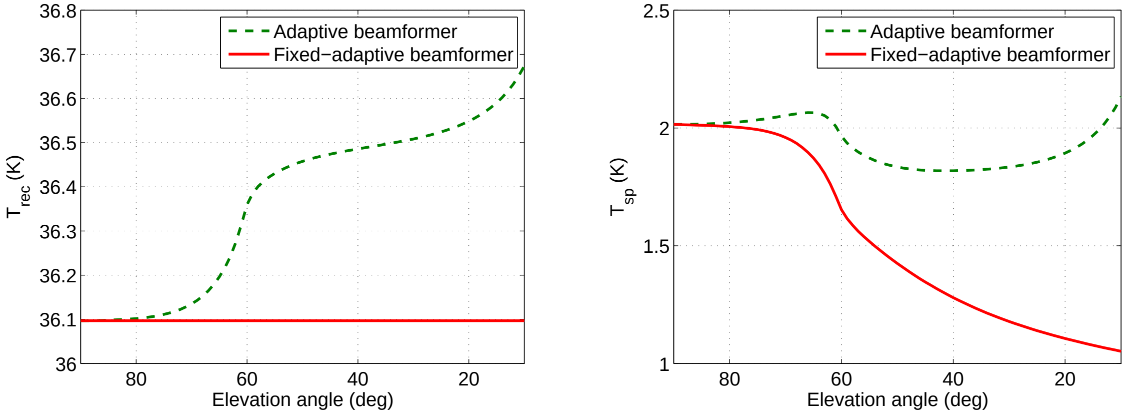

(a) LNA noise temperature.

(b) Spillover noise temperature.

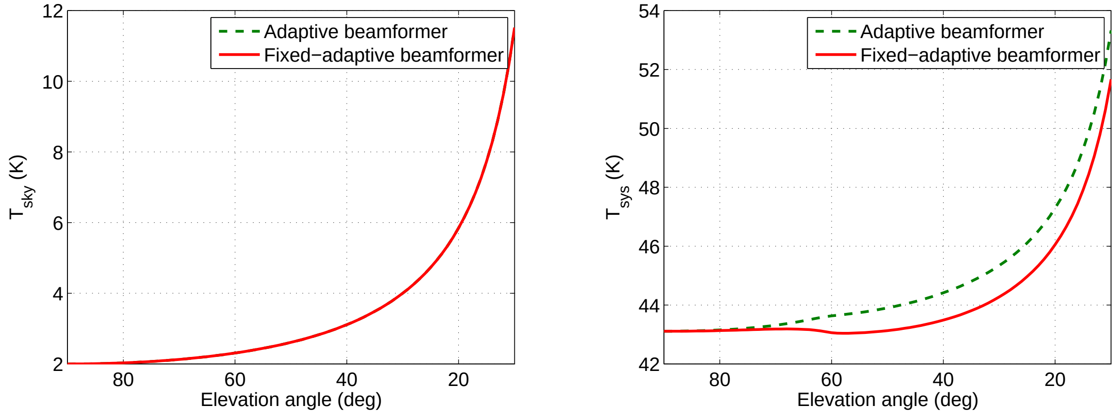

(c) Main beam noise temperature.

(d) Total noise temperature.

Figure 4.15: Noise temperature as a function of dish elevation angle for each individual noise source. The model assumes active impedance matching for the array and a T_{min} of 33K for each LNA. The adaptive beamformer (dashed line) varies its response to each noise source to improve sensitivity. The system noise temperature is greater for the adaptive beamformer case, indicating an improvement in aperture efficiency, which boosts the signal temperature and gives an increase in sensitivity. T_{sky} is the same for both adaptive and fixed-adaptive beamformers.

the received spillover noise occurs as more of the spillover region is redirected from a warm ground toward a cooler sky (Figure 4.15(b)). The main beam noise varies according to the secant curve described in Section 2.2.3, increasing as the dish elevation angle approaches the horizon (Figure 4.15(c)).

The results of the adaptive beamformer are not so predictable because it works to reduce the output power as much as possible given the current noise state. Receiver noise is independent of dish elevation angle, i.e., constant, implying that any adaptation results from

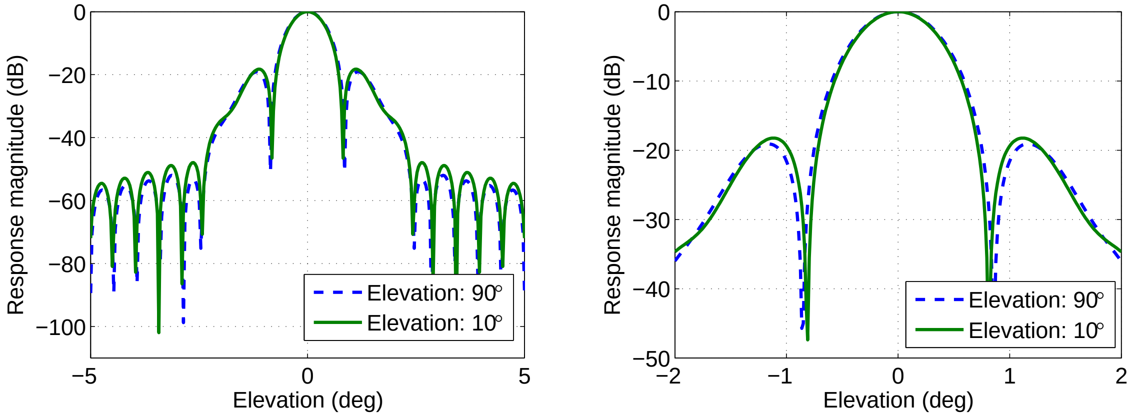

(a) Main beam and first side lobes.

(b) Closer look at the main beam.

Figure 4.16: The MVDR adaptive beamformer produces changes in the beam pattern as the reflector dish is tipped. As the adaptive beamformer responds to the increasing main beam noise the result is a more narrow main beam and higher side lobe levels. The plots compare the beam patterns at zenith and 10° above the horizon. $T_{\text{min}} = 33$ K.

attempts to counter the effects of variations in the other noise sources. Near zenith, the sky noise is quite small so the adaptive beamformer minimizes the total output power by reducing the spillover and receiver noise. As the dish tips toward the horizon, the increasing main beam sky noise becomes a more dominant factor in determining the beamformer weights. This is evident in Figure 4.15 by the fact that, below 70° elevation, both the receiver and spillover noise temperatures are greater than in the fixed-adaptive case. In order to minimize the output power due to the increasing sky noise, the adaptive beamformer reduces the width of the main beam and allows the side lobe levels to rise. This effect is apparent in the far field pattern shown in Figure 4.16. This does not, however, affect the main beam sky noise temperature since it is an extended source. Coincidentally, this action results in a net increase in the total system noise as the spillover and receiver noise levels rise (see Figure 4.15(d) for a comparison of T_{sys} values for each beamformer). The reduction in main beam width results from an effort to maximize sensitivity, which leads to increased aperture efficiency as seen in Figure 4.17(a). Figure 4.17(b) shows the overall benefit of adaptive beamforming as an increase in system sensitivity compared to the case of the fixed beamformer.

Comparisons to a single horn feed system are shown in Figure 4.18. Such a system is representative of a typical radio telescope in operation today. The horn feed is modeled

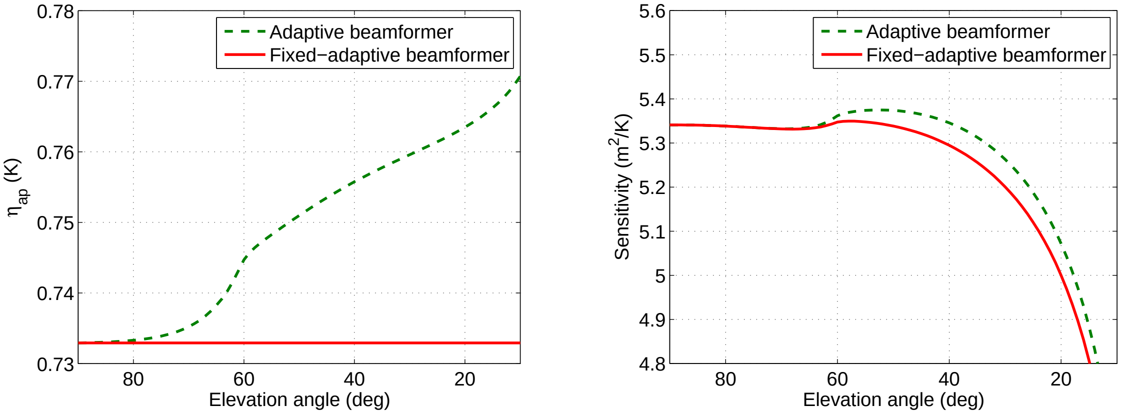

(a) Aperture efficiency.

(b) Sensitivity.

Figure 4.17: Modeled aperture efficiency and sensitivity for $T_{\text{min}} = 33$ K. The adaptive beamformer provides an improvement in aperture efficiency and an increase in sensitivity as it responds to changes in the noise field.

as having a $\cos^q(\phi)$ radiation pattern, with the value of $q = 2.28$ chosen to maximize the sensitivity at zenith, and ϕ defined as the angular offset from boresight. This comparison shows that the array with adaptive beamforming is capable of achieving a sensitivity at least as good as the single horn feed system.

Further insight can be obtained by studying the changes in the array illumination pattern. As shown in figure 2.2 and discussed in Section 2.2.2, the beam pattern of the array illuminates the surface of the reflector dish and the surrounding spillover region. The ideal array illumination pattern results in uniform gain across the reflector dish and zero gain in the spillover region. In practice, however, the best we can achieve is an approximation of this behavior. Figure 4.19 shows the modeled array illumination patterns of a 19 element with the 33 K LNAs, for zenith and 10° elevation. The illumination pattern is computed as an elevation cut from $-90°$ to $+90°$, with 0° being defined as the boresight direction. The pattern changes as the adaptive beamformer adjusts to the variations in the noise field associated with the tipping reflector dish. The symmetry of the patterns is due to the orientation of the array relative to the horizontal cut. The modeled array sees the same amount of spillover on both sides of the dish. As the reflector tips toward the horizon, the spillover noise decreases and the side lobes of the adaptive beamformer illumination pattern

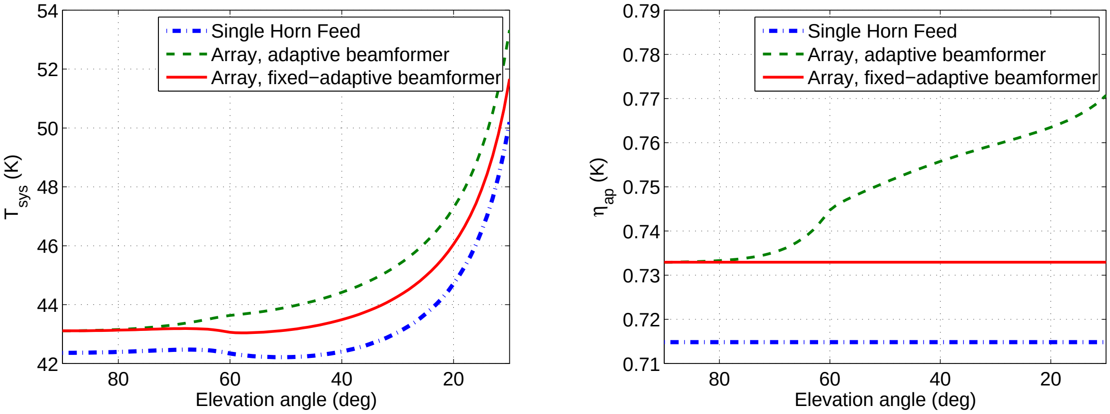

(a) Total noise temperature.

(b) Aperture efficiency.

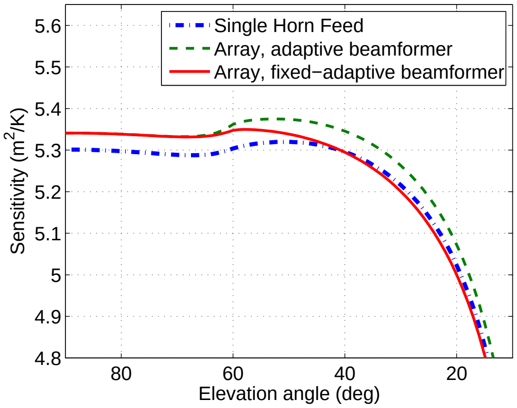

(c) Sensitivity.

Figure 4.18: Modeled comparison of the 19-element array to a single horn feed. The array performs at least as good as the horn over all elevation angles.

are allowed to increase in order to constrain the far field pattern in response to increasing sky noise.

4.4 Conclusion

PAFs for radio astronomical observations need beamformers that can provide both high sensitivity and beam pattern stability. The max-SNR and numerically optimized equiripple beamformers presented in this chapter represent the extreme cases: one offering the best sensitivity available and the other full beam pattern shape control. In order to satisfactorily manage the tradeoff between these conflicting goals we have introduced a

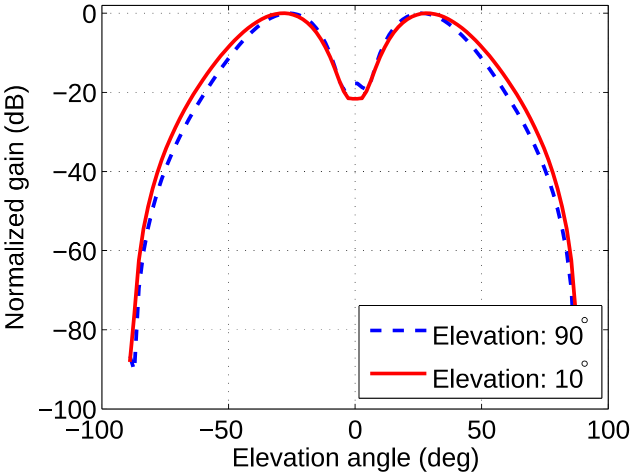

Figure 4.19: Modeled array illumination patterns for elevation angles of 90° and 10° above the horizon. As the dish tips the sky noise increases, the spillover decreases and the illumination pattern picks up more of the spillover region. This tradeoff allows for more attention to be given toward reducing the sky noise.

hybrid beamforming method. The hybrid beamformer uses a numerical optimizer and the weighting parameter γ to provide both adequate sensitivity and the necessary pattern control. A comparison between the max-SNR, equiripple, and hybrid beamformers has shown the usefulness of the hybrid approach.

Due to unpredictable PAF operating conditions which limit the accuracy of simulation models, beamformers designed using modeled calibration data introduce substantial distortions in the final measured beam pattern. A transformation step was introduced to mitigate these effects. Such a transformation requires a grid of measured calibration vectors, which can be used themselves to directly compute the desired beamformer. It has been shown that there is no benefit to designing beamformers with a model because the outcome of the required transformation is so highly dependent upon the limitations of the measured calibration data.

The max-SNR and numerically optimized equiripple beamformers presented in this analysis are representative of any number of beamformers that could be used in PAF systems. They were used as simple examples in demonstrating solutions to general PAF beamforming challenges.

Noise characteristics for a radio telescope are elevation-dependent. To recover the desired cosmic signals acquired by such a telescope, the strong noise field in which they are concealed must be suppressed as much as possible. Adaptive beamforming techniques consider the changes in the noise field and minimize its contribution to the received signal, increasing the SNR and allowing these faint astronomical signals to be retrieved. The three main contributors to the system noise are receiver noise, spillover noise, and main beam sky noise. While the receiver noise is a constant, the spillover and main beam sky noise vary as a function of elevation. As the incoming signal changes, the adaptive beamformer applies its available degrees of freedom toward balancing the output power contribution of each noise source in the most effective manner.

The accuracy of the model explained here, which was created to mimic the effects of the elevation-dependent noise, was verified through comparisons with experimental data. Using the model allowed for further analysis of the effects of the noise field and the capabilities of adaptive beamformers. The adaptation of the beamformer to elevation variations is most apparent in the change to the aperture efficiency and far field beam pattern of the radio telescope system, and the array beam patterns illuminating the reflector dish. While the benefits of adaptive beamforming are plentiful, the computational cost may be overly burdensome for large radio telescope systems. A fixed-adaptive beamformer, with pre-computed weights based on calibration data, could offer considerable computational savings while still providing some of the benefits of fully-adaptive beamformers.

Chapter 5

Weak Source Detection and Imaging

Since PAFs operate by electronically steering beams within the FOV, variations between PAF beams result in additional noise being added to the imaging measurements. This may obscure weak astronomical sources of interest that are already well below the noise floor. For a single feed system this is not an issue because the beam pattern is identical for each measurement. In order for PAFs to be competitive with a traditional system the beam-to-beam variations must be reduced to an acceptable level.

Single dish radio astronomical PAF images have been created on relatively bright sources [28,53,54,65]. The results in [28,53] show some artifacts of beam-to-beam variations. In these publications some mosaic images are presented combining the imaged FOVs for several pointings of the telescope. Hard transition lines are observed between the separate FOVs because of the variations between the beams.

When imaging with a multibeam receiver system, each of the beams forms a different far-field pattern on the sky with different signal and noise responses, which must be accounted for when processing the data. In [66] the differing signal responses in each pixel of an image from the ALFA array are smoothed by making several passes of the same part of the sky after adding a slight offset to the array pointing. Averaging multiple overlapping observations is also discussed in [67], where data from the Parkes 21-cm multibeam array are used to do a neutral hydrogen sky survey. Signal processing for a multibeam receiver is different from that of a PAF, but beam variations are a common problem in both systems and these examples confirm the need to address this problem in PAF imaging.

This chapter discusses reducing beam-to-beam variations in order to improve detection and imaging with a PAF-fed radio telescope. Variations in noise response between beamformers can be eliminated by a normalization to a local and contemporary noise field estimate. This estimate may be improved by averaging multiple measurements, which is shown to reduce the beam-to-beam variations. Normalization introduces a tradeoff between signal response variations and noise response levels between beams. A dual constraint beamformer design is presented, which is capable of maintaining uniform signal response while simultaneously meeting a uniform noise response constraint. This is particularly important for imaging purposes, when accurate relative power measurements are needed for each pixel. A rim constraint beamformer is introduced that uses eigenvector constraints to control the spillover response of the array dish illumination pattern and reduce the beam-to-beam variations. A combination beamformer that incorporates both the dual constraint and rim constraint techniques is shown (using a model as well as experimental data) to offer reduced beam-to-beam variations while meeting signal and noise response constraints.

5.1 Sinc-matched Beamformer

Stable and predictable beam patterns are needed to reduce beam variations. The deterministic equiripple beamformer presented in Section 4.1.2 is predictable, but it is not constrained appropriately to maintain pattern shape as the beams are steered within the FOV. Consequently, a beamformer has been designed to match a two-dimensional sinc pattern, which provided much greater stability among the steered far field beam patterns.

Like the equiripple beamformer, this deterministic beamformer utilizes an iterative optimization algorithm to solve a constrained minimization to construct a beamformer that produces a pattern which closely matches a two-dimensional sinc pattern. This method was chosen to provide stable, well-formed far-field patterns.

