

Development of an Experimental Phased Array Feed System
and Algorithms for Radio Astronomy

Jonathan C. Landon

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 2011

Copyright © 2011 Jonathan C. Landon

All Rights Reserved

ABSTRACT

Development of an Experimental Phased Array Feed System and Algorithms for Radio Astronomy

Jonathan C. Landon

Department of Electrical and Computer Engineering, BYU
Doctor of Philosophy

Phased array feeds (PAFs) are a promising new technology for astronomical radio telescopes. While PAFs have been used in other fields, the demanding sensitivity and calibration requirements in astronomy present unique new challenges. This dissertation presents some of the first astronomical PAF results demonstrating the lowest noise temperature and highest sensitivity at the time (66 Kelvin and $3.3\text{m}^2/\text{K}$, respectively), obtained using a narrowband (425 kHz bandwidth) prototype array of 19 linear co-polarized L-band dipoles mounted at the focus of the Green Bank 20 Meter Telescope at the National Radio Astronomy Observatory (NRAO) in Green Bank, West Virginia. Results include spectral line detection of hydroxyl (OH) sources W49N and W3OH, and some of the first radio camera images made using a PAF, including an image of the Cygnus X region. A novel array Y-factor technique for measuring the isotropic noise response of the array is shown along with experimental measurements for this PAF. Statistically optimal beamformers (Maximum SNR and MVDR) are used throughout the work. Radio-frequency interference (RFI) mitigation is demonstrated experimentally using spatial cancelation with the PAF.

Improved RFI mitigation is achieved in the challenging cases of low interference-to-noise ratio (INR) and moving interference by combining subspace projection (SP) beamforming with a polynomial model to track a rank 1 subspace. Limiting factors in SP are investigated including sample estimation error, subspace smearing, noise bias, and spectral scooping; each of these factors is overcome with the polynomial model and prewhitening. Numerical optimization leads to the polynomial subspace projection (PSP) method, and least-squares fitting to the series of dominant eigenvectors over a series of short term integrations (STIs) leads to the eigenvector polynomial subspace projection (EPSP) method. Expressions for the gradient, Hessian, and Jacobian are given for use in numerical optimization. Results are given for simulated and experimental data, demonstrating deeper beampattern nulls by 6 to 30dB.

To increase the system bandwidth toward the hundreds of MHz bandwidth required by astronomers for a fully science-ready instrument, an FPGA digital backend is introduced using a 64-input analog-to-digital converter running at 50 Msamp/sec and the ROACH processing board developed at the University of California, Berkeley. International efforts to develop digital backends for large antenna arrays are considered, and a road map is proposed for development of a hardware correlator/beamformer at BYU using three ROACH boards communicating over 10 gigabit Ethernet.

Keywords: phased array feeds, radio astronomy, isotropic noise response, array Y-factor method, interference canceling, adaptive array processing, adaptive beamforming, subspace projection, subspace tracking, covariance estimation, CASPER, ROACH, FPGA, astronomical back ends

ACKNOWLEDGMENTS

I am deeply grateful for the generous support of Professors Brian Jeffs and Karl Warnick who gave me the opportunity to do exciting research in exciting locations. More than just guiding my research, they encouraged me to believe in myself. For these experimental results I am indebted to students, faculty, and staff at Brigham Young University (BYU), the National Radio Astronomy Observatory in Green Bank, West Virginia (NRAO-GB), and the University of California, Berkeley.

At BYU, Drs. Jeffs and Warnick directed the research. James Nagel laid the foundation of this work, designing the original downconverter boxes and 7-element array feed. As a student at the University of Utah, James granted generous access to the anechoic chamber critically important to this work. My thanks to Micah Lillrose for the software correlator and digital back end; Jacob Waldron for array construction, hot/cold ground screen construction, and early planning for the Green Bank 20 Meter Telescope tests; Jiyoung Son for help with experiments, David Jones for his mutual coupling work, Alan Stemmons who stayed up through the night to prepare for the initial FEB mounting, Michael Elmer for his dish tipping work and experimental help, Vikas Asthana for his help during the anechoic chamber tests and CASPER development, David Carter who correlated much of the 2008 data, and Taylor Webb for his insights.

At NRAO-GB, I am grateful for Roger Norrod's mentoring through my time as a summer student and for his involvement in all of our experiments at the observatory. Drs. Rick Fisher and Richard Bradley did the first PAF work at NRAO in the mid-1990's, opening the way for this work. Dr. Fisher was also extremely helpful during the planning and execution of these experiments, providing valuable insights from an astronomer's perspective as well as pointing of the telescope itself. Thanks to technicians Bill Shank, Bob Simon, Ken Ward, and Dave Woody for their many hours preparing for and helping with experiments; Pete Chestnut, Bob Anderson, Mike Hedrick, and Gary Anderson for their help in constructing parts of the experimental platform and preparing for the outdoor antenna range test.