A pattern magnitude response equality constraint in direction Ω_k is specified as

$$|\mathbf{v}^H(\Omega_k)\mathbf{w}| = c_k \tag{5.1}$$

where c_k is the specified response referenced to a unity gain main lobe peak and set to the value of the corresponding point in the sinc function. Additionally, the phase of the first element in $\mathbf{w}$ is set to zero, eliminating an ambiguous degree of freedom. The least squares difference between the desired and actual beam pattern vectors for a given weight vector is given by

$$F(\mathbf{w}) = \Big|\Big|\, \mathbf{f}_{\text{sinc}} - |\mathbf{w}^H\mathbf{V}|\, \Big|\Big|_2^2 \tag{5.2}$$

where the columns of $\mathbf{V} = [\mathbf{v}(\Omega_1), \cdots, \mathbf{v}(\Omega_K)]$ are the calibration vectors associated with the K calibration angles, and $\mathbf{f}_{\text{sinc}} = [c_1, \cdots, c_K]$ is the vectorized desired magnitude sinc response. The complex valued sinc-matched beamformer $\mathbf{w}_{\text{sinc}}$ is obtained by numerical optimization of the objective function

$$\begin{gathered}\mathbf{w}_{\text{sinc}} = \arg\min_{\mathbf{w}} F(\mathbf{w}) \text{ subject to} \\ |\mathbf{v}^H(\Omega_k)\mathbf{w}| = c_k \;\; \{\forall\, k, 1 \le k \le C\} \text{ and } \angle[\mathbf{w}]_1 = 0\end{gathered} \tag{5.3}$$

where C is the number of point response constraints applied to the beam.

5.2 Noise Response Normalization

In general, PAF beamformers designed with a specified consistent signal response from beam to beam have varying noise responses. There are also relative gain differences in the receiver chains of the different array elements. These factors contribute to large variations of noise floor levels between pixels k, $1 \le k \le K$, in the observed field of view. This variation can be reduced by normalizing the beams to a local estimated noise field covariance $\hat{\mathbf{R}}_\eta$. To compute a normalized noise response beamformer $\tilde{\mathbf{w}}_k = \zeta_k \mathbf{w}_k$ from an arbitrary beamformer weight vector $\mathbf{w}_k$ to produce a desired noise response $\tilde{\mathbf{w}}_k^H \hat{\mathbf{R}}_\eta \tilde{\mathbf{w}}_k = p_\eta$, solve

$$(\zeta_k \mathbf{w}_k)^H \hat{\mathbf{R}}_\eta (\zeta_k \mathbf{w}_k) = p_\eta \tag{5.4}$$

to obtain the scale factor

$$\zeta_k = \sqrt{\frac{p_\eta}{\mathbf{w}_k^H \hat{\mathbf{R}}_\eta \mathbf{w}_k}}. \tag{5.5}$$

This normalization scales all K beamformers to have the desired response to the noise field described of $\hat{\mathbf{R}}_\eta$. It is assumed that the dish observed a source-free FOV with uniform (isotropic) noise distribution when estimating $\hat{\mathbf{R}}_\eta$, and that this noise distribution matches that of the noise floor in the formed beams for the field of interest. Normalizing in this manner achieves roughly the same purpose as on-off subtraction commonly used to remove noise contributions in single beam, single feed observations. The difference is that normalizing results in uniform noise response while subtraction maintains the uniform signal response.

5.3 Improving Relative Beam Variations

With a traditional radio telescope the pixels of a power image are obtained by directing the telescope in a raster scan pattern to each point in the sky for which a measurement is desired. This process can be improved by employing a phased array feed telescope as a "radio camera" to record multiple simultaneous measurements for each telescope pointing. In either case the ability to detect a weak astronomical signal is dependent on the ability to distinguish the source from the surrounding noise environment. Ideally the variations between measurements can all be attributed to a varying intensity of sources in the sky. However, this is not the case since spillover noise and main beam sky noise are both dependent on the pointing direction of the telescope as depicted in Figure 4.11.

When obtaining a single measurement per pointing, the beam pattern used to measure each pixel remains constant and the variability between pointings is primarily due to the changing spillover and sky noise. A PAF introduces an added amount of variability since steered beams each have differing far-field and reflector dish illumination patterns. An example of this problem is shown in Figure 5.1 where the beamformed power of a noise-only azimuthal slice through the sky exhibits much greater variation for electronically steered beams compared to physically steering the telescope. In both cases max-SNR beamformers are used and the beam-steered power for the kth pointing is computed as

$$p_k = \tilde{\mathbf{w}}_k^H \mathbf{R}_{\mathrm{x}} \tilde{\mathbf{w}}_k \tag{5.6}$$

where $\tilde{\mathbf{w}}$ is the noise response normalized version of $\mathbf{w}$ as described in Section 5.2. In order for PAFs to be considered for weak source detection and imaging, this additional relative variation across electronically steered beams must be reduced.

One option for further reducing beam variations is to average multiple noise field estimates to be used as the off-pointing reference $\hat{\mathbf{R}}_\eta$ in (2.19). As a radio telescope points

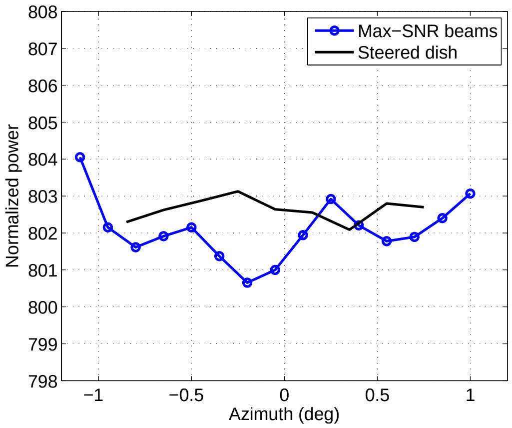

Figure 5.1: Comparison of steering the dish and steering beams through noise-only sky. There is much more variation in the electronically steered beams than there is when physically steering the telescope to each point.

to different parts of the sky there are changes in the noise field, specifically the spillover and main beam noise. For this reason it is difficult to get an accurate estimate of the true noise field corresponding to any particular place in the sky, which causes degradation in performance. It has been suggested that a detailed model of the spatial noise distribution may be used to improve the estimate of the true noise field by making a prediction based on an off-source noise measurement. An elevation noise model exists but it is insufficient to include azimuthal variations, nor does it explain all elevation variations. In lieu of such a model and method, techniques such as acquiring a noise estimate at the same elevation as the source, and averaging multiple noise measurements are used to improve the estimate. In addition to not having the true noise field in the off-source measurement, it is likely that a part of the sky that is believed to be source-free may in fact contain a source that will corrupt the noise estimate. By averaging across multiple noise fields the differences that exist between them are reduced while their common characteristics dominate, leaving an improved estimate of the true noise field.

Additional alternatives include designing beamformers that have a specified noise response and using calibrators on the rim of the reflector dish to control spillover response. The motivation for these techniques along with a description of each is presented below.

5.3.1 Dual Constraint Beamformer

A common constraint in PAF beamformer design is to require a specified signal response in the main beam. However, this produces a set of beams with different noise responses which causes variability between pixels in the power images and may conceal weak sources of interest. Normalizing the beams as described in Section 5.2 causes the noise responses to be the same but the tradeoff is significant variations between the signal responses. This is troublesome because radio camera images have artifacts due to signal response variations and absolute differences between the resulting pixels in an image cannot be measured.

We design a beamformer that provides both specified signal and noise response. The dual constraint beamformer is the solution to

$$\mathbf{w}_{\mathrm{c}} = \arg\min_{\mathbf{w}} \mathbf{w}^H\mathbf{w}, \text{ subject to} \tag{5.7a}$$

$$\mathbf{C}^H\mathbf{w} = \mathbf{f} \text{ and } \mathbf{w}^H\mathbf{R}_\eta\mathbf{w} = p_\eta \tag{5.7b}$$

where $\mathbf{C}$ is the $P \times D$ constraint matrix whose columns are the response vectors corresponding to the D directions in which a constraint is desired and $\mathbf{f}$ is a vector of the associated constraint values. The main beam response is specified by the linear constraints of (5.7b), which may also be used to constrain the beam pattern in additional directions. A noise power quadratic constraint is used to dictate the overall noise response of the beamformer. Since the noise variance is constrained in this way, minimum variance cannot be used as the optimization criterion as it is in the LCMV solution [38]. This linearly-constrained variance-constrained problem is ill-posed and the minimum norm optimization criterion was selected to provide a unique solution without additional complexity. The dual constraint beamformer is then the minimum norm solution that simultaneously satisfies the main beam response and noise response constraints. The problem is simplified by applying the techniques of the generalized side lobe canceller [38] to remove the linear constraints. This gives

$$\begin{aligned}\mathbf{w}_{\mathrm{c}} &= \mathbf{w}_0 + \mathbf{v} \\ &= \mathbf{w}_0 + \mathbf{C}_{\mathrm{n}}\mathbf{w}_{\mathrm{n}}\end{aligned} \tag{5.8}$$

where

$$\mathbf{w}_0 = (\mathbf{C}^H)^\dagger \mathbf{f} \tag{5.9}$$

is the least squares solution ensuring the constraint is satisfied and $\mathbf{v} = \mathbf{C}_\mathrm{n}\mathbf{w}_\mathrm{n}$ is orthogonal to $\mathbf{w}_0$; $\mathbf{C}_\mathrm{n}$ is a $P \times (P - D)$ matrix spanning the null space of $\mathbf{C}^H$ and is obtained through the singular value decomposition (SVD) or some other method; and $(\mathbf{C}^H)^\dagger = (\mathbf{C}\mathbf{C}^H)^{-1}\mathbf{C}$ is the left matrix pseudoinverse of $\mathbf{C}^H$. This leaves $\mathbf{w}_\mathrm{n}$ as the only unknown, which is obtained by solving

$$\mathbf{w}_\mathrm{n} = \arg\min_{\mathbf{w}} [\mathbf{w}_0 + \mathbf{C}_\mathrm{n}\mathbf{w}]^H [\mathbf{w}_0 + \mathbf{C}_\mathrm{n}\mathbf{w}] \tag{5.10a}$$

$$\text{subject to } [\mathbf{w}_0 + \mathbf{C}_\mathrm{n}\mathbf{w}]^H \mathbf{R}_\eta [\mathbf{w}_0 + \mathbf{C}_\mathrm{n}\mathbf{w}] = p_\eta \tag{5.10b}$$

where $(P - D) \times 1$ $\mathbf{w}$ is the dummy variable of optimization.

The problem is simplified by temporarily removing the known constant $\mathbf{w}_0$ from the constraint which eliminates the cross terms, leaving

$$\mathbf{w}_\mathrm{n} = \arg\min_{\mathbf{w}} [\mathbf{w}_0 + \mathbf{C}_\mathrm{n}\mathbf{w}]^H [\mathbf{w}_0 + \mathbf{C}_\mathrm{n}\mathbf{w}] \tag{5.11a}$$

$$\text{subject to } [\mathbf{C}_\mathrm{n}\mathbf{w}]^H \mathbf{R}_\eta [\mathbf{C}_\mathrm{n}\mathbf{w}] = p_\mathrm{n} \tag{5.11b}$$

where

$$\begin{gathered} p_\mathrm{n} = p_\eta - \mathbf{w}_0^H \mathbf{R}_\eta \mathbf{w}_0 - \\ \mathbf{w}_0^H \mathbf{R}_\eta \mathbf{C}_\mathrm{n} \mathbf{w}_\mathrm{n} - \mathbf{w}_\mathrm{n}^H \mathbf{C}_\mathrm{n}^H \mathbf{R}_\eta \mathbf{w}_0 \end{gathered} \tag{5.12}$$

which is obtained by expanding (5.10b). Since p_n depends on $\mathbf{w}_\mathrm{n}$, it is initially treated as an unknown arbitrary constant used to solve for the form of $\mathbf{w}_\mathrm{n}$. With this parameterization of the problem a solution is later obtained for the jointly optimum $\mathbf{w}_\mathrm{n}$ and p_n which satisfy (5.10a) and (5.10b).

Since $\mathbf{w}_0$ and $\mathbf{C}_\mathrm{n}$ are orthogonal, the cross terms of (5.11a) are zero and the method of Lagrange multipliers leads to the generalized eigenvalue equation

$$\mathbf{C}_\mathrm{n}^H \mathbf{C}_\mathrm{n} \mathbf{w} = \lambda \mathbf{C}_\mathrm{n}^H \mathbf{R}_\eta \mathbf{C}_\mathrm{n} \mathbf{w} \tag{5.13}$$

where λ is the Lagrange multiplier and eigenvalue parameter, and the solution does not depend on the value of p_n. Since the arguments of the eigenvalue problem are Hermitian symmetric, the eigenvalues are all non-negative. The eigenvector of interest then, is the one that corresponds to the minimum eigenvalue of the matrix pencil $\{\mathbf{C}_\mathrm{n}^H \mathbf{C}_\mathrm{n}, \mathbf{C}_\mathrm{n}^H \mathbf{R}_\eta \mathbf{C}_\mathrm{n}\}$ since the optimization function is being minimized.

The minimum eigenvalue λ_min from (5.13) corresponds to the maximum eigenvalue $\rho_\mathrm{max} = 1/\lambda_\mathrm{min}$ when (5.13) is reformatted as

$$\mathbf{C}_\mathrm{n}^H \mathbf{R}_\eta \mathbf{C}_\mathrm{n} \mathbf{w} = \rho \mathbf{C}_\mathrm{n}^H \mathbf{C}_\mathrm{n} \mathbf{w}. \tag{5.14}$$

This yields

$$\check{\mathbf{w}}_\mathrm{n} = \wp \left\{ \mathbf{C}_\mathrm{n}^H \mathbf{R}_\eta \mathbf{C}_\mathrm{n}, \mathbf{C}_\mathrm{n}^H \mathbf{C}_\mathrm{n} \right\} \tag{5.15}$$

where $\check{\mathbf{w}}_\mathrm{n}$ is proportional to the optimal weight vector solution to (5.11) for an unspecified power constraint p_n and $\wp\{\cdot\}$ is the operator that identifies the principal eigenvector of a matrix pencil, i.e., the eigenvector associated with the largest eigenvalue of the generalized eigenvalue problem (5.14).

Since eigenvectors are indeterminate to within a single complex scale factor, the $\mathbf{w}_\mathrm{n}$ which satisfies (5.11b) is

$$\mathbf{w}_\mathrm{n} = \mu \check{\mathbf{w}}_\mathrm{n} \tag{5.16}$$

where μ is chosen to satisfy constraint function (5.11b), which with substitution is

$$\left[\mathbf{C}_\mathrm{n} \mu \check{\mathbf{w}}_\mathrm{n})\right]^H \mathbf{R}_\eta \left[\mathbf{C}_\mathrm{n} \mu \check{\mathbf{w}}_\mathrm{n})\right] = p_\mathrm{n} \tag{5.17}$$

yielding

$$\mu = \sqrt{\frac{p_\mathrm{n}}{\check{\mathbf{w}}_\mathrm{n}^H \mathbf{C}_\mathrm{n}^H \mathbf{R}_\eta \mathbf{C}_\mathrm{n} \check{\mathbf{w}}_\mathrm{n}}}. \tag{5.18}$$

Expansion of the constraint (5.10b) yields

$$\begin{split}\mathbf{w}_0^H\mathbf{R}_\eta\mathbf{w}_0 + \mathbf{w}_0^H\mathbf{R}_\eta\mathbf{C}_\mathrm{n}\mathbf{w}_\mathrm{n} + \\ \mathbf{w}_\mathrm{n}^H\mathbf{C}_\mathrm{n}^H\mathbf{R}_\eta\mathbf{w}_0 + \mathbf{w}_\mathrm{n}^H\mathbf{C}_\mathrm{n}^H\mathbf{R}_\eta\mathbf{C}_\mathrm{n}\mathbf{w}_\mathrm{n} = p_\eta.\end{split} \tag{5.19}$$

The last term on the left side of (5.19) is recognized to be p_n and (5.16) and (5.18) are substituted into (5.19). Treating p_n as the unknown, the quadratic formula is used to obtain

$$p_\mathrm{n} = -\beta + \frac{1}{2}\sqrt{\beta^2 + 4\left(p_\eta - \mathbf{w}_0^H\mathbf{R}_\eta\mathbf{w}_0\right)} \tag{5.20}$$

where

$$\beta = \frac{\mathbf{w}_0^H\mathbf{R}_\eta\mathbf{C}_\mathrm{n}\check{\mathbf{w}}_\mathrm{n} + \check{\mathbf{w}}_\mathrm{n}^H\mathbf{C}_\mathrm{n}^H\mathbf{R}_\eta\mathbf{w}_0}{\sqrt{\check{\mathbf{w}}_\mathrm{n}^H\mathbf{C}_\mathrm{n}^H\mathbf{R}_\eta\mathbf{C}_\mathrm{n}\check{\mathbf{w}}_\mathrm{n}}} \tag{5.21}$$

which is then used to obtain μ from (5.18), $\mathbf{w}_\mathrm{n}$ from (5.16), and finally $\mathbf{w}_\mathrm{c}$ from (5.8) using only known constant values to complete a closed-form solution to (5.7).

5.3.2 Spillover Control with Rim Calibrators

Deterministic beamformers designed for PAF use are generally computed to produce a desired far-field response with secondary consideration given to the array pattern that illuminates the dish surface. The illumination pattern in the side lobe region may not be well controlled and can vary significantly from beam to beam. Since statistically optimal beamformers are designed to minimize the total noise response, they only indirectly control this part of the array pattern to reduce spillover noise and illuminate the dish in an efficient manner. In either case differences between the illumination patterns of distinct beamformers contribute to added variability in noise levels between pixels in the FOV since each beam is observing the non-isotropic and potentially highly structured spillover noise field through a different spillover illumination pattern structure. Incorporating specific reflector dish illumination pattern constraints is challenging because a) it is difficult to obtain a set of response vectors that also consider interaction with the telescope feed support structure and b) the beamformer has only a limited number of degrees of freedom to constrain the illumination pattern while meeting the previously desired design specifications. We propose implementing eigenvector constraints [38, 89] to minimize the illumination pattern response at the edge of the reflector dish using response vectors measured with calibration sources mounted around the rim of the reflector.

Illumination pattern response vectors can be obtained as in (2.17) using signals transmitted by sources located at the dish rim. The eigenvector approach allows us to control response over multiple points using a reduced-rank constraint matrix. The max-SNR beamformer can be modified to include illumination pattern eigenvector constraints as follows.

A constrained optimization alternative definition of (2.19) can be shown to be

$$\mathbf{w}_{\mathrm{m}} = \arg\min_{\mathbf{w}} \mathbf{w}^H \mathbf{R}_\eta \mathbf{w} \text{ subject to } \mathbf{C}^H \mathbf{w} = \mathbf{f} \tag{5.22}$$

where $\mathbf{C}$ for this special single constraint case contains only one response vector $\mathbf{v}_{\mathrm{s}}$ to steer a beam in the desired direction with a response level of $\mathbf{f} = 1$. The general solution to this problem for an arbitrary number of constraints is

$$\mathbf{w} = (\mathbf{R}_\eta)^{-1} \mathbf{C} [\mathbf{C}^H \mathbf{R}_\eta^{-1} \mathbf{C}]^{-1} \mathbf{f} \tag{5.23}$$

which is the LCMV solution of Section 2.4.1.

To effectively control or reduce spillover illumination, a relatively large number M, of constraint points are needed. The procedure described below can be used to significantly reduce the corresponding number of linear constraints required in $\mathbf{C}$.

For M rim calibrator sources construct the illumination pattern constraint matrix $\mathbf{D}$ whose columns are the response vectors for signals coming from the direction of the rim sources such that $\mathbf{D} = [\mathbf{v}(\Omega_1), \cdots, \mathbf{v}(\Omega_M)]$, with corresponding response values in vector $\mathbf{g}$. Compute the SVD of $\mathbf{D}$ to obtain

$$\mathbf{D} = [\mathcal{U}_{\mathrm{r}} \, \overline{\mathcal{U}}_{\mathrm{r}}] \begin{bmatrix} \Sigma_{\mathrm{r}} & 0 \\ 0 & \overline{\Sigma}_{\mathrm{r}} \end{bmatrix} \begin{bmatrix} \mathcal{V}_{\mathrm{r}}^H \\ \overline{\mathcal{V}}_{\mathrm{r}}^H \end{bmatrix} \tag{5.24}$$

where Σ_{r} comprises the Q largest singular values of $\mathbf{D}$ and $\mathcal{U}_{\mathrm{r}}$ and $\mathcal{V}_{\mathrm{r}}$ contain the corresponding left and right singular vectors, respectively. According to the structure of the constraint equation in (5.22)

$$\mathbf{D}_{\mathrm{r}}^{H}\mathbf{w} = \left(\mathcal{U}_{\mathrm{r}}\Sigma_{\mathrm{r}}\mathcal{V}_{\mathrm{r}}^{H}\right)^{H}\mathbf{w} = \mathbf{g} \tag{5.25}$$

and the reduced rank constraint equation corresponding to the rim calibrator sources becomes

$$\mathcal{U}_{\mathrm{r}}^{H}\mathbf{w} = \Sigma_{\mathrm{r}}^{-1}\mathcal{V}_{\mathrm{r}}^{H}\mathbf{g} = \mathbf{g}_{\mathrm{r}}. \tag{5.26}$$

The rim constraint beamformer is then obtained according to (5.23) with

$$\begin{aligned} \mathbf{C} &= [\mathbf{v}_{\mathrm{s}}\ \mathcal{U}_{\mathrm{r}}], \\ \mathbf{f} &= \begin{bmatrix} 1 \\ \mathbf{g}_{\mathrm{r}} \end{bmatrix} \end{aligned} \tag{5.27}$$

where $\mathbf{v}_{\mathrm{s}}$ steers a beam in the desired direction with unity gain.

Setting $\mathbf{g} = 0$ avoids the need to specify a phase value, which uses an additional degree of freedom. Also, the rim sources are positioned just beyond the edge of the reflector dish to improve illumination of the dish surface.

An obvious improvement on the dual constraint beamformer design is to add more linear constraints to (5.7). If the eigenvector constraints in (5.26) are used, a combination beamformer is produced that provides uniform signal response, uniform noise response, and reduced beam-to-beam variation through control of the dish illumination pattern. The downside is that the far-field pattern of the dual constraint beamformer is mostly uncontrolled. Additional constraints are needed to ensure that it does not give an undesirable pattern, but with a limited number of available degrees of freedom these will need to be chosen intelligently.

5.4 Experimental Results

The results presented here are based on data collected in 2011 at the NRAO facility in Green Bank, WV. Experiments were conducted on the 20-Meter dish, using the 19-element Carter array shown in Figure 3.9.

Three rim sources were mounted at the base of three of the four dish feed support arms (see Figure 3.12). This is not ideal placement for computing rim constraints with $\mathbf{g} = 0$ as in (5.25) since it limits the illumination of the dish surface. Also, $M = 3$ is too small for effective control. This configuration was originally used for the calibration stability tests and was not specifically setup for computing the rim constraint beamformer. The dish rim is at a 60° angle from the vertex with respect to the array and the support arms are estimated to be at about 45°. Each antenna was pointed directly at the bare array and aligned with the polarization direction of the array elements for maximum efficiency. A signal transmission from each of the horns was measured as part of the array calibration. This involved sending a series of sine-wave calibration signals from a single function generator, through a remotely-controlled coaxial switch used to select each of the horns in sequence. From these data, response vectors were obtained which allowed us to incorporate array pattern constraints.

Normalizing a set of beamformer weights to have a specified response to a local noise field as described in Section 5.2 is necessary to reduce the beam-to-beam variations in order to observe anything in the field of view. There is considerable improvement obtained when the noise is normalized out, but as shown in Figure 5.2 there is still a substantial amount of variation. The image shows power measurements of a noise field after max-SNR beams are normalized for uniform response, based on a separate nearby noise field. If the estimated noise field matched the imaging noise field then the image appears “flat,” but the variations affirm that they are not. It is impossible to detect a weak source in this field of view whose intensity is on the same level as these variations. Although necessary, this normalization generally has the undesirable side effect of producing non-uniform signal responses in the main beam. This effect is shown in Figure 5.3, which compares the original max-SNR, normalized max-SNR, and dual constraints beams. The original max-SNR beams have uniform unity response main beams, but the response to noise varies between beams. After normalization the noise response is uniform but the variation now appears in the main beam signal response. Whether the signal or the noise responses vary, the result is an inability to identify the absolute power levels in a PAF image. The dual constraint beamformer is able to provide both specified signal and noise response and is therefore more useful for PAF imaging.

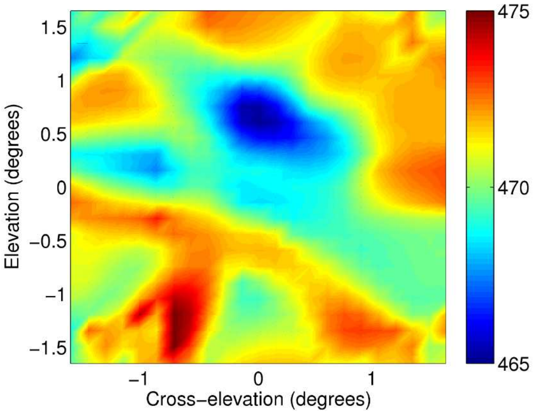

Figure 5.2: Image of noise field after normalizing max-SNR beams to a separate nearby noise field (linear scale). The image is flat if the noise fields were the same, but the variation in the image shows that there is a considerable amount of difference between the two fields.

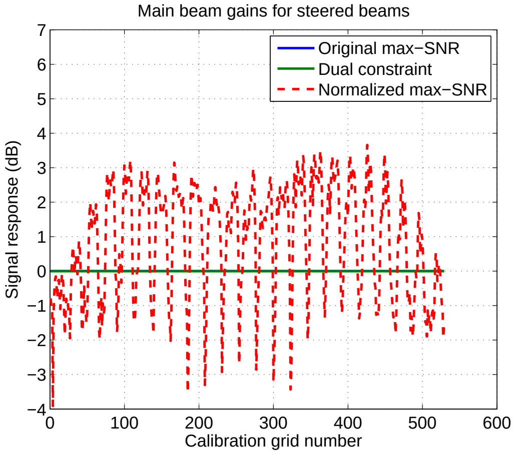

Figure 5.3: Signal response for different beamformers. The max-SNR beamformers are designed to have unity response in the direction of the main beam. Normalizing the beams to have the same noise response causes differences between the signal response of each beam. The noise response constrained beamformer achieves a uniform noise response across all beams while maintaining unity signal response.

To offer another perspective into the problem of varying main beam signal responses let us consider presenting the max-SNR problem of (5.22) as

$$\mathbf{w}_{\mathrm{m}} = \arg\max_{\mathbf{w}} \mathbf{w}^H \mathbf{v}_{\mathrm{s}} \text{ subject to } \mathbf{w}^H \mathbf{R}_\eta \mathbf{w} = p_\eta \tag{5.28}$$

which maximizes the main beam gain after meeting the specified noise constraint, i.e., it normalizes the unity gain max-SNR beams to have the same response to $\mathbf{R}_\eta$. Applying the method of Lagrange multipliers, the solution to (5.28) is

$$\mathbf{w}_\mathrm{m} = \sqrt{\frac{p_\eta}{\mathbf{v}_\mathrm{s}^H \mathbf{R}_\eta \mathbf{v}_\mathrm{s}}} \mathbf{R}_\eta^{-1} \mathbf{v}_\mathrm{s}. \tag{5.29}$$

The form of this solution is equivalent to that obtained by solving the eigenvalue problem (2.20) for a rank one matrix $\hat{\mathbf{R}}_\mathrm{s}$. The difference is that the solution of (2.20) scales each beamformer to have the same main beam response while (5.29) scales each beamformer to provide the same noise response.

It is unsettling to think that as the different beams are steered around the sky the signal response of each is different. For the purposes of detection this may be acceptable since maximum sensitivity is desired to observe very weak sources, but for imaging an extended space object it is important to have both fixed signal and noise response. To further demonstrate this point a comparison of the max-SNR and dual constraint beamformers shows the effect of the main beam response variations and noise response variations. Figure 5.4 shows the results of the max-SNR beamformers designed with a local noise field estimate and to provide uniform signal response. On the left, Figure 5.4(a) is a simulated image with arbitrary units, used for convenience in this demonstration. Each pixel in the image has been multiplied by the signal response of the beam that points in the corresponding direction. With uniform signal response between the max-SNR beams there is no distortion seen in this image. On the right, Figure 5.4(b) shows observed noise field power measurements obtained by using the uniform signal response max-SNR beams (i.e., there was no normalization to a nearby noise field). Note the significant spatial structure in this noise floor.

Figure 5.5 shows the same set of plots after normalizing the beamformers to the local noise estimate. The normalization affects the signal responses significantly and causes distortion in the simulated image. However, the noise response is now relatively flat. Figure 5.6 shows the performance of the dual constraint beamformer. The simulated image remains undistorted and a flat measured noise field is achieved.

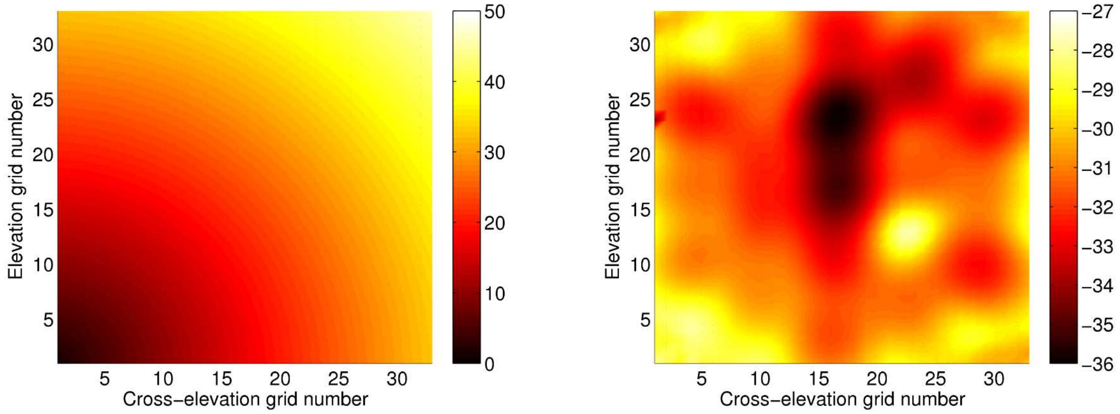

(a) Signal response image.

(b) Noise field image.

Figure 5.4: Signal response and noise field images for the uniform signal response max-SNR beams. (a) shows simulated variations across a field of view that are undistorted by the uniform responses between beams. (b) is a measured noise field power image showing significant variation between the noise response of each beam (dB scale). Ideally the noise response is flat and the image is a solid color.

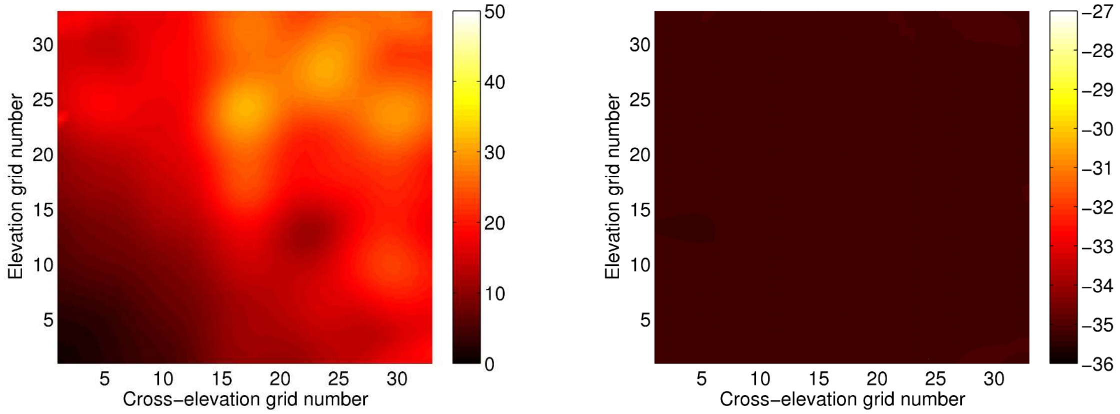

(a) Signal response image.

(b) Noise field image.

Figure 5.5: Signal response and noise field images for the normalized noise response max-SNR beams. (a) shows distortion in the image due to the signal response variations of each beamformer. (b) is the noise field image that now appears relatively flat after the beams were normalized to a separate nearby noise field (dB scale).

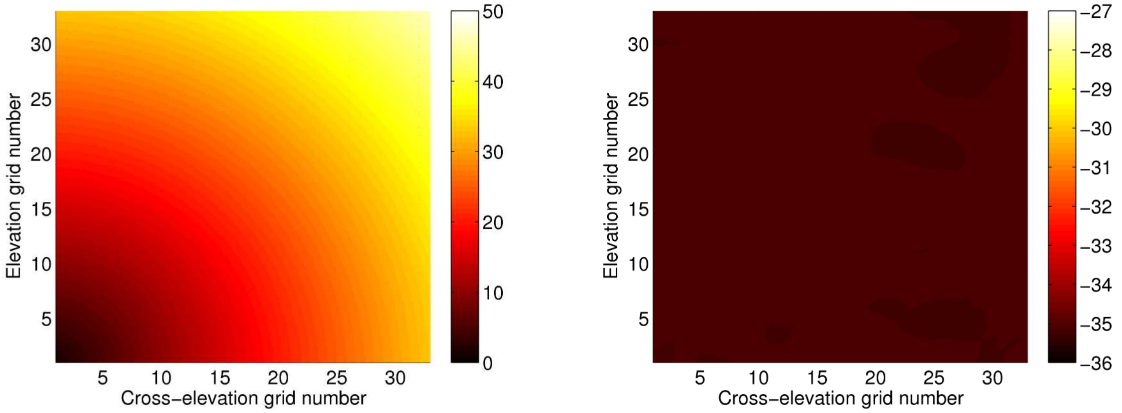

(a) Signal response image.

(b) Noise field image.

Figure 5.6: Signal response and noise field images for the dual constraint beams. (a) shows simulated variations across a field of view that are undistorted by the uniform signal responses between beams. (b) is a measured noise field power image that appears relatively uniform throughout the field of view (dB scale). The benefit of the dual constraint beamformer is that it can provide both uniform signal response and flat noise response.

As explained in Section 5.3, the variations present after noise response normalization can be reduced by averaging multiple noise field estimates to flatten the variations depicted in Figure 5.2. The result of averaging two noise field estimates is shown in Figure 5.7. Measured power is shown for both max-SNR and deterministic sinc-matched beamformers, as well as the "steered dish" measurements obtained with the boresight pointed max-SNR beam. For the "steered dish" result the telescope was steered through a slice of 9 points in a source free region of the sky to simulate performance of a single horn feed telescope scanning through the same sky pixels covered by the electronically steered beams. In the case of the steered beams the telescope was pointed to the middle slice point and then beams were electronically steered in azimuth at a constant elevation. The noise fields used for averaging were located at ±5° in azimuth from the middle slice point. The figure shows that averaging reduces both the amount of noise and the variability between the steered beams. This same behavior is observed for all of the beamformers considered in the paper. Table 5.1 quantifies the variation between the power measurements in the figure, showing the standard deviation of the measurements in units of Kelvin. Averaging two noise fields results in a 25% improvement for the max-SNR beamformers and a 50% improvement for

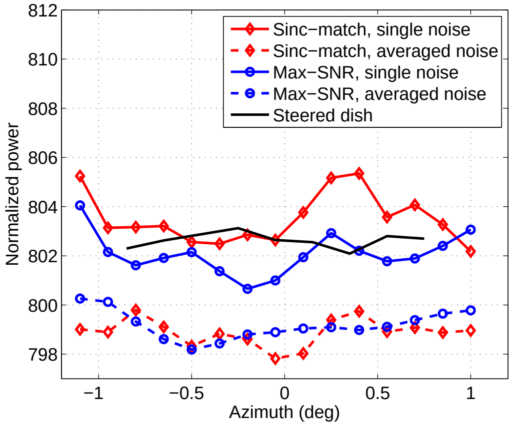

Figure 5.7: Averaging multiple off-steered (outside the FOV) noise field estimates reduces the amount of observed noise and the standard deviation of the steered beam measurements.

the sinc-matched beamformers. While this represents a significant reduction in the standard deviation of the measurements, it still does not compete with that of the steered dish.

We previously hypothesized that at least some of the variability between beams was the result of using the max-SNR beamformers, since max-SNR beams do not have any pattern constraints and are free to adjust the beam response in order to minimize noise. For this reason the sinc-matched beamformer was developed to provide a known and spatially stable beam pattern. However, as seen in Figure 5.7 the sinc-matched beamformers do not perform any better than the max-SNR beamformers. This suggests that the problem lies with the array pattern as it illuminates the dish and spillover region, and was the motivation behind the development of the spillover controlled beams described in Section 5.3.2. The fixed beams were not designed to have any spillover control while the max-SNR is working to reduce the spillover noise, which explains their comparable performance. As expected,

Table 5.1: Standard deviation of noise measurements in Kelvins.

Beamformer	**Standard Deviation (K)**
Sinc-matched, single noise	0.091
Sinc-matched, averaged noise	0.048
Max-SNR, single noise	0.074
Max-SNR, average noise	0.052
Max-SNR, steered dish	0.027

the “steered dish” max-SNR beamformer has the least variations, and is representative of performance of a conventional single horn feed dish. This is the desired target level so that a PAF radio camera can compete with a conventional telescope scanned image.

A comparison of the different beamformer designs presented in Section 5.3 is shown in Figure 5.8. The rim constraint beamformer was developed using all of the singular vectors of (5.24) with a desired response of zero at the rim. Since there were only three available rim sources, the compression possible with the eigenvector methods was unnecessary. Later simulation results show the benefit of the reduced rank approach using more rim sources. The figure plots the standard deviation in Kelvins of the normalized beam response of all the beams within a one square degree area in the center of the calibration grid (64 different beams). All beams were normalized to the first point of a 51 point slice through a 10 Jy source with each point separated by 0.2°. Standard deviation across the slice is computed for each set of beams (corresponding to the different beamformers) steered around that point. The dual constraint beamformer exhibits the largest variation, but this is reduced by adding rim constraints to get the combination beams. The dual constraint beamformer does not necessarily improve the beam variability but offers a fixed response. The rim constraint beamformer generally performs better than the max-SNR beamformer. This is surprising considering the small number of rim sources used in the experiment, but it is encouraging and suggests that such a beamformer design could be useful for PAF imaging. It may be especially beneficial if implementing additional rim sources reduced the variation of the combination beamformer. It is shown in the simulation results that follow that indeed the combination beamformer with more rim constraints performs very well. Another improvement is likely to come from additional linear constraints on the dual constraint beamformer design in order to better control the far-field pattern.

The results of Figure 5.8 indirectly show a varying sensitivity (SNR) between the beams as they pass over the source. By definition the max-SNR beamformers provide maximum sensitivity with respect to the chosen off-pointing and they provide a reference to consider the sensitivity penalty that results from using any of the other suggested beamforming approaches. Sensitivity values associated with the different beamforming techniques is provided in Table 5.2 for both boresight directed beams and beams steered off boresight by 0.9°.

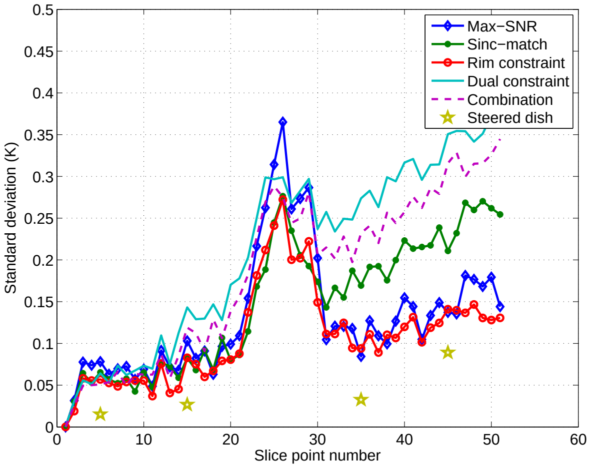

Figure 5.8: Standard deviation of beam powers for different beamformers as a function of distance from normalization noise field. The slice passes over a 10 Jy source. The max-SNR beamformer provides the greatest sensitivity but there is some improvement seen in the beam variations by constraining the spillover noise response. The steered dish variation is shown as well for comparison. The steered dish standard deviation was computed over a set of points to which the dish was steered, and only the boresight max-SNR beam was used. The slice points that include the source were excluded.