At UC Berkeley, I am grateful to the late Professor Don Backer for sponsoring me for a six month visiting studentship and for his kindness and belief in me; Dr. Aaron Parsons and Jason Manley for their many hours teaching me about the CASPER tools and basic Linux usage; Dr. Dan Wertheimer, Matt Dexter, and Dave MacMahon for their friendship and assistance in managing the PAPER correlator. Jack Hickish at Oxford has provided tremendous assistance getting the x64adc working, providing several example designs and the yellow-block interface itself (with help from Dave George at the Karoo Array Telescope).

This dissertation is for my dad who decided that instead of finishing his dissertation he would move to Salt Lake City, get another Masters degree, and work late supporting his family. He has kept my car running so I could focus on school. And it is for my mom who has been very supportive of the adventures along the way – research trips to Green Bank, West Virginia; Berkeley, California; The Netherlands; and Puerto Rico. Thank you!

Table of Contents

List of Tables	xvii
List of Figures	xxi
Chapter 1 Introduction	1
1.1 Background on Astronomical Radio Telescopes	2
1.1.1 Current Radio Telescopes	2
1.1.2 Next Generation Reflector Telescopes: Multi-Beam Feeds	4
1.2 Literature Review	7
1.2.1 Scientific Motivation for PAFs in Radio Astronomy	7
1.2.2 Previous Developments for PAFs in Radio Astronomy	8
1.2.3 Calibration and Beamforming	9
1.2.4 PAF Experiments	9
1.2.5 RFI Mitigation and Subspace Projection	11
1.3 Dissertation Outline	12
1.4 Contributions and Publications	13
1.4.1 Contributions	14
1.4.2 Publications	16
Chapter 2 Background	19
2.1 Notation	19

2.2	Astronomy Background	20
2.2.1	Fundamentals	20
2.2.2	Spectral Line Sources	22
2.2.3	Telescope Parameters	23
2.3	Beamforming Background	26
2.3.1	Signal Model	26
2.3.2	Adaptive PAF Beamforming	27
2.3.3	Active Interference-Canceling Beamformers	29
2.4	Location: NRAO – Green Bank, WV	29
Chapter 3 Development of an Experimental Phased Array Feed for Radio Astronomy		33
3.1	Introduction	33
3.2	System Layout	34
3.3	RF Front End	36
3.3.1	Green Bank 20 Meter Telescope	36
3.3.2	Phased Array Feed	38
3.3.3	Analog Front End Electronics	40
3.3.4	Front End Box (FEB)	41
3.4	Telescope Pedestal	43
3.4.1	IF and LO	43
3.4.2	Digital Back End	44
3.5	Reference Antenna	46
3.6	Jansky Lab	46
3.7	Conclusions	47

Chapter 4	First PAF Results on the Green Bank 20 Meter Telescope	49
4.1	Introduction	49
4.2	Calibration	51
4.3	Proof-of-Concept PAF Demonstration	52
4.3.1	OH Spectral Line Detection	53
4.3.2	Beam Sensitivity	54
4.3.3	Dish Illumination Control	56
4.3.4	PAF De-focusing	57
4.3.5	Adaptive Spillover Noise Control	58
4.3.6	RFI Mitigation Demonstration	58
4.4	Array Noise Response Characterization	60
4.4.1	Significance of $R_{\text{ext,iso}}$ in Connecting Active Arrays and Passive Antennas	61
4.4.2	A Novel Y-factor Approach to Obtain $R_{\text{ext,iso}}$	62
4.4.3	Experimental $R_{\text{ext,iso}}$ Measurements	64
4.5	Mutual Coupling Measurements	68
4.5.1	Element Radiation Patterns with Mutual Coupling	68
4.5.2	Increased System Noise Temperature Due to Mutual Coupling	69
4.6	Sensitivity	73
4.6.1	Using $R_{\text{ext,iso}}$ in Sensitivity Calculations	74
4.6.2	Measured Sensitivity Results and Comparisons to Single Feed Telescopes .	75
4.6.3	Sensitivity Degradation for Steered Beams	76
4.7	Radio Camera Images	77
4.7.1	Multiple Tile Image Mosaics	78
4.7.2	Demonstrating RFI Mitigation in a Radio Camera Image	79
4.8	Conclusions and Future Work	81