It is interesting to note that the sensitivity for the rim constraint beamformer remains about the same between the two cases while all of the others experience a significant decrease, which may suggest that the rim constraint beams experience the least amount of change as they are steered. For the boresight case, using a dual constraint beamformer with uniform signal and noise response only results in a 0.9 dB loss in sensitivity.

Table 5.2: Comparison of measured sensitivity values (m^2/K)

Beamformer	**Boresight beam**	**Steered beam** (0.9°)
Max-SNR	3.17	2.98
Sinc-matched	2.35	2.07
Rim constraint	2.50	2.53
Dual constraint	2.38	1.85
Combination	2.38	1.85

5.4.1 Far-field Beam Patterns

When comparing different beamformers it is usually helpful to examine their respective beam patterns to understand any oddities that have been introduced while meeting the desired beamformer objective. Strange or unpredictable behavior in the beam pattern is generally undesired since it may affect the sensitivity of the beams and limit the accuracy of measurements within the FOV. This section discusses the differences between the far-field beam patterns corresponding to the beam formers included in this analysis.

Figure 5.9 shows boresight far-field beam patterns for the max-SNR, sinc-matched, rim constraint, dual constraint, and combination beamformers. As is usually the case the max-SNR pattern is non-symmetric and largely unpredictable. The side lobes are at a reasonable level and knowing that the beamformer is achieving maximum sensitivity helps to mitigate any reservations about their nonuniform structure. The sinc-matched beam pattern is more acceptable. With a nearly symmetric appearance and a low, uniformly level side lobe, it is the only one of the beamformers that offers a predictable and predetermined pattern. The rim constraint, dual constraint, and combination beam patterns are all relatively similar with unpredictable variations and an obvious lack of control. The high side lobe point in the upper right of the dual constraint pattern is disturbing, but is suppressed in the combination beam pattern. Additional linear constraints in (5.7) could be used to provide more control in the far-field pattern and address this concern. Interestingly, the rim constraint and combination beam patterns look very similar, which highlights the relationship between the array illumination pattern and the far-field pattern, and suggests that much of the noise variation can indeed be attributed to differences in spillover response.

5.4.2 Weak Source Images

One of the main goals of this research is to improve the ability of a PAF to detect weak sources. Without the right data set it is not possible to demonstrate whether this is achieved. However, a comparison of the imaged weak source fields provides some idea of how the different beamformers perform. The source at the center of the images in Figure 5.10 is 3c309, which has an intensity value of 7 Jy. For the 20-Meter Telescope this source is seen as a point source and it is not possible to image any fine detailed structure. An image

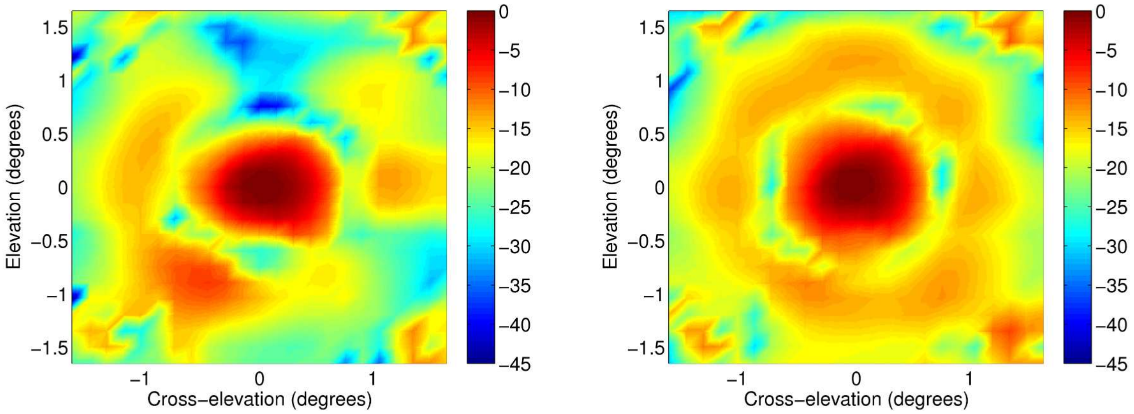

(a) Max-SNR beam pattern.

(b) Sinc-matched beam pattern.

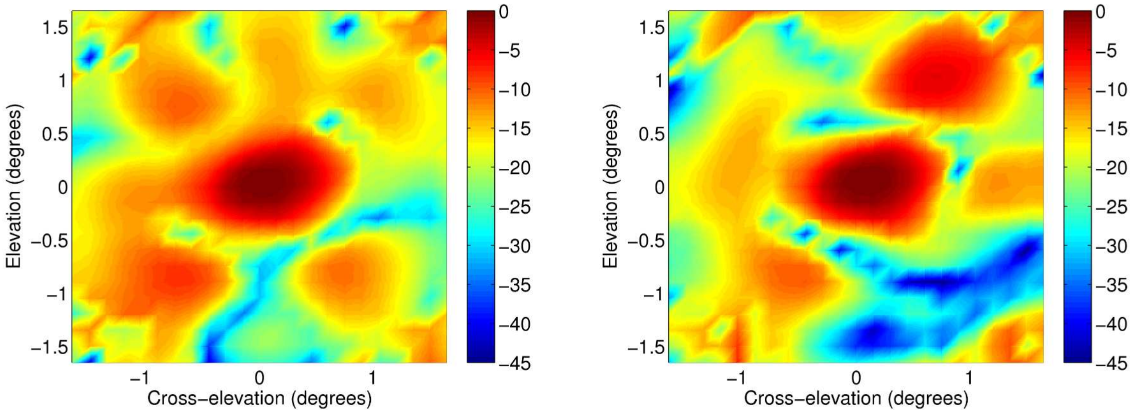

(c) Rim constraint beam pattern.

(d) Dual constraint beam pattern.

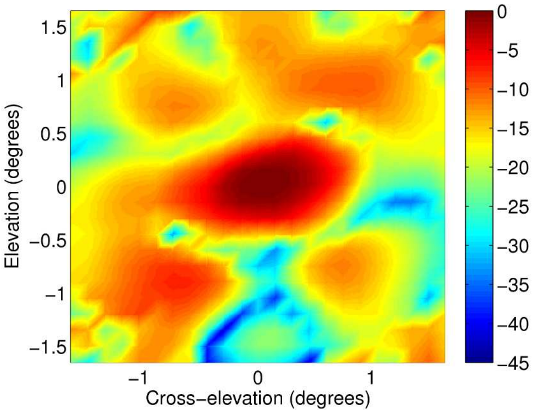

(e) Combination beam pattern.

Figure 5.9: Comparison of boresight beam patterns for max-SNR, sinc-matched, rim constraint, dual constraint, and combination beamformers. The sinc-matched is the only beamformer designed specifically to meet a desired far-field pattern structure. The dual constraint pattern has an undesirably high side lobe that may be eliminated with additional linear constraints. The rim constraint and combination beam patterns are very similar.

is presented using sets of beams corresponding to each of the five beamformers discussed previously. In each case the beamformers were normalized to a nearby noise field positioned at 5° away in azimuth from the source.

The max-SNR and rim constraint beamformers are able to clearly identify the source in the middle of the FOV. The max-SNR image does appear to measure a slightly higher value for the source power but at the same time is reports higher values in the noisy region surrounding the source. It is actually unclear whether the structure surrounding the source is real or simply noise, which highlights the problem of relative beam variations and weak source identification. The sinc-matched, dual constraint, and combination beamformer images exhibit an increased amount of noise. The source is visible in each case but only because it is expected to be there based on the results of the max-SNR and rim constraint beamformers. There is also a decreased sensitivity in each of these three beamformers, which is apparent from the lesser amount of power detected in the direction of the signal of interest.

As suggested in Section 5.3 the beamformers can be improved by normalizing to an averaged noise estimate. In the case of Figure 5.11 two noise field estimates were averaged and used in the normalization. The data used for these images and those of Figure 5.10 is part of a 15-point slice across 3c309 with two noise measurements obtained at ±5° from the source. Because of the rotation of the Earth there is a slight elevation change from the beginning of the slice to the end. This variation leads to a change in the spillover and main beam sky noise (see Section 4.3). The two noise fields vary in both elevation and azimuth and together provide an improved estimate of the noise field seen at the various points of the slice. These images obtained with the averaged noise estimate show a great improvement over the single noise measurement results. For all beamformers the source is clearly visible at the center of the image and the noise artifacts are significantly reduced.

For reference, the standard deviation between beams when normalizing with averaged noise estimates is shown in Figure 5.12. The improvements noted between the image plots above is seen here also in comparison to the single noise equivalent plot of Figure 5.8. The averaging provides a reduced amount of variation in the middle segment of the plot. To verify that the improvement by averaging is not just a matter of reduced estimation error

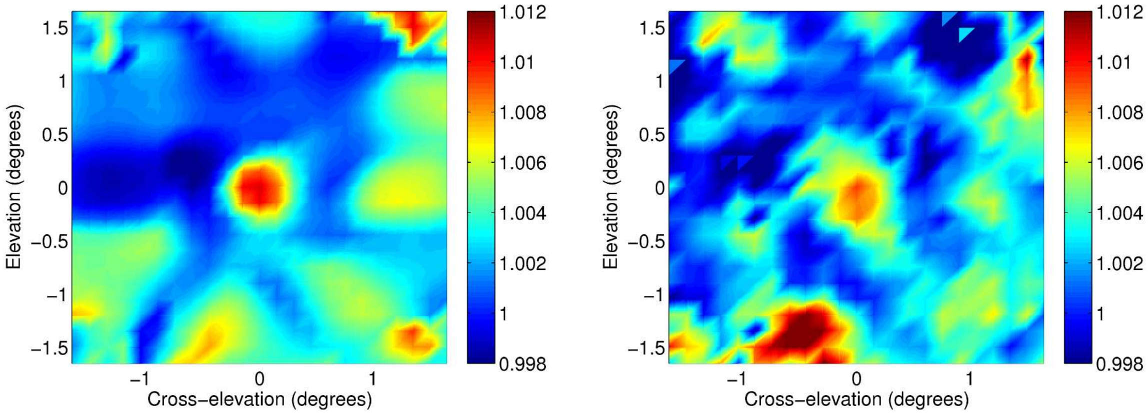

(a) Max-SNR beamformers.

(b) Sinc-matched beamformers.

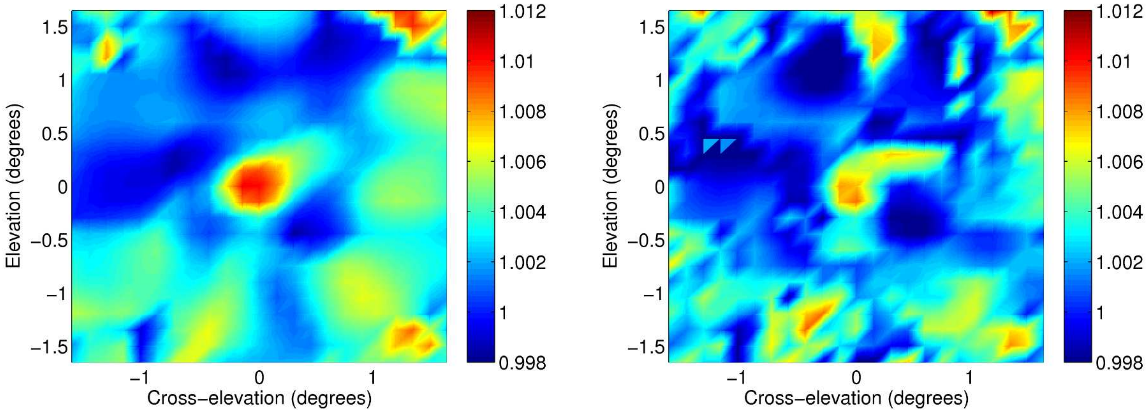

(c) Rim constraint beamformers.

(d) Dual constraint beamformers.

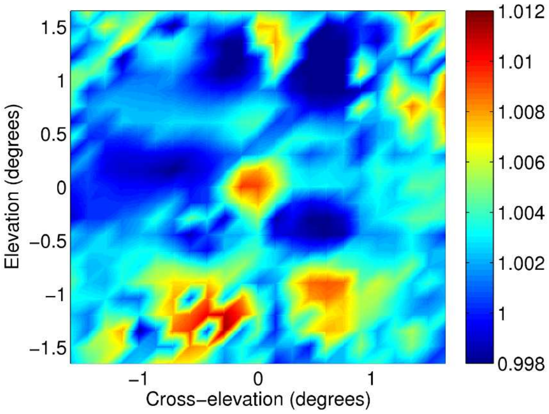

(e) Combination beamformers.

Figure 5.10: Images of the field surrounding the source 3c309, which has an intensity of about 7 Jy. The images are generated using the max-SNR, sinc-matched, rim constraint, dual constraint, and combination beamformers normalized to a single nearby noise field.

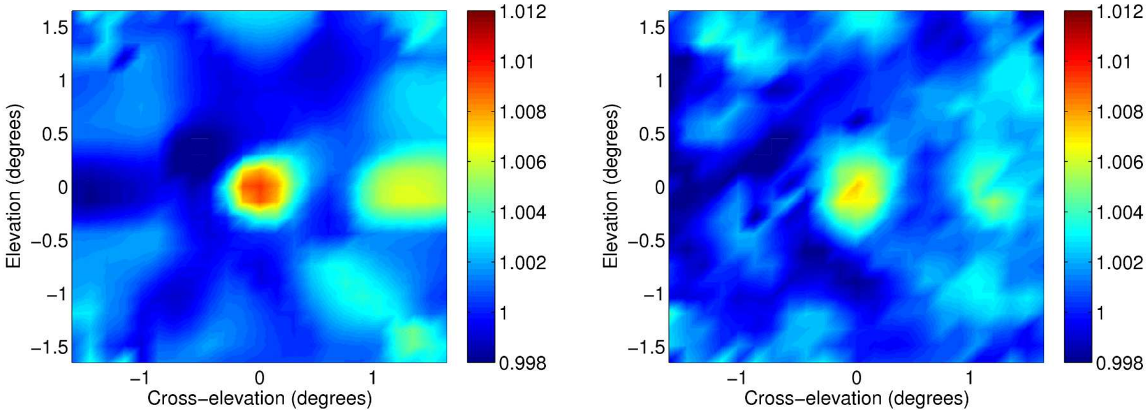

(a) Max-SNR beamformers.

(b) Sinc-matched beamformers.

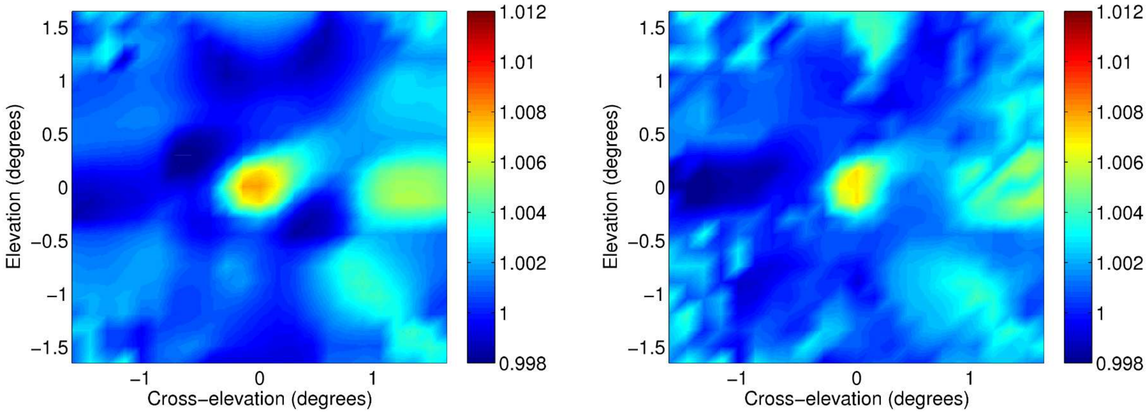

(c) Rim constraint beamformers.

(d) Dual constraint beamformers.

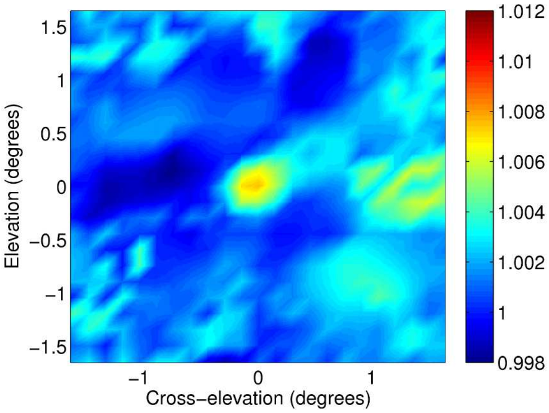

(e) Combination beamformers.

Figure 5.11: Images of the field surrounding the source 3c309, which has an intensity of about 7 Jy. The images are generated using the max-SNR, sinc-matched, rim constraint, dual constraint, and combination beamformers. In each case the beamformers were normalized to the average of two nearby noise fields.

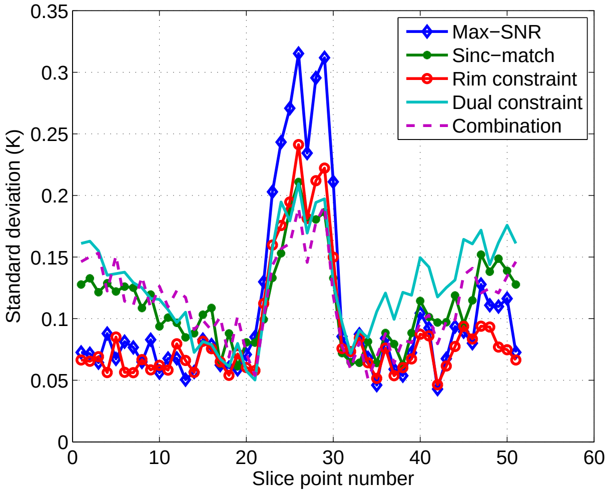

Figure 5.12: Standard deviation of beam powers for different beamformers as a function of distance from normalization noise field. The slice passes over a 10 Jy source. The beams were all normalized to the average of two nearby noise fields. Compare to Figure 5.8.

(by including additional samples), the tests were repeated using only half of the samples for each noise field measurement and the results did not change.

5.5 Simulation Results

Since the experimental data are limited to just three rim sources in the design of the rim constraint and combined beamformers the full benefits of such an approach are unclear. It is therefore worthwhile to explore this idea further using detailed electromagnetic models [74]. The model incorporates the array and reflector dish, as well as the full noise field in which the radio telescope operates. Whereas in a real system it is a challenge to obtain bare array measurements that allow us to view the array illumination pattern, this information is readily available from the model, allowing us to investigate different scenarios to identify the best approach to implementing a system of rim sources.

For the rim constraint beamformer model, 15 calibrator sources are placed at an angle of 10° above the dish rim as seen from the PAF. Five eigenvector constraints corresponding to the largest singular values of the constraint matrix were used to drive the illumination pattern response toward zero at the dish rim. By moving the sources beyond the rim of

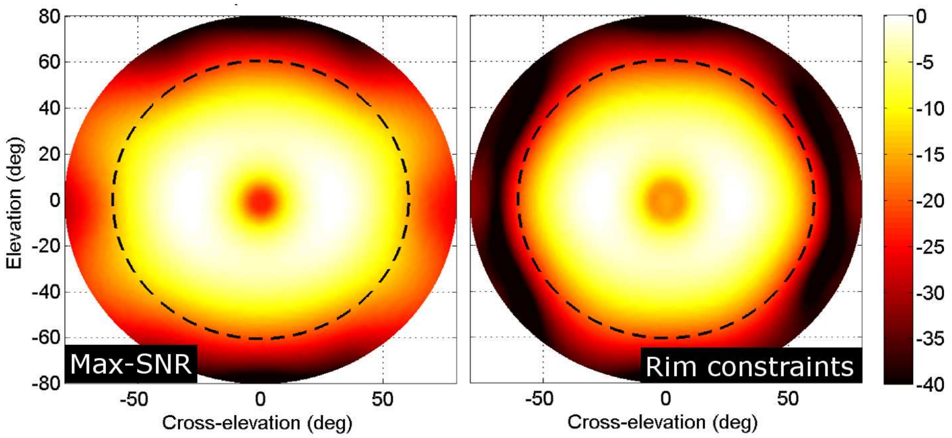

Figure 5.13: Simulated illumination patterns with and without illumination pattern constraints for a boresight steered far-field beam. The left-half image is the unconstrained max-SNR pattern. The right-half image is the rim constrained beamformer with 15 rim calibrator sources positioned at 10° above the rim with respect to the array. The added constraints provide greatly improved illumination of the dish with reduced spillover.

the dish, the array pattern more efficiently covers the dish surface while still limiting the response in the spillover region. The modeled results are for a dish steered to 55° elevation.

Figure 5.13 compares the max-SNR and rim constraint beamformers dish illumination patterns corresponding to a boresight-pointed far-field beam. The dish rim is indicated by the dashed line. The benefit of the eigenvector constraints is clear from this picture; the max-SNR pattern spills over the edge of the dish while the rim constraint beamformer provides almost uniformly low response levels at the rim. The effect on the far-field pattern is depicted in Figure 5.14, where very little distortion is observed in the rim constraint pattern compared to the max-SNR pattern.

The results of steering the far-field beam to an angle of 1° off-boresight are shown in Figs 5.15 and 5.16. As before, some relatively high gain portion of the max-SNR illumination pattern lies beyond the rim of the reflector dish. The rim constraint beamformer is again able to provide control to limit the response in the spillover region. The far-field pattern exhibits some distortion but it is relatively minor considering the constraints introduced in the beamformer and their effect on the illumination pattern.

Figure 5.17 is the modeled version of Figure 5.8 (without a 10 Jy source) and shows the standard deviation of modeled powers obtained for different beams within a 1° box

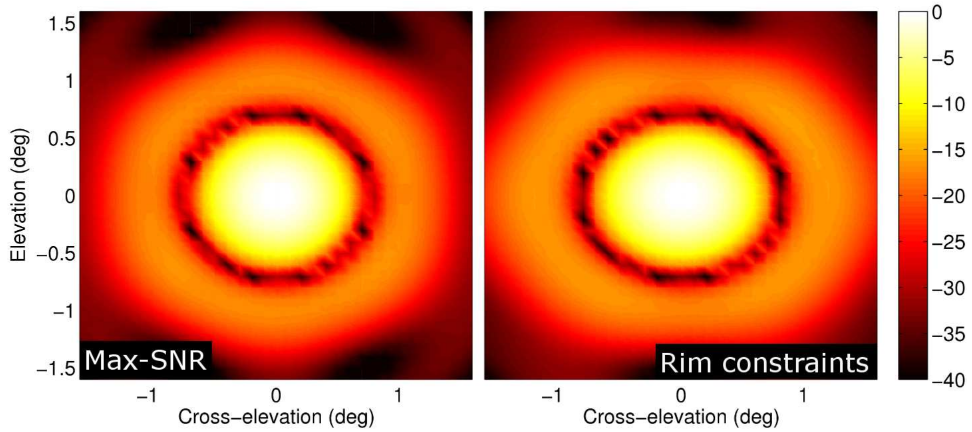

Figure 5.14: Simulated far-field patterns with and without illumination pattern constraints for a boresight steered far-field beam. The left-half image is the unconstrained max-SNR pattern. The right-half image is the rim constrained beamformer with 15 rim calibrator sources positioned at 10° above the rim with respect to the array. The added constraints cause only slight distortions in the far-field pattern.

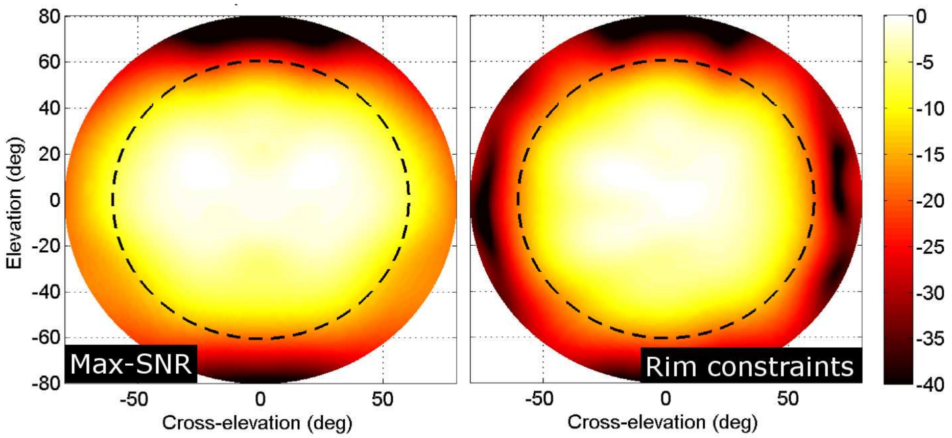

Figure 5.15: Simulated illumination patterns with and without illumination pattern constraints for a 1° off-boresight steered far-field beam. The left-half image is the unconstrained max-SNR pattern. The right-half image is the rim constrained beamformer with 15 rim calibrator sources positioned at 10° above the rim with respect to the array. The added constraints provide greatly improved illumination of the dish with reduced spillover.

around the center of the calibration grid. The beams are each normalized to the first of the 51 slice points. There is a 4° elevation change between the first and last points to match the experimental data slice. The most significant result shown in this plot is the reduction in variation for the combination beamformer. This beamformer provides known uniform main beam signal response, constrained noise response, and still achieves beam-to-beam variations

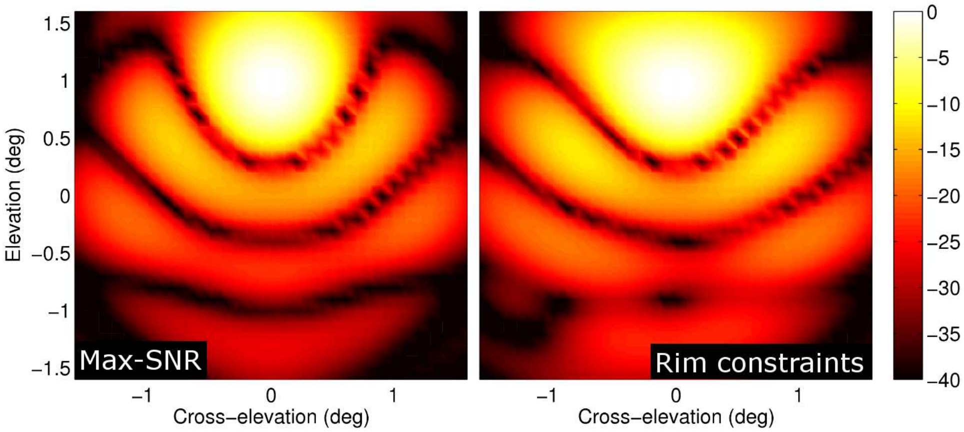

Figure 5.16: Simulated far-field patterns with and without illumination pattern constraints for a 1° off-boresight steered far-field beam. The left-half image is the unconstrained max-SNR pattern. The right-half image is the rim constrained max-SNR with 15 rim calibrator sources positioned at 10° above the rim with respect to the array. The added constraints cause only slight distortions in the far-field pattern.

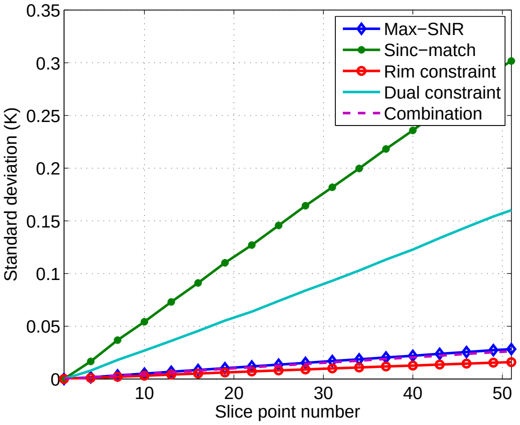

Figure 5.17: Standard deviation of beam powers for different modeled beamformers as a function of distance from normalization noise field. The slice covers a 4° elevation change to match the experimental data. The combination beamformer reduces beam-to-beam variations just as well as the max-SNR and rim constraint beamformers. The line showing the performance of the combination beamformer is barely visible, just below that of the max-SNR beamformer.

on the same order as those of the max-SNR and rim constraint beamformers. This is a very promising result that justifies experimental work with a more extensive system of rim sources for constraining the illumination pattern.

5.6 Conclusion

The primary advantage of a PAF system for radio astronomy is the ability to see a much larger portion of the sky with each distinct pointing of the telescope. As a beam is electronically steered to different points within the PAF field of view an image of the sky can be created much more rapidly compared to a conventional single feed system that only measures one image pixel at a time. For this to be beneficial the beam-to-beam variations must be minimized so that the weak sources of interest are not obscured in the image. This chapter has addressed the issue by presenting several options for reducing beam-to-beam variations in order to obtain accurate radio astronomical images.

Normalization to a contemporary off-pointed noise field estimate outside of but local to the FOV is used to reduce the noise variation over the FOV. However, while this normalization removes differences between beam noise responses it introduces differences in the main beam signal responses of each beamformer. This distorts the relative image pixel values so that absolute power levels cannot be measured. This can produce discontinuities and artifacts in a radio camera image of an extended source. A dual constraint beamformer has been introduced that is capable of providing both uniform specified signal response and noise response. The penalty is a reduction in sensitivity since it raises the noise response to a level that can be met by all the beams. The dual constraint beamformer does not improve beam-to-beam variations, which are due to differences between spillover noise structure in the off-pointed noise reference and that found in the FOV observation.

A first step in reducing beam variations is to average multiple noise field estimates before normalizing. This has been shown to improve variations by 25-50%, depending on the beamformer. A significant contributor is the changing spillover response between beams since different beams see differing amounts of spillover noise. This has been shown to be reduced with rim calibrator sources and eigenvector constraints to control the illumination pattern at the dish rim. The rim constraint beamformer combined with the dual constraint beamformer offers a beamformer with desired signal response, uniform noise response, and reduced variations between beams that is consistent with the variations seen for the max-SNR and rim constraint beamformers. Real data experimental results are limited, but encouraging. In simulation the combination beamformer appears to offer a very promising approach to PAF imaging.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

As radio astronomical PAF research and development continues to progress toward a science-ready instrument there are practical issues to be addressed with regards to operational procedures. This dissertation addresses some of these concerns by providing analysis and solutions on the topics of calibration, beamforming, and imaging for single dish observations.

The first step in working with a PAF-fed radio telescope is to obtain accurate calibration data, so that proper beamformers may be computed. The most accurate calibrators come from measurements obtained close in time and proximity to the astronomical field of interest. However, this can be an impractical expectation since obtaining a set of calibration vectors is a very time consuming process that may require several hours to record data from hundreds or thousands of points. Additionally, available calibration sources are scarce and sparsely located in the sky.

It has been proposed that temporal variations in calibrators due to electronic gain drift and other time-varying sources may be corrected by measuring the change in response to calibration sources at the vertex and on the rim of a radio telescope. This dissertation has shown this to be unnecessary at this time because directionally dependent sources of variation are much stronger than any temporal variations. This dissertation has shown this to be unnecessary at this time because directionally dependent sources of variation are much stronger than any temporal variations. A study of long-term calibration drift has shown remarkable stability over periods of days and weeks. This stability was measured by variations in sensitivity and beam pattern structure. In both cases, the variation was found to be affected more by other factors than by the age of the calibrators from which the beamformer was derived. Until the source of the larger variations can be identified and accounted for, the penalty for using an old beamformer is negligible.

After acquiring a set of calibration array response vectors, beamformers can be designed to meet the desired observational goals. Due to the low SNR environment in which a radio telescope operates, something like the max-SNR beamformer seems most appropriate. However, since this beamformer relies on the statistics of the current observed noise field it yields an unknown and unpredictable beam pattern. This may create a problem when observing near strong sources or creating images where beam-to-beam stability is needed. At the other end of the beamformer design spectrum is the class of fully deterministic beamformers that have known and predictable spatial response, but exhibit lower sensitivity.

A hybrid beamforming method has been proposed which combines both the max-SNR and deterministic beamforming methods. A weighting parameter is used to control the contribution of each so that the user may design a beamformer with the maximum sensitivity for a given amount of beam pattern control. This beamformer has been shown to provide increased control with only minimal degradation to the sensitivity performance.

Beamformers used for imaging a desired field of view must have a limited amount of variation between beam responses in order to distinguish a weak source from among these variations. A traditional radio telescope with a single horn feed does not have this problem since the beam pattern remains constant as the dish is steered around to each pixel in the desired image. The advantage of a PAF is that it can be used to obtain multiple image pixels per telescope pointing (which makes it particularly useful for rapidly evolving sources compared to a single beam system), but if the variation introduced by the different beams is too large then the advantage is forfeited. I note that with a single feed antenna there are still some variations between pixels due to the variability of the noise fields as the dish is steered. For a PAF to be a reasonable option for imaging, the beam-to-beam variations must be reduced to something near those observed with a single feed system.

Several approaches were introduced to minimize the beam-to-beam variations in a PAF image. Normalizing beamformer weights to an average measurement of multiple, distinct noise fields provides a considerable amount of reduction without much effort. A dual constraint beamformer that simultaneously meets a specified main beam gain and noise response allows relative power measurements between pixels to be measured with accuracy. Applying eigenvector constraints based on data measured from calibration sources placed at the dish rim provides control of the array's illumination pattern on the dish and reduces the spillover noise variation between beams. By combining this rim constraint beamformer with the dual constraint beamformer another option is made available that has reduced beam-to-beam variation while meeting main beam gain and noise response specifications.

6.2 Future Work

The field of research that aims to incorporate PAFs for radio astronomy use is still relatively new. The work presented in this dissertation represents preliminary results to important problems, however these are not necessarily fully resolved. Ideas for future work related to the topics of calibration, beamforming, and imaging for PAF single dish radio astronomy are presented here.

The long-term calibration stability analysis showed that a significant portion of the variation in sensitivity and beam pattern measurements is due to directionally dependent spillover and main beam noise. It is unclear, however, how each noise source contributes to the total variability, and how this affects the calibration measurements. Additional work is needed to identify and analyze the possible sources of this variation and derive methods for correcting the instability. Once the large variations are reduced, it may be necessary to develop techniques for updating old calibrators with rim calibration source data or by some other means.

Improved beamformer design is needed to provide stable patterns and high sensitivity. The deterministic beamformers used in this work are representative of a large class of such beamformers and were largely selected for their simplicity. More sophisticated beamformers can be designed to more fully address issues of high noise and pattern stability. Also, more sophisticated electromagnetic and noise models may allow for beamformers to be designed in simulation and then applied to real-life measurements. Further research is needed to identify whether this is even possible or if there are fundamental issues that prevent such a high level of accuracy.

Beam-to-beam variability causes problems for PAF imaging, since it may obscure a weak source in the field of view. Some beamformers were introduced to help address this issue, but more can still be done. An apparent next step in this research is to construct an experimental rim constraint system that is similar to the one used to obtain the modeled results. There were some good experimental results obtained using just three rim sources but it appears that this can be greatly improved by deploying additional sources. Also, improvement may be made to the dual constraint beamformer by adding additional far-field constraints or by using a different minimization function.

Full demonstration of sensitive, flat noise field, radio camera imaging has yet to be demonstrated with real data. New experimental data collection runs are needed to try some of the methods proposed here, which have only been run in simulation. Weak source detection with a PAF needs to be demonstrated. Additionally, beam-to-beam variability needs to be added as a term in the derivation of noise fluctuation for PAFs.

A fixed-adaptive beamformer was mentioned as an alternative to fully-adaptive beamforming. With pre-computed weights based on calibration data, such an approach could offer considerable computational savings while still providing some of the benefits of fully-adaptive beamformers. Further research in this area is required to thoroughly understand the quantitative benefits of doing so.