Chapter 5	Model-Based Subspace Projection Beamforming for Deep Interference Nulling	83
5.1	Introduction	83
5.2	Beamforming Issues for Moving Interference	85
5.2.1	Signal Model	85
5.2.2	Subspace Projection Interference Canceling	86
5.2.3	Factors Limiting Cancelation Performance	88
5.3	A Polynomial Covariance Motion Model for Subspace Projection	93
5.3.1	Background	93
5.3.2	Problem Formulation	94
5.3.3	Solution by Numerical Optimization	95
5.3.4	Optimizer Initialization and Approximate Solution	98
5.3.5	Cramer-Rao Lower Bound (CRB) Analysis	100
5.4	Results	105
5.4.1	Simulation Performance Analysis	105
5.4.2	Comparison with Subspace Tracking Spatial Projections	108
5.4.3	Real Data Experiments	111
5.5	Conclusions	114
Chapter 6	CASPER-Based Digital Back End for the BYU PAF	117
6.1	Introduction	117
6.1.1	International CASPER Collaboration	118
6.1.2	BYU Involvement with CASPER	118
6.2	Current Beamformer/Correlator Development Efforts at BYU	119
6.2.1	System Description	119

6.2.2	BYU Real-Time Beamformer Capabilities	123
6.3	Feasibility Assessment of a CASPER Platform AO40 Back End	127
6.3.1	Background	127
6.3.2	System Performance Specifications	128
6.3.3	Overview	128
6.3.4	Feasibility Assessment	129
6.3.5	Alternative Hardware Approaches and Related Development Efforts	132
6.3.6	Conclusion	133
6.4	Future Development at BYU of a 64-Input Correlator/Beamformer Using CASPER Hardware	134
6.4.1	Phase 1: Real-Time Spectrometer	136
6.4.2	Phase 2: Streaming Raw Samples to Disk	137
6.4.3	Phase 3: Beamformer Development	140
6.4.4	Phase 4: Hardware Correlator Development	142
6.4.5	Phase 5: Imaging Software Development and Streaming Raw Beams to Disk	144
6.4.6	Tradeoffs and Alternative Architectures	145
6.5	Conclusions	146
Chapter 7 Conclusions and Future Work		149
7.1	Conclusions	149
7.2	Future Work	151
Bibliography		155

List of Tables

4.1	Noise temperature measurements (all values given in Kelvin)	71
4.2	Sliding dipole noise temperature measurements, in Kelvin	72
4.3	Measured noise temperatures in Kelvin, for array with LNA on center dipole and various terminations on other dipoles	73
4.4	Measured and modeled peak beam sensitivity, system temperature, and aperture efficiency for the 19 element prototype dipole array	75
4.5	Sensitivity, system temperature, and aperture efficiency for several single feed radio telescopes	76
6.1	Summary of existing array correlator projects using CASPER FPGA systems . . .	130

List of Figures

1.1	The primary advantage of FPA telescopes is increased field of view provided by multiple, simultaneously formed beams.	5
2.1	Overhead view of the site.	31
3.1	Receiving system block diagram, showing communication among the various sub-systems.	35
3.2	The Green Bank 20 Meter Telescope.	37
3.3	The BYU PAF mounted at the focus of the Green Bank 20 Meter Telescope.	38
3.4	Dipole dimensions and numbering.	39
3.5	RF receiver electronics inside the FEB.	40
3.6	Frequency downconverter box.	41
3.7	The FEB viewed from behind with IF cables.	42
3.8	Shielded rack holding the digital back end and signal generators for LO signals.	44
3.9	3.6m reference antenna.	47
4.1	Spectral line detection of OH source W49N.	53
4.2	Beam sensitivity comparison for center dipole and MaxSNR formed beam.	55
4.3	Slice across CygA comparing center element with fixed MaxSNR formed beam.	57
4.4	Elevation slice for spillover optimization.	59
4.5	Demonstration of RFI mitigation using SP beamforming to remove a CW tone.	60
4.6	Overhead view of ground screen.	65