Bibliography

[1] G. Westerhout, "The early history of radio astronomy," *Annals of the New York Academy of Sciences*, vol. 198, no. 1, pp. 211–218, 1972. 1

[2] W. T. Sullivan, *Classics in Radio Astronomy*, ser. Studies in the History of Modern Science. Reidel Publishing Company, 1982. 1

[3] ——, *The Early Years of Radio Astronomy: Reflections Fifty Years After Jansky's Discovery*. Cambridge University Press, 2004. 1

[4] R. D. Ekers, "Observational advances in radio astronomy," in *Seeing Through the Dust: The Detection of HI and the Exploration of the ISM in Galaxies*, ser. Astronomical Society of the Pacific Conference Series, A. R. Taylor, T. L. Landecker, and A. G. Willis, Eds., vol. 276, Dec. 2002, p. 37. 1

[5] J. S. Hey, *The Evolution of Radio Astronomy*. Science History Publications, 1973. 1

[6] J. D. Kraus, M. Tiuri, and A. V. Räisänen, *Radio astronomy*, 2nd ed. Cygnus-Quasar Books, 1986. 1

[7] K. I. Kellermann and J. M. Moran, "The development of high-resolution imaging in radio astronomy," *Annual Review of Astronomy and Astrophysics*, vol. 39, pp. 457–509, 2001. 2

[8] A. R. Thompson, J. M. Moran, and J. George W. Swenson, *Interferometry and Synthesis in Radio Astronomy*, 2nd ed. John Wiley & Sons, 2008. 2

[9] T. L. Wilson, K. Rohlfs, and S. Huttemeister, *Tools of Radio Astronomy*, 5th ed. Springer, 2009. 2, 13

[10] P. J. Napier, A. R. Thompson, and R. D. Ekers, "The very large array: Design and performance of a modern synthesis radio telescope," *Proceedings of the IEEE*, vol. 71, no. 11, pp. 1295 – 1320, Nov. 1983. 2

[11] P. J. Napier, D. S. Bagri, B. G. Clark, A. E. E. Rogers, J. D. Romney, A. R. Thompson, and R. C. Walker, "The very long baseline array," *Proceedings of the IEEE*, vol. 82, no. 5, pp. 658 –672, May 1994. 2

[12] J. W. M. Baars and B. G. Hooghoudt, "The synthesis radio telescope at Westerbork. general lay-out and mechanical aspects," *Astronomy and Astrophysics*, vol. 31, pp. 323–331, Mar. 1974. 2

[13] P. Caselli, P. C. Myers, and P. Thaddeus, "Radio-astronomical spectroscopy of the hyperfine structure of N2H+," *The Astrophysical Journal Letters*, vol. 455, no. 1, p. L77, 1995. 2

[14] B. P. Wakker, "Recent developments concerning high-velocity clouds," *Astrophysics and Space Science*, vol. 289, pp. 381–390, Feb. 2004. 2

[15] J. W. T. Hessels, S. M. Ransom, I. H. Stairs, P. C. C. Freire, V. M. Kaspi, and F. Camilo, "A radio pulsar spinning at 716 hz," *Science*, vol. 311, no. 5769, pp. 1901–1904, 2006. 2

[16] T. Chang, U. Pen, K. Bandura, and J. B. Peterson, "An intensity map of hydrogen 21-cm emission at redshift = 0.8," *Nature*, vol. 466, pp. 463–465, July 2010. 2

[17] W. D. Watson, "Interstellar molecule reactions," *Reviews of Modern Physics*, vol. 48, pp. 513–552, Oct. 1976. 2

[18] J. M. Cordes, P. C. C. Freire, D. R. Lorimer, F. Camilo, D. J. Champion, D. J. Nice, R. Ramachandran, J. W. T. Hessels, W. Vlemmings, J. van Leeuwen, S. M. Ransom, N. D. R. Bhat, Z. Arzoumanian, M. A. McLaughlin, V. M. Kaspi, L. Kasian, J. S. Deneva, B. Reid, S. Chatterjee, J. L. Han, D. C. Backer, I. H. Stairs, A. A. Deshpande, and C.-A. Faucher-Gigure, "Arecibo pulsar survey using ALFA I. survey strategy and first discoveries," *Astrophysical Journal*, vol. 637, no. 1, pp. 446–455, Jan. 2006. 2

[19] L. Staveley-Smith, W. E. Wilson, T. S. Bird, M. W. Sinclair, and R. D. Ekers, "The Parkes 21 cm multibeam receiver," *ASP Conference Series, Multi-feed systems for radio telescopes*, vol. 75, pp. 136–144, 1995. 2

[20] L. Staveley-Smith, W. E. Wilson, T. S. Bird, M. J. Disney, R. D. Ekers, K. C. Freeman, R. F. Haynes, M. W. Sinclair, R. A. Vaile, R. L. Webster, and A. E. Wright, "The Parkes 21 cm multibeam receiver," *Publications Astronomical Society of Australia*, vol. 13, no. 3, pp. 243–248, Nov. 1996. 2

[21] B. Veidt and P. Dewdney, "A phased-array feed demonstrator for radio telescopes," in *Proc. URSI General Assembly*, 2005. 2

[22] W. A. van Cappellen, J. G. Bij de Vaate, M. V. Ivashina, L. Bakker, and T. Oosterloo, "Focal plane arrays evolve," in *URSI General Assembly*, Chicago, Illinois, USA, Aug. 2008. 2, 5, 29

[23] R. P. Millenaar, T. Oosterloo, and M. Brentjens, "The radio observatory at ASTRON: News from the WSRT and LOFAR," in *URSI General Assembly*, Chicago, Illinois, USA, Aug. 2008. 2

[24] J. R. Fisher, K. F. Warnick, B. D. Jeffs, G. Cortes-Medellin, R. D. Norrod, F. J. Lockman, J. M. Cordes, and R. Giovanelli, "Phased array feeds," *Astro2010 Technology Development White Paper*, 2009, national Radio Astronomy Observatory, Charlottesville, VA, USA. 2

[25] S. G. Hay, J. D. O'Sullivan, J. S. Kot, C. Granet, A. Grancea, A. R. Forsyth, and D. H. Hayman, "Focal plane array development for ASKAP (Australian SKA Pathfinder)," in *Antennas and Propagation, 2007. EuCAP 2007. The Second European Conference on*, Nov. 2007, pp. 1–5. 2

[26] S. G. Hay and J. D. O'Sullivan, "Analysis of common-mode effects in a dual-polarized planar connected-array antenna," *Radio Science*, vol. 43, p. RS6S04, 2008. 2

[27] R. Nan, D. Li, C. Jin, Q. Wang, L. Zhu, W. Zhu, H. Zhang, Y. Yue, and L. Qian, "The five-hundred-meter aperture spherical radio telescope (FAST) project," *International Journal of Modern Physics D*, vol. 20, no. 6, pp. 989–1024, 2011. 2

[28] J. Landon, M. Elmer, J. Waldron, D. Jones, A. Stemmons, B. D. Jeffs, K. F. Warnick, J. R. Fisher, and R. D. Norrod, "Phased array feed calibration, beamforming, and imaging," *The Astronomical Journal*, vol. 139, no. 3, pp. 1154–1167, 2010. [Online]. Available: http://stacks.iop.org/1538-3881/139/1154 2, 7, 8, 15, 16, 20, 21, 56, 58, 63, 84

[29] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. W. Lazio, "The Square Kilometre Array," *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1482–1496, Aug. 2009. 2

[30] D. R. DeBoer, R. G. Gough, J. D. Bunton, T. J. Cornwell, R. J. Beresford, S. Johnston, I. J. Feain, A. E. Schinckel, C. A. Jackson, M. J. Kesteven, A. Chippendale, G. A. Hampson, J. D. O'Sullivan, S. G. Hay, C. E. Jacka, T. W. Sweetnam, M. C. Storey, L. Ball, and B. J. Boyle, "Australian SKA Pathfinder: A high-dynamic range wide-field of view survey telescope," *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1507–1521, Aug. 2009. 2

[31] M. Ivashina, J. Simons, and J. Bij de Vaate, "Efficiency analysis of focal plane arrays in deep dishes," *Experimental Astronomy*, vol. 17, no. 1-3, pp. 149–162, June 2004. 2

[32] R. Braun, "The concept of the square kilometer array interferometer," in *High-sensitivity Radio Astronomy*, N. Jackson and R. Davis, Eds. Cambridge University Press, 1997, pp. 260–268. 2

[33] K. F. Warnick and M. A. Jensen, "Effects of mutual coupling on interference mitigation with a focal plane array," *Antennas and Propagation, IEEE Transactions on*, vol. 53, no. 8, pp. 2490–2498, Aug. 2005. 3, 16

[34] M. V. Ivashina, M. Kehn, P.-S. Kildal, and R. Maaskant, "Control of reflection and mutual coupling losses in maximizing efficiency of dense focal plane arrays," in *Antennas and Propagation, 2006. EuCAP 2006. First European Conference on*, Nov. 2006, pp. 1–6. 3

[35] M. Ng Mou Kehn, M. V. Ivashina, P.-S. Kildal, and R. Maaskant, "Definition of unifying decoupling efficiency of different array antennas–case study of dense focal plane array feed for parabolic reflector," *AEU - International Journal of Electronics and Communications*, vol. In Press, Corrected Proof, 2009. 3

[36] J. E. Evans, D. F. Sun, and J. R. Johnson, "Application of advanced signal processing techniques to angle of arrival estimation in ATC navigation and surveillance systems," *Technical Report, MIT Lincoln Laboratory*, June 1982. 5

[37] B. Ottersten, "Array processing for wireless communications," in *Statistical Signal and Array Processing, 1996. Proceedings., 8th IEEE Signal Processing Workshop on (Cat. No.96TB10004*, June 1996, pp. 466–473. 5

[38] B. D. Van Veen and K. M. Buckley, "Beamforming: a versatile approach to spatial filtering," *ASSP Magazine, IEEE [see also IEEE Signal Processing Magazine]*, vol. 5, no. 2, pp. 4–24, Apr. 1988. 5, 13, 22, 23, 24, 25, 57, 89, 93

[39] H. L. V. Trees, *Detection, Estimation, and Modulation Theory, Part IV, Optimum Array Processing.* John Wiley and Sons, 2002. 5, 13, 23, 30

[40] M. Ng Mou Kehn and L. Shafai, "Characterization of dense focal plane array feeds for parabolic reflectors in achieving closely overlapping or widely separated multiple beams," *Radio Science*, vol. 44, no. 3, p. RS3014, 2009. 5

[41] K. J. Maalouf and E. Lier, "Theoretical and experimental study of interference in multibeam active phased array transmit antenna for satellite communications," *Antennas and Propagation, IEEE Transactions on*, vol. 52, no. 2, pp. 587–592, Feb. 2004. 5

[42] D. B. Hayman, A. Chippendale, R. Qiao, J. D. Bunton, R. J. Beresford, P. Roberts, and P. Axtens, "Gain calibration of phased array feeds," in *Electromagnetics in Advanced Applications (ICEAA), 2010 International Conference on*, Sept. 2010, pp. 418–421. 5, 29

[43] O. M. Smirnov and M. V. Ivashina, "Element gain drifts as an imaging dynamic range limitation in PAF-based interferometers," in *URSI General Assembly*, Istanbul, Turkey, Aug. 2011. 5, 8, 52

[44] M. Sarcione, J. Mulcahey, D. Schmidt, K. Chang, M. Russell, R. Enzmann, P. Rawlinson, W. Guzak, R. Howard, and M. Mitchell, "The design, development and testing of the THAAD (Theater High Altitude Area Defense) solid state phased array (formerly ground based radar)," in *Phased Array Systems and Technology, 1996., IEEE International Symposium on*, Oct. 1996, pp. 260 –265. 6

[45] R. Sorace, "Phased array calibration," *Antennas and Propagation, IEEE Transactions on*, vol. 49, no. 4, pp. 517–525, apr 2001. 6

[46] D. B. Hayman, T. S. Bird, K. P. Esselle, and P. J. Hall, "Experimental demonstration of focal plane array beamforming in a prototype radiotelescope," *Antennas and Propagation, IEEE Transactions on*, vol. 58, no. 6, pp. 1922–1934, June 2010. 6, 29

[47] W. A. van Cappellen and M. V. Ivashina, "Temporal beam pattern stability of a radio astronomy phased array feed," in *Antennas and Propagation (EUCAP), Proceedings of the 5th European Conference on*, Apr. 2011, pp. 926–929. 6, 29, 52

[48] B. Veidt, G. Hovey, T. Burgess, R. Smegal, R. Messing, A. Willis, A. Gray, and P. Dewdney, “Demonstration of a dual-polarized phased-array feed,” *Antennas and Propagation, IEEE Transactions on*, vol. 59, no. 6, pp. 2047 –2057, June 2011. 6, 7

[49] Y. T. Lo, S. W. Lee, and Q. H. Lee, “Optimization of directivity and signal-to-noise ratio of an arbitrary antenna array,” *Proceedings of the IEEE*, vol. 54, no. 8, pp. 1033–1045, Aug. 1966. 7, 56

[50] T. Oosterloo, W. van Cappellen, and L. Bakker, “First results with APERTIF,” in *Calibration and Imaging Workshop, CALIM2008*, Perth, Australia, Apr. 2008, http://calim2008.atnf.csiro.au/twiki/pub/Main/WorkshopProgram/OosterlooCalim.pdf. 7, 56

[51] B. D. Jeffs, K. F. Warnick, M. Elmer, J. Landon, J. Waldron, D. Jones, R. Fisher, and R. Norrod, “Calibration and optimal beamforming for a 19 element phased array feed,” in *Calibration and Imaging Workshop, CALIM2008*, Perth, Australia, Apr. 2008, http://calim2008.atnf.csiro.au/twiki/pub/Main/WorkshopProgram/JeffsCalim.pdf. 7, 56

[52] M. A. W. Verheijen, T. A. Oosterloo, W. A. van Cappellen, L. Bakker, M. V. Ivashina, and J. M. van der Hulst, “Apertif, a focal plane array for the WSRT,” in *The Evolution of Galaxies Through the Neutral Hydrogen Window*, R. Minchin and E. Momjian, Eds., vol. 1035. AIP, 2008, pp. 265–271. 7, 56

[53] K. F. Warnick, B. D. Jeffs, J. Landon, J. Waldron, D. Jones, J. R. Fisher, and R. Norrod, “Beamforming and imaging with the BYU/NRAO L-band 19-element phased array feed,” in *Antenna Technology and Applied Electromagnetics and the Canadian Radio Science Meeting, 2009. ANTEM/URSI 2009. 13th International Symposium on*, Feb. 2009, pp. 1–4. 7, 8, 56, 84

[54] W. A. van Cappellen, L. Bakker, and T. A. Oosterloo, “Experimental results of a 112 element phased array feed for the Westerbork Synthesis Radio Telescope,” in *Antennas and Propagation Society International Symposium, 2009. APSURSI ’09. IEEE*, June 2009, pp. 1–4. 7, 8, 56, 84

[55] M. V. Ivashina, O. A. Iupikov, R. Maaskant, W. A. van Cappellen, L. Bakker, and T. Oosterloo, “Off-axis beam performance of focal plane arrays for the Westerbork Synthesis Radio Telescope - initial results of a prototype system,” in *Antennas and Propagation Society International Symposium, 2009. APSURSI ’09. IEEE*, June 2009, pp. 1–4. 7, 56

[56] B. D. Jeffs, K. F. Warnick, J. Landon, J. Waldron, D. Jones, J. R. Fisher, and R. D. Norrod, “Signal processing for phased array feeds in radio astronomical telescopes,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 5, pp. 635–646, Oct. 2008. 7, 26, 27, 56, 57

[57] C. K. Hansen, "Beamforming techniques and interference mitigation using a multiple feed array for radio astronomy," M.S. Thesis, Brigham Young University, Provo, Utah, Apr. 2004. 7, 57

[58] K. F. Warnick, B. Woestenburg, L. Belostotski, and P. Russer, "Minimizing the noise penalty due to mutual coupling for a receiving array," *Antennas and Propagation, IEEE Transactions on*, vol. 57, no. 6, pp. 1634–1644, June 2009. 7, 16, 38, 57, 77

[59] J. S. Waldron, "Nineteen-element experimental phased array feed development and analysis on effects of focal plane offset and beam steering on sensitivity," M.S. Thesis, Brigham Young University, Provo, Utah, Aug. 2008. 7, 20, 34, 57

[60] M. V. Ivashina, O. Iupikov, R. Maaskant, W. A. van Cappellen, and T. Oosterloo, "An optimal beamforming strategy for wide-field surveys with phased-array-fed reflector antennas," *Antennas and Propagation, IEEE Transactions on*, vol. 59, no. 6, pp. 1864–1875, June 2011. 7

[61] M. Voronkov and T. Cornwell, "On the calibration and imaging with eigenbeams," *ATNF SKA Memo 12*, Jan. 2007, http://www.atnf.csiro.au/projects/mira/newdocs/eigenbeams.pdf. 7, 56

[62] T. Willis, "Simulations of synthesis telescope antennas equipped with focal plane arrays," in *Calibration and Imaging Workshop, CALIM2009*, Socorro, New Mexico, Mar. 2009, https://safe.nrao.edu/wiki/pub/Software/CalIm09Program/agw_calim09.pdf. 8, 56

[63] A. Schinckel, J. Bunton, A. Chippendale, R. Gough, G. Hampson, S. Hay, C. Jackson, K. Jeganathan, J. O'Sullivan, J. Reynolds, R. Shaw, and C. Wilson, "Recent progress in the australian ska pathfinder (askap)," in *Microwave Conference Proceedings (APMC), 2011 Asia-Pacific*, Dec. 2011, pp. 1178 –1181. 8

[64] S. G. Hay, "Comparison of single-ended and differential beamforming on the efficiency of a checkerboard phased array feed in offset- and front-fed reflectors," in *Antennas and Propagation (EuCAP), 2010 Proceedings of the Fourth European Conference on*, Apr. 2010, pp. 1 –5. 8

[65] J. Welch, D. Backer, L. Blitz, D. C. Bock, G. C. Bower, C. Cheng, S. Croft, M. Dexter, G. Engargiola, E. Fields, J. Forster, C. Gutierrez-Kraybill, C. Heiles, T. Helfer, S. Jorgensen, G. Keating, J. Lugten, D. MacMahon, O. Milgrome, D. Thornton, L. Urry, J. van Leeuwen, D. Werthimer, P. Williams, M. Wright, J. Tarter, R. Ackermann, S. Atkinson, P. Backus, W. Barott, T. Bradford, M. Davis, D. DeBoer, J. Dreher, G. Harp, J. Jordan, T. Kilsdonk, T. Pierson, K. Randall, J. Ross, S. Shostak, M. Fleming, C. Cork, A. Vitouchkine, N. Wadefalk, and S. Weinreb, "The allen telescope array: The first widefield, panchromatic, snapshot radio camera for radio astronomy and seti," *Proceedings of the IEEE*, vol. 97, no. 8, pp. 1438 –1447, Aug 2009. 8, 84

[66] R. Giovanelli, M. P. Haynes, B. R. Kent, P. Perillat, A. Saintonge, N. Brosch, B. Catinella, G. L. Hoffman, S. Stierwalt, K. Spekkens, M. S. Lerner, K. L.

Masters, E. Momjian, J. L. Rosenberg, C. M. Springob, A. Boselli, V. Charmandaris, J. K. Darling, J. Davies, D. G. Lambas, G. Gavazzi, C. Giovanardi, E. Hardy, L. K. Hunt, A. Iovino, I. D. Karachentsev, V. E. Karachentseva, R. A. Koopmann, C. Marinoni, R. Minchin, E. Muller, M. Putman, C. Pantoja, J. J. Salzer, M. Scodeggio, E. Skillman, J. M. Solanes, C. Valotto, W. van Driel, and L. van Zee, "The arecibo legacy fast alfa survey. i. science goals, survey design, and strategy," *The Astronomical Journal*, vol. 130, no. 6, p. 2598, 2005. [Online]. Available: http://stacks.iop.org/1538-3881/130/i=6/a=2598 9, 84

[67] D. G. Barnes, L. Staveley-Smith, W. J. G. De Blok, T. Oosterloo, I. M. Stewart, A. E. Wright, G. D. Banks, R. Bhathal, P. J. Boyce, M. R. Calabretta, M. J. Disney, M. J. Drinkwater, R. D. Ekers, K. C. Freeman, B. K. Gibson, A. J. Green, R. F. Haynes, P. Te Lintel Hekkert, P. A. Henning, H. Jerjen, S. Juraszek, M. J. Kesteven, V. A. Kilborn, P. M. Knezek, B. Koribalski, R. C. Kraan-Korteweg, D. F. Malin, M. Marquarding, R. F. Minchin, J. R. Mould, R. M. Price, M. E. Putman, S. D. Ryder, E. M. Sadler, A. Schrder, F. Stootman, R. L. Webster, W. E. Wilson, and T. Ye, "The H i Parkes all sky survey: southern observations, calibration and robust imaging," *Monthly Notices of the Royal Astronomical Society*, vol. 322, no. 3, pp. 486–498, 2001. [Online]. Available: http://dx.doi.org/10.1046/j.1365-8711.2001.04102.x 9, 84

[68] F. Haddock, "Introduction to radio astronomy," *Proceedings of the IRE*, vol. 46, no. 1, pp. 3 –12, Jan 1958. 13

[69] W. Kummer, "Basic array theory," *Proceedings of the IEEE*, vol. 80, no. 1, pp. 127–140, Jan 1992. 13

[70] H. Cox, R. Zeskind, and M. Owen, "Robust adaptive beamforming," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 35, no. 10, pp. 1365 – 1376, Oct 1987. 13

[71] T. Parks and J. McClellan, "Chebyshev approximation for nonrecursive digital filters with linear phase," *Circuit Theory, IEEE Transactions on*, vol. 19, no. 2, pp. 189–194, Mar. 1972. 13

[72] S. Ellingson and G. Hampson, "A subspace-tracking approach to interference nulling for phased array-based radio telescopes," *Antennas and Propagation, IEEE Transactions on*, vol. 50, no. 1, pp. 25–30, Jan. 2002. 13, 72

[73] K. F. Warnick and B. D. Jeffs, "Efficiencies and system temperature for a beamforming array," *Antennas and Wireless Propagation Letters, IEEE*, vol. 7, pp. 565–568, 2008. 15, 26, 27, 60

[74] D. Jones, "Characterization of a phased array feed model," M.S. Thesis, Brigham Young University, Provo, Utah, Aug. 2008. 16, 26, 27, 74, 108

[75] K. Warnick and B. Jeffs, "Gain and aperture efficiency for a reflector antenna with an array feed," *Antennas and Wireless Propagation Letters, IEEE*, vol. 5, no. 1, pp. 499–502, Dec. 2006. 19, 26

[76] D. Roddy, *Satellite Communications*, 4th ed. McGraw-Hill, 2006. 20

[77] J. R. Nagel, "A prototype platform for array feed development," M.S. Thesis, Brigham Young University, Provo, Utah, Dec. 2006. 20, 74, 139

[78] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-time signal processing*, 2nd ed. Prentice Hall, 1999. 25

[79] J. Rissanen, "A universal prior for integers and estimation by minimum description length," *The Annals of Statistics*, vol. 11, no. 2, pp. 416–431, 1983. 30

[80] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 33, no. 2, pp. 387–392, Apr. 1985. 30

[81] J. C. Landon, "Development of an experimental phased array feed system and algorithms for radio astronomy," Ph.D. Dissertation, Brigham Young University, Provo, Utah, Aug. 2011. 34, 40, 74, 136

[82] D. E. Carter, "Active impedance matching and sensitivity optimized phased array feed design for radio astronomy," M.S. Thesis, Brigham Young University, Provo, Utah, Dec. 2011. 38, 75, 129, 136

[83] National Radio Astronomy Observatory, "20 meter Geodetic Telescope," https://science.nrao.edu/facilities/gbt/other-telescopes/20meter, [Online, accessed July 2012]. 55

[84] P. A. Fridman and W. A. Baan, "RFI mitigation methods in radio astronomy," *Astronomy & Astrophysics*, vol. 378, pp. 327–344, 2001. 56

[85] D. Hayman, R. Beresford, J. Bunton, C. Cantrall, T. Cornwell, A. Grancea, C. Granet, J. Joseph, M. Kesteven, J. O'Sullivan, J. Pathikulangara, T. Sweetnam, and M. Voronkov, "The NTD interferometer: A phased array feed test bed," in *Workshop on Applications of Radio Science*, Queensland, Australia, Feb. 2008. 56

[86] C. Hansen, K. Warnick, B. Jeffs, J. Fisher, and R. Bradley, "Interference mitigation using a focal plane array," *Radio Science*, vol. 40, June 2005. 72

[87] J. Raza, A.-J. Boonstra, and A. van der Veen, "Spatial filtering of RF interference in radio astronomy," *Signal Processing Letters, IEEE*, vol. 9, no. 2, pp. 64–67, Feb. 2002. 72

[88] J. Landon, B. Jeffs, and K. Warnick, "Model-based subspace projection beamforming for deep interference nulling," *Signal Processing, IEEE Transactions on*, vol. PP, no. 99, p. 1, 2011. 72

[89] Y. Zhao, W. Liu, and R. Langley, "Adaptive wideband beamforming with frequency invariance constraints," *Antennas and Propagation, IEEE Transactions on*, vol. 59, no. 4, pp. 1175 –1184, Apr 2011. 93

[90] National Astronomy and Ionosphere Center, "Arecibo Observatory," http://www.naic.edu/general/, [Online, accessed July 2012]. 129

[91] V. Asthana, "Development of l-band down converter boards and real-time digital back-end for phased array feeds," M.S. Thesis, Brigham Young University, Provo, Utah, June 2012. 133, 136

[92] G. Cortes-Medellin, G. Rajagopalan, P. Perillat, A. Vishwas, K. Warnick, B. Jeffs, M. Elmer, D. Carter, V. Asthana, and T. Webb, "Field of view characterization of arecibo radio telescope with a phased array feed," in *Antennas and Propagation (APSURSI), 2011 IEEE International Symposium on*, July 2011, pp. 847 –850. 136

[93] G. Cortes-Medellin and D. Campbell, "Arecibo focal phased array feasibility study and instrument concept design," in *General Assembly and Scientific Symposium, 2011 XXXth URSI*, Aug. 2011, pp. 1 –4. 136

[94] K. Warnick, D. Carter, T. Webb, J. Landon, M. Elmer, and B. Jeffs, "Design and characterization of an active impedance matched low-noise phased array feed," *Antennas and Propagation, IEEE Transactions on*, vol. 59, no. 6, pp. 1876 –1885, June 2011. 136

[95] K. Warnick, D. Carter, T. Webb, B. Jeffs, J. Landon, V. Asthana, M. Elmer, R. Norrod, D. Roshi, and J. Fisher, "Towards a high sensitivity cryogenic phased array feed antenna for the green bank telescope," in *General Assembly and Scientific Symposium, 2011 XXXth URSI*, Aug. 2011, pp. 1 –4. 136

Appendix A

Glossary

1. $\mathbf{a}$: Bold, lowercase indicates a column vector.
2. $\mathbf{A}$: Bold, uppercase indicates a matrix.
3. a: Non-bold, italicized, uppercase and/or lowercase indicates a scalar.
4. $\mathbf{A}_k$: Italicized subscript indicates an index variable.
5. $\mathbf{A}_\mathrm{k}$: Non-italicized subscript indicates a fixed quantity definition.
6. $E\{\mathbf{a}\}$: Expected value of the random variable argument $\mathbf{a}$.
7. $\mathbf{A}^H$: Complex conjugate (Hermitian) transpose of $\mathbf{A}$.
8. $\mathbf{A}^T$: Transpose of $\mathbf{A}$.
9. $\mathbf{A}^\dagger$: Moore-Penrose pseudoinverse of $\mathbf{A}$.
10. $a[n]$: Temporal index of the stochastic process a.
11. $\hat{\mathbf{A}}$: Estimated value of $\mathbf{A}$. Also used for measured data, which is generally an estimate of the true value.
12. $a \propto b$: a is proportional to b.
13. $\mathbf{a}^*$: Conjugate of $\mathbf{a}$.
14. $\bar{\mathbf{a}}$: Sample mean of $\mathbf{a}$.
15. $|a|$: Magnitude of complex a.
16. $\angle a$: Phase of complex a.
17. $||\mathbf{a}||_2$: The l_2 norm of vector $\mathbf{a}$.
18. $||\mathbf{a}||_\infty$: The l_∞ norm of vector $\mathbf{a}$.
19. $||\mathbf{A}||_\mathbf{F}$: The Frobenius norm of matrix $\mathbf{A}$.
20. $\mathrm{diag}\{\mathbf{a}\}$: Diagonal matrix with elements of $\mathbf{a}$ along the diagonal.
21. $[\mathbf{a}]_k$: The kth element of the vector $\mathbf{a}$.

22. $\mathbf{a} \,./\, \mathbf{b}$: Element-wise division (applies to vectors and matrices).
23. $\mathbf{a} \odot \mathbf{b}$: Element-wise multiplication (applies to vectors and matrices).
24. $\mathbf{tr}(\mathbf{A})$: Trace of matrix $\mathbf{A}$.
25. $\grave{\mathbf{a}}$: Modeled value.
26. $\wp\{\cdot\}$: Operator that identifies the principal eigenvector (i.e., the eigenvector associated with the largest eigenvalue).

Appendix B

Experimental Platform Development and Observation Campaigns

In addition to participation in the Green Bank experiments described previously and numerous system tests performed on the roof of the BYU Clyde Building, two major experimental efforts were completed which contributed to the goals of the Radio Astronomy Systems Research (RASR) group. The first was a PAF feasibility study for the Arecibo telescope. The second was a redesign of the existing receivers in anticipation of a new 64-input FPGA-based data acquisition system capable of a 50 MHz sample rate and 20 MHz bandwidth per input channel. The details of these two activities are presented in this appendix.

B.1 Arecibo Telescope Feasibility Study

In 2009-2010 the RASR group participated in a PAF feasibility study on the largest single dish radio telescope in the world, the Arecibo Telescope. Development of a PAF-based system would allow the observatory to remain relevant in the radio astronomy community. Our role was to provide the PAF array, receiver electronics, and data acquisition system for installation on the telescope. We also performed data analysis to assess feasibility of developing a permanent PAF for Arecibo. The 2010 experiments served a dual purpose since they provided an opportunity to test a new array design, a new data acquisition system, and new receiver hardware.

The Arecibo Telescope, shown in Figure B.1(a), is the largest of its kind in the world and is located in Arecibo, Puerto Rico. It consists of a 305-meter stationary spherical reflector dish with a feed housed on a platform which is suspended above the dish surface by three concrete towers. The feed platform contains secondary and tertiary Gregorian reflectors which focus the received signal on the feed. These reflectors are situated within the “Gregorian dome” attached to the platform. Several feeds are mounted to a rotating floor (see Figure B.1(b)) which is remotely controlled to set the desired feed at the focus position. The platform is capable of movement that mimics elevation and azimuthal steering of a non-stationary reflector. See [90] for more information about the Arecibo Telescope and the Arecibo Observatory.

The 90-foot diameter Gregorian dome consists of three levels (stories of about eight-foot height) and housed all of the electronics and hardware needed for the experiment. The lower level contains the various feeds attached to the underside of the rotating floor. The BYU active impedance matched array was mounted here and used for the majority of the experiments. A dual pol version [82] was also used but only for a limited number of experiments. A small rack containing the receiver hardware was mounted to the top side of


(a) Arecibo Telescope. (b) Array mounted on rotating floor.

Figure B.1: (a) The Arecibo telescope consisting of the reflector dish and feed platform. (b) the BYU array mounted to the underside of the rotating feed platform.

the rotating floor in an air conditioned room on the middle level of the dome. Finally, the non-cooled upper level held the shielded rack containing the data acquisition computers and equipment.

All equipment within the racks, as well as the racks themselves, were firmly mounted to shelving and the floor, respectively. This was because of the motion of the dome in response to changes in elevation and azimuth commands. Data acquisition was remotely controlled using TCP/IP from a control room located on the ground.

The feasibility study involved many individuals from both BYU and the Arecibo Observatory. Extensive planning and coordination was required to successfully execute the desired experiment plan. An overview of the system upgrades that occurred in preparation for the experiment is provided below, along with the experiment plan and some results of the feasibility study.

B.1.1 Data Acquisition System

As a new 38-input dual polarized array was being constructed, a corresponding data acquisition (DAQ) system with 40 input channels was also developed. Instead of building a duplicate 20-channel system to synchronize and combine with the original 20-channel National Instruments (NI) system, the decision was made to acquire new computers and analog-to-digital converter (ADC) cards, and construct an entirely new system.

The 20-channel NI system consisted of a single server computer with five National Instruments PCI ADC cards installed. Each card was capable of sampling four inputs at a rate of 1.25 Msamples/second. These NI PCI-6115 cards were not designed for synchronous sampling across multiple computers. National Instruments manufactures a PXI-6115 card that is capable of providing this functionality but a joint decision with collaborators from the Green Bank Observatory (in anticipation of future Green Bank experiments) was made to use PCI-9812 ADC cards from Adlink Technology.

Delays with the development of the 40-channel Adlink DAQ system necessitated updates to the 20-channel NI system to allow communication with the Arecibo control computer. The NI system was used as the primary instrument throughout the first of the two Arecibo experiments which occurred in June of 2010. Development of the Adlink system continued during this time and it was available for use during the second experiment in August of 2010.

40-channel Adlink DAQ Systems

Two new systems, each capable of acquiring data across 40 input channels were constructed concurrently at BYU and the Green Bank Observatory for later use in experiments in Green Bank. They were built independently using the same Adlink ADC cards and similar computer architectures, with the potential to be used together to acquire up to 80 channels of data. Each used a central hub computer to control the data acquisition and store data files aggregated from several "node" computers. The Green Bank system contains 10 node computers, each with a single Adlink card installed. The BYU system shown in Figure B.2 contains five node computers with two Adlink cards installed in each. The PCI-9812 Adlink cards are capable of simultaneously sampling four inputs at a maximum sample rate of 20 MHz. In both cases 10 ADC cards are synchronized using common clock and trigger signals. The computers purchased were custom built by ASA computers. The Arecibo experiments provided motivation to complete the 40-channel BYU system in order to obtain data using the new dual-polarized array.

Three different teams of students participated in the 40-channel DAQ system development project over the course of a year and half before the project was turned over to the RASR group. The system hardware included five node computers, a large server computer to act as the data hub, 10 Adlink ADC cards, five 10-gigabit ethernet cards, and a 10-gigabit ethernet switch. *C* code was developed to communicate through the 10-gigabit switch and with a remote host computer to synchronously acquire data across all 10 cards at a sample rate of 1.25 MHz. Many hours were spent by myself and others leading up to and throughout the initial experimental campaign to finish the system. The 40-channel system was ready for use during the second of the two Arecibo experiments, but was limited to small acquisition times of 60 seconds or less.

Some of the issues identified during the experiment were lost samples due to conflicts with memory buffer settings, non-synchronous triggering and failed triggers, and problems associated with synchronizing the clock signals. Many of the problems were attributed to the use of the Adlink ADC cards and the ASA computers. In the case of both the cards and computers there was limited technical support provided by the manufacturers. In retrospect it may have been better to have purchased the trusted NI PXI-6115 cards and synchronized them between two brand name computers.

Adapting the 20-Channel NI System

Since at the time of the first Arecibo deployment the Adlink system was not yet reliable, the 20-channel NI system was updated and prepared to serve as our main DAQ system. This involved modifying the NI breakout boxes to couple a sinusoidal tone burst to the first input of each PCI-6115 card. The tone burst signal was a 312.5 KHz sinusoid

Figure B.2: Front and back views of the 40-channel data acquisition system, consisting of a central control computer and five node computers. Two ADC cards were installed on each node and synchronous acquisition was controlled by the hub computer.

that was turned on briefly at the beginning of each acquisition to provide a reference for identifying half sample phase misalignment between channels due to an ambiguity in the clocking, and was corrected during correlation by applying the necessary time delay filter. The tone burst signal was provided by an Agilent 33250A arbitrary waveform generator which was remotely controlled over GPIB (general purpose interface bus) by adding the appropriate commands to the data acquisition LabView virtual instrument (vi) file.

The LabView vi was also updated to communicate with the Arecibo telescope control computer that was to be used as the master data acquisition computer at the observatory. The vi was further modified to work with two additional ADC cards (eight additional input channels) to provide an option for acquiring limited data with the dual polarized array. The updated 20-channel system was used to acquire the majority of the Arecibo data. The two additional cards were never used since the 40-channel system was available for the dual polarized array experiments.

B.1.2 1.25 MHz Receiver Cards

In preparation for the Arecibo experiment the receiver electronics were upgraded from SMA connectorized modular components in the bulky receiver boxes to surface mount and through hole components on a printed circuit board (PCB). The primary motivation for this change was to save space and provide reasonable means for expansion from 20 receiver chains

Figure B.3: Rack mount card cage holding the receiver cards. The ten receiver cards were installed on rails so that they could easily slide in and out of position. Ample room was provided to install cables on both the front and back side.

to 40. The performance requirements remained the same. The boards (shown in Figure 3.11) were designed by Vikas Asthana using Advanced Design System. They were constructed as grounded co-planar waveguide (GCPW) to provide increased isolation between components and the multiple channels on each PCB. When possible, the surface mount equivalent of the original plug-in component was used for the receiver cards. The receiver boxes contained the electronics for two separate receiver channels excluding the final IF filter and amplifier. In contrast, each PCB comprised the electronics for four separate receiver channels including the final IF filters and amplifiers. Four-way splitters were added to each PCB to distribute the LO signals. Because of the low frequency of the final IF, the two final amplifier stages were changed to use op-amp circuits that allowed for the possibility of easily varying the gain of each channel to a uniformly desired value. A more detailed description of the development of the receiver cards, including schematics and a parts list, is found in Vikas Asthana's thesis [91].

A 6U rack mount card cage was purchased from Elma Electronic Inc. to hold the receiver cards. It contains 10 card guides uniformly spaced across its width to hold the cards and it is open in the front and back to allow input and output cables to easily be connected. It is lightweight and can be bolted to the vertical rails of a standard rack. A picture of the cage holding the 10 receiver cards is shown in Figure B.3.

B.1.3 Receiver Rack

The receiver rack did not need to be shielded because it did not contain noisy digital components that would interfere with the experimental observations. The rack pictured in Figure B.4 contained the ten receiver boxes (used as backup spares); the LO distribution network, which consists of amplifiers and splitters to deliver the LO signals to each receiver; the receiver cards in the rack mount card cage; and a series of power supplies to provide DC power to the array and receiver components. There was a 15 volt power supply dedicated

Figure B.4: The receiver rack mounted to the base of the rotating floor. The left image shows the closed rack with IF cables heading up through the ceiling to the data acquisition rack. The right image shows the receiver cards, power supplies, and receiver boxes within the rack.

to the LNAs, a 12 volt supply for the LO amplifiers, a 5 volt supply for the receiver cards, and another 12 volt supply for the receiver boxes. The receiver boxes also use a 5 volt power supply, but since the receiver boxes and cards would never be in operation at the same time a single 5 volt supply was sufficient. Non-switching power supplies were chosen so that no additional interference would be introduced to the array.

The data signals at the output of the receiver cards were carried up to the data acquisition rack by 100-foot lengths of LMR-195 coaxial cable (SMA male to BNC male). The two LO mixing signals from the data acquisition rack were provided to the receiver rack by 100-foot lengths of LMR-400 coaxial cable (N-type male to N-type male).

B.1.4 Shielded Data Acquisition Rack

The double-wide shielded rack shown in Figure B.5 was purchased with the intent that it would hold both the 40-channel and the 20-channel DAQ systems. The 20-channel system was originally meant to be used as a backup in the event that there was a problem with the 40-channel system and it was necessary to switch between the two. As mentioned, the 40-channel system was ready for the second campaign and we were able to conveniently switch from the 20-channel system and acquire 40 channels of data using the dual polarized array. The rack contained two halves that connected at the center but could be easily separated for easier transportation.

The data acquisition rack contained two Agilent 8648D signal generators to supply the LO mixing signals; three Agilent 33250A arbitrary waveform generators to provide clock, trigger, and tone burst signals; the Dell server computer, breakout boxes, and IF trays for the 20-channel system; the ASA server computer, five node computers, and 10 gigabit ethernet switch for the 40-channel system; an uninterrupted power supply (UPS); and an ethernet fiber modem. The tone burst signal was setup to work with either system. The UPS was included to avoid losing power to the equipment since the Arecibo Observatory was notorious

Figure B.5: The double-wide data acquisition rack contains both the 40-channel system and old 20-channel system for backup purposes. It was located in the data acquisition room in the Gregorian dome. Parts of the 40-channel system were removed to continue software development.

for having frequent power losses. Digital communication between computers in the control room and those in the Gregorian dome was to be done over fiber optic cable, so a fiber modem was necessary to convert the signal to the RJ-45 ethernet cable used by the computers.

Bulkhead feed-through connectors were used to pass signals through the walls of the rack in order to maintain shielding. These included 40 female BNC to female BNC connectors for the data signals, two female N-type to female N-type connectors for the LO signals, two female to female fiber optic connectors for each strand of the fiber optic cable, and an AC power feed-through capacitor. With the exception of the AC filter, which was installed at the base of the rack, all other connectors were installed on one of the top panels of the rack. The top panel was used because of the limited space in the dome's data acquisition room and the uncertainty regarding the placement of the rack relative to other equipment and walls.

Since the rack was shielded and was located in a non-cooled room, fans were used to circulate air within the rack. The rack came with two fans installed on the bottom of each half. Additional fans were mounted to the air holes at the top of each door to increase air flow within the rack.

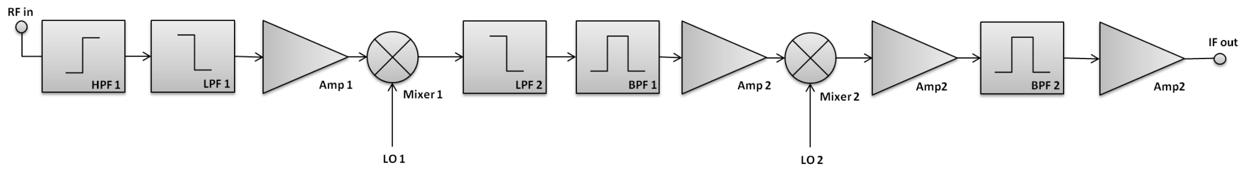

Figure B.6: Block diagram of the 50 MHz receivers.