4.7	Array feed with copper mesh clamped down around it.	66
4.8	Absorber being lowered into the ground screen.	67
4.9	Antenna range measurements.	69
4.10	Sliding dipole tests assembly.	72
4.11	Sensitivity of the BYU PAF as a function of beam steering angle.	77
4.12	Mosaic image of W49N.	78
4.13	Cygnus X region.	79
4.14	RFI mitigation for a radio camera image of W3OH.	80
5.1	IRR and residual interference power dependence on STI length and INR level for a moving interferer.	89
5.2	Limiting factors in SP.	91
5.3	Spectral scooping illustration.	92
5.4	Phase normalization examples.	99
5.5	CRB vs Monte Carlo trials of eigenvector polynomial algorithm.	104
5.6	CRB and Monte Carlo trials for different set of polynomial coefficients.	104
5.7	Comparison of SP, EPSP and PSP.	106
5.8	Performance of SP and EPSP for various input INRs and motion rates.	107
5.9	Comparison of PASTd, SP, EPSP and PSP.	110
5.10	The 19-element array mounted on a rotating table in the anechoic chamber.	111
5.11	Real data experiments showing SP improvement for a slow-moving low-INR CW tone interferer.	113
5.12	Real data experiments showing SP improvement for a fast-moving high-INR interferer.	113
6.1	Interior view of the 64-input x64adc connected to a CASPER ROACH FPGA board.	120
6.2	Transformer circuit with anti-alias capacitor and block diagram of baseband sub-sampling.	121

6.3	Close-up front view of SMA bulkhead for the ROACH x64adc system.	123
6.4	Equipment rack configuration with ROACH and 64 RF input connections for the x64adc, a second ROACH which will perform beamforming operations, and host computers and data file server.	124
6.5	Design environment snapshot.	124
6.6	Straw man block diagram of possible 300 MHz AO40 back end.	131
6.7	Overview of correlator/beamformer full system using the x64adc and three ROACH boards.	135
6.8	Initial hardware demonstration: real-time spectrometer using the x64adc.	137
6.9	Complex-valued frequency channelized data is sent without accumulation to the 10GbE packetizer for streaming to disk.	138
6.10	With a rack of receiving computers to receive the frequency channelized data, the full bandwidth can be streamed to disk.	139
6.11	ROACH 1 streams data over XAUI to ROACH 3 where beams are formed.	141
6.12	Beamformer detail.	142
6.13	ROACH 1 passes channelized data over XAUI to ROACH 2 for real-time correlation used in finding beamformer calibrators, and to ROACH 3 for beamforming. . .	143

Chapter 1

Introduction

This dissertation presents an experimental phased array feed (PAF) system for radio astronomy imaging and radio frequency interference (RFI) mitigation. PAFs are a promising new technology in radio astronomy because of their ability to form multiple simultaneous beams, improving sky survey speed as compared to a traditional radio telescope with a single horn feed that produces a single image pixel for each telescope pointing. In a PAF several small antennas are placed at the focus of a large reflector antenna and signals received at each array element are combined to form multiple beams. In addition to improving sky survey time, PAFs offer the flexibility to adaptively shape the beams to optimize sensitivity and steer beam pattern nulls to cancel RFI.

Two major goals are accomplished in this dissertation: moving PAF technology from purely theoretical to actual demonstration of astronomical detections, and improving RFI mitigation by producing beam pattern nulls an order of magnitude deeper than previously achieved. The dissertation organization mirrors the signal flow through a radio telescope, starting with the RF front end electronics in Chapter 3 and ending with the digital back end in Chapter 6. The middle chapters discuss experiments run with the PAF and algorithm development for canceling RFI in difficult signal scenarios.

The experimental results presented here are at the forefront of PAF development in radio astronomy. Although PAFs have been used in communications and radar for many years, the technology is still new to radio astronomy because of the extreme requirements for telescope sensitivity and calibration. Very stable and well-characterized beam patterns are required in astronomy for the techniques used in detecting weak sources and differentiating a weak source in the main beam from a strong source in a sidelobe. Beamforming techniques adequate for other applications are inadequate for astronomy in terms of stability, polarization purity, dish illumination, and radiation pattern shape. For example, noise sources such as spillover noise from warm ground seen by

the feed and entering the receiver through over-illuminating the reflector are important for astronomy but can be neglected in other applications. Experimental demonstration of PAF beamforming techniques are required for astronomers to gain confidence in this new technology.

The RFI mitigation algorithm development of Chapter 5 is an important advancement in the understanding of subspace projection limitations in dealing with weak or moving interference. Since astronomical signals of interest (SOIs) are below the receiver noise floor and may require integrations on the order of minutes or hours for detection, low interferer-to-noise ratio (INR) sources must be canceled to well below the noise floor. The algorithm presented here considers such sources and presents results in simulated and experimental data.

As the astronomical community works toward incorporating PAFs into the next generation of scientific instruments, serious consideration must be given to the hardware requirements of real-time data processing for observing bandwidths on the order of hundreds of MHz, with an independent receiver chain for each PAF element. Chapter 6 presents initial work on an FPGA-based digital back end for real-time PAF beamforming. The CASPER [1] gateway design tools described here are used by several groups around the world, with hardware correlator progress described in that chapter, along with alternative hardware platforms in development for other astronomical arrays.

This chapter begins with an introduction to astronomical radio telescopes PAFs