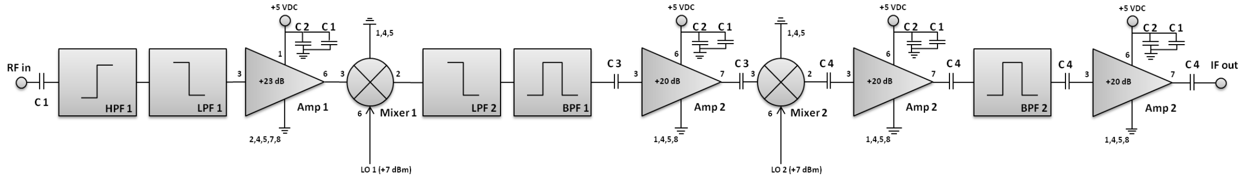

Figure B.7: Circuit diagram of the 50 MHz receivers showing pin assignments, amplification values, and blocking and coupling capacitors.

B.1.5 Experiments

The main goal of the experiment was to collect PAF data with the array positioned at a series of pointings so that the data could be combined and used to evaluate the expected performance of a much larger array and determine the associated field of view. This objective was achieved and additional results were also obtained. Among these are verification of the performance of an active impedance matched PAF, PAF sensitivity measurements, calibrator and weak source images, FOV measurements, and steered PAF beams. These results and additional details about the feasibility study can be found in [82, 91–95].

B.2 50 MHz Receiver Cards

In 2010 work began on a 64-channel FPGA-based data acquisition system that included a correlator and beamformer. This new system was to interface with existing and future BYU antenna arrays. The specifications for the 64-channel system are outlined in [81,91] and include a sample rate of 50 MHz with a per-input-channel bandwidth of 20 MHz. The 1.25 MHz receivers were designed for a much lower sample rate and consequently a new receiver was needed.

The 50 MHz receiver cards were designed as an upgrade of the original GCPW receivers. For this reason the general layout remained the same and many of the same components were used. The gain requirements did not change and many of the same design principles used in the previous receiver continued. However, with the increase in sample rate and bandwidth it was necessary to replace some of the filters and amplifiers. The new data acquisition system can accept 64 inputs, so 16 four-channel receiver boards were built. In addition to this quantity, two more boards were built to be used as backups in the event that there was a problem with one of the 16 boards. A block diagram of the 50 MHz receiver is shown in Figure B.6, and a circuit diagram is given in Figure B.7. A component list is provided in B.1.

Table B.1: Component list for the 50 MHz receivers.

Identifier	**Part Number**	**Distributor**	**Manufacturer**	**Description**
Amp1	VNA-28	Mini-Circuits	Mini-Circuits	$500 - 2500$ MHz amplifier, 22.8 dB gain, 5V DC
Amp2	AD8354ACPZ	Digi-Key	Analog Devices	$1 - 2700$ MHz amplifier, 20 dB gain, 5V DC
BPF1	SF480	Vanlong	Vanlong	480 MHz SAW filter, 22 MHz bandwidth
BPF2	3010	KR Electronics	KR Electronics	37.5 MHz bandpass filter, 20 MHz bandwidth
C1	490-1512-1-ND	Digi-Key	Murata Electronics	0.01μF capacitor
C2	490-3302-1-ND	Digi-Key	Murata Electronics	4.7μF capacitor
C3	490-1521-1-ND	Digi-Key	Murata Electronics	0.033μF capacitor
C4	490-3300-1-ND	Digi-Key	Murata Electronics	0.47μF capacitor
HPF1	HFCN-1200+	Mini-Circuits	Mini-Circuits	1180 MHz highpass filter
LPF1	LFCN-1500+	Mini-Circuits	Mini-Circuits	1825 MHz lowpass filter
LPF2	LFCN-1000+	Mini-Circuits	Mini-Circuits	1300 MHz lowpass filter
Mixer1	ADE-35	Mini-Circuits	Mini-Circuits	+7 dBm mixer, RF: 1.6-3.5 GHz, IF: DC-1.5 GHz
Mixer2	ADE-2+	Mini-Circuits	Mini-Circuits	+7 dBm mixer, RF: 5-1000 MHz, IF: DC-1.0 GHz
Splitter1	BP4P1+	Mini-Circuits	Mini-Circuits	1500-2500 MHz 4-way power splitter
Splitter2	SCA-4-10+	Mini-Circuits	Mini-Circuits	5-1000 MHz 4-way power splitter
Atten1	PAT121CT-ND	Digi-Key	Susumu	1 dB attenuator
Atten2	PAT122CT-ND	Digi-Key	Susumu	2 dB attenuator
Atten3	PAT123CT-ND	Digi-Key	Susumu	3 dB attenuator
DCsocket	538-39-30-1020	Mouser	Molex	Molex connector socket
DCplug	538-39-01-2020	Mouser	Molex	Molex connector plug
RFconnector	ACX1232-ND	Mouser	Amphenol Connex	SMA right-angle female through-hole connector

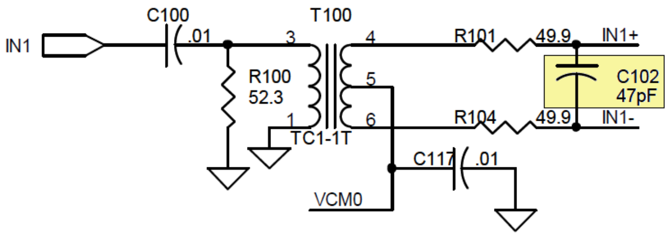

Figure B.8: Schematic of the pre-ADC circuitry for a single ADC input. The capacitor in the yellow box must be removed to allow operation in a higher Nyquist zone.

The boards were designed, fabricated and assembled in 2011-2012. The design was completed using Advanced Design System (ADS). Fabrication was done by Quick Turn Circuits in Salt Lake City, UT. Assembly of an initial test board was provided by Automated Assembly in Lehi, UT. Assembly of the 18 revised, final boards was provided by Trident Manufacturing in Salt Lake City, UT.

B.2.1 Nyquist Sub-Sampling

For a 50 MHz sample rate the Nyquist frequency is 25 MHz. This puts the center of the passband, or the first Nyquist zone, at 12.5 MHz. The second, third, and fourth Nyquist zones are centered at 37.5, 62.5, and 87.5 MHz , respectively. Ideally we would like to operate in the first Nyquist zone but the design of the anti-aliasing filter becomes easier as we get farther away from baseband.

Operating in a higher Nyquist zone may reduce costs but there are limitations based on the characteristics of the ADCs. Sampling a high frequency signal, relative to the ADC sample rate specifications, that is changing too rapidly may lead to inaccurate conversion and undesirable signal attenuation. The 64-input ADC board has a built in anti-aliasing filter that must be removed before testing for this problem. A schematic showing the pre-ADC circuitry for a single input is shown in Figure B.8, where the capacitor to be removed is indicated by the yellow box. This is repeated for each of the 64 ADC inputs. After making this modification, signals from DC to 100 MHz were sampled using the ADCs and there was only 1-2 dB attenuation noted for the largest signal, confirming the ability to operate in any of the above mentioned Nyquist zones.

B.2.2 Changes from 1.25 MHz Receiver Design

Much of the existing receiver design carried over to the new design. However, a few changes were needed to accommodate the increased sample rate and bandwidth. The SAW filter, some of the amplifiers, the final IF filter, and a mixer needed to be changed. Also, some pads for surface mount attenuators were added.

SAW Filter

The Vanlong SF400 used in the 1.25 MHz receiver design, as well as in the receiver boxes, has a 3 dB bandwidth of 5 MHz and would not satisfy the 20 MHz bandwidth

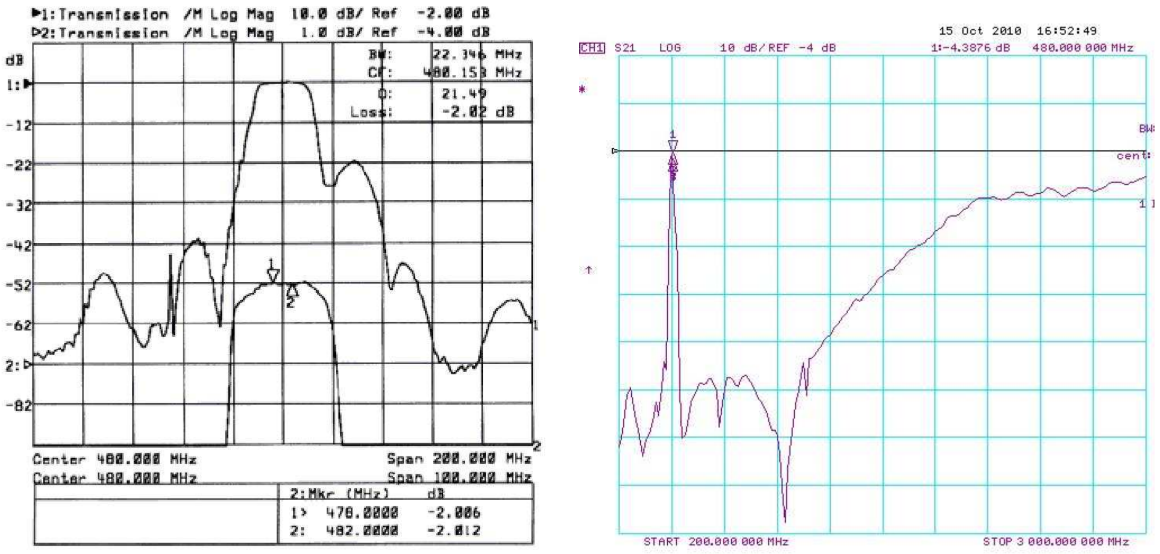

(a) Standard performance plot for the SF480.

(b) Higher frequency performance of the SF480.

Figure B.9: Frequency response plots for the SF480 SAW filter. (a) is taken from the Vanlong data sheet. (b) was requested to confirm that a low pass filter was needed to block mixing signals above 3000 MHz.

specification of the 64-channel system. To maintain familiarity and to limit the effect on other parts of the receiver, it was desired to keep the center frequency of this bandpass filter around 400 ± 100 MHz. The Vanlong SF480 SAW filter operates at a center frequency of 480 MHz and has a 3 dB bandwidth of 22 MHz. It is housed in a through-hole package and is relatively inexpensive.The standard performance plot for the SF480, shown in Figure B.9(a), shows that it will pass the desired frequencies around 480 MHz, but there was some concern about its ability to block the sum of the mixing frequencies which would be around $3000 - 4000$ MHz. Additional measurements seen in Figure B.9(b) show that around 1100 MHz the response of the filter begins to increase and that it is unable to sufficiently block frequencies above 3000 MHz as needed. This same behavior was present with the SF400 in the 1.25 MHz design and was handled by adding a low pass filter in conjunction with the SAW filter [77]. Likewise this new design includes a low pass filter in order to avoid passing the higher frequencies that result from the first mixing stage.

Amplifiers

The final two amplifier stages in the 1.25 MHz PCB receivers used op-amp circuits. For this new design these were not able to give the necessary gain at the final IF frequency. Conveniently, the second stage amplifier (Analog Devices part number AD8354ACPZ) provides the necessary performance for both of the final amplifier stages. It operates from 1 MHz up to 2700 MHz and has a nominal gain of $19 - 20$ dB. This same amplifier then is used in the 2nd, 3rd, and 4th stages and reduces complexity by eliminating the capacitors and resistors needed for the op-amp circuits.

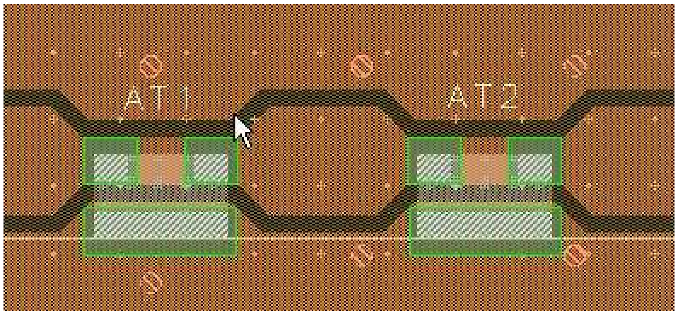

Figure B.10: ADS layout of the two attenuator pads that were added to each receiver chain to be able to adjust the gains of all channels to be uniform.

IF Filter

The final IF filter needed to be replaced to operate at the higher sampling frequency. A custom designed filter centered at 37.5 was built by KR Electronics. The specifications for each filter are 3 dB maximum insertion loss and 40 dB attenuation at ±15 MHz from the center of the band. The part is a surface mount package with dimensions of $1.5'' \times 0.5'' \times 0.3''$ (L×W×H). This filter is the biggest component on the board in terms of both its footprint and height from the surface of the board.

Pads for Attenuators

Due to slight variations in the gain of each component in the receiver design, each receiver chain will have a slightly different overall gain. In the 1.25 MHz PCB design this could be taken into account by changing the gain of the op-amp circuits. In order to make this option available on the 50 MHz receiver cards, two sets of solder pads were added to the end of each receiver chain. The pads were left connected on the signal path in the design. An attenuator may be soldered to the board after manually cutting the trace between the pads. With two sets of pads, some combination of available attenuators may be added to give the desired level of gain. The ADS software layout of the two attenuator pads is shown in Figure B.10.

B.2.3 Tests

A prototype board, shown in Figure B.11, was first tested to make sure there was not a flaw in the design, fabrication, or assembly, and to measure the channel gains and cross-coupling. The test was setup using three Agilent 8648D signal generators to provide the input signal and each of the two LO signals. For each test a single input channel was used and the output signal was viewed on a spectrum analyzer, with all other inputs and outputs terminated with 50Ω terminators. The input signal was set to −70 dBm at 1600 MHz, with the first LO set to 13 dBm at 2080 MHz and the second LO set to 13 dBm at 442.5 MHz. The results of the test are given in Table B.2 with the first channel receiving the input signal. Channel 3 was excluded from the tests because the first mixer was not working properly. It was later discovered that two ground vias connecting the top and bottom of the

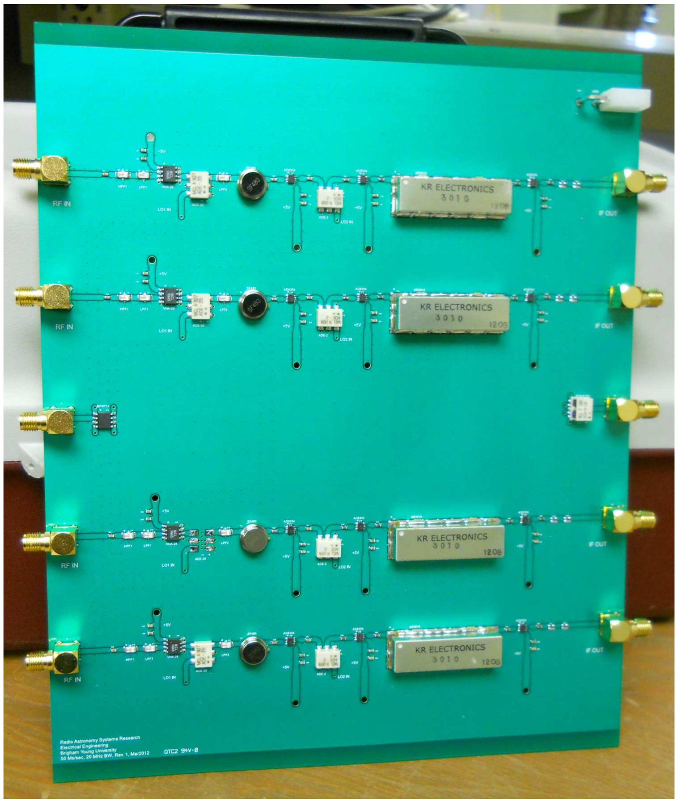

Figure B.11: The 50 MHz test receiver card. A single board was assembled then tested to verify that the design was done correctly. Test results were positive and showed that gain and cross-coupling expectations were met.

Table B.2: Receiver board gain and cross-coupling test results with input on channel 1.

	channel 1	**channel 2**	**channel 3**	**channel 4**
output power (dBm)	-11.15	-50.9	not tested	-53.5
net power gain (dBm)	58.85	19.1	not tested	16.5

board were placed on the LO signal trace. These needed to be drilled out on each of the boards before each channel 3 worked correctly.

The results of the test show that there is a gain of nearly 60 dBm in channel 1 and isolation between channels of about 40 dBm. On the spectrum analyzer the 20 MHz pass band centered on 37.5 MHz is very noticeable (see Figure B.12). The output power for the channels not receiving the input signal is very near to the noise floor in the pass band, which was consistently measured at at about −54 dBm. Similar results were obtained with all the channels tested.

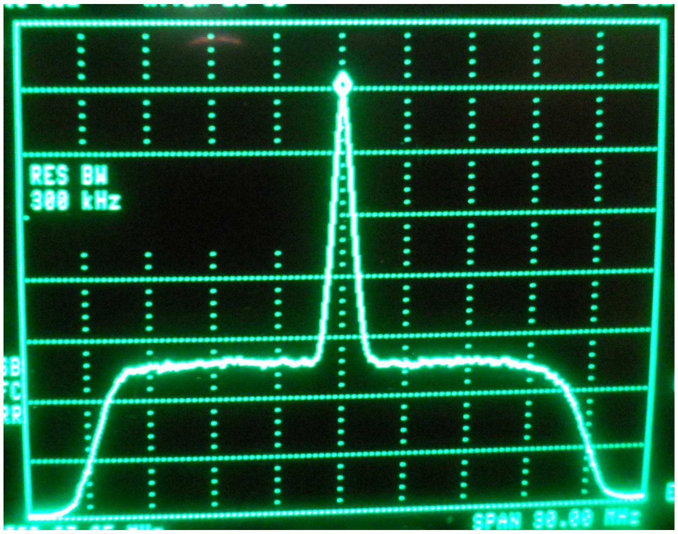

Figure B.12: Frequency spectrum of the 50 MHz receiver card showing amplified signal and flat noise response. The signal appears at the output of channel 1. The noise floor sits at about −54 dBm and has the required bandwidth.

Appendix C

Procedure for Aligning Multiple Calibration Grids

When completing an analysis involving multiple calibration grids it is important to confirm that they are properly aligned, or to correctly identify a misalignment so that it can be accounted for. Misalignment is possible in the case of telescope pointing errors, or if different initial elevation and azimuth settings were applied the grid scanning software when the calibration data was collected. There are many perceivable ways of checking the alignment, one of which is described here.

For a rough alignment check between two calibration grids we measure the difference between array response vectors from the two grids. First, identify the $m \times m$ response vectors that form a small grid around the supposed center of the first grid (to which a second grid is to be aligned), and designate them as the columns of a matrix $\mathbf{A}$. Next, identify the $m \times m$ response vectors from the second grid, centered on the supposed center, and designate them as the columns of a matrix $\mathbf{B}$. Compute the Frobenius norm of the difference between matrices $\mathbf{A}$ and $\mathbf{B}$. Repeat this computation for different sets of response vectors from the second grid, centered on different points immediately surrounding the supposed center of the grid. Identify the set of response vectors from the second grid that results in the smallest normed difference to determine the amount of misalignment present between the two calibration grids. If a misalignment is identified, the center of the second grid must be shifted accordingly so that the response vectors for the same point in each grid are actually steering in the same direction.

Identifying misalignment on a smaller scale can be done by interpolating between response vectors and then repeating the above procedure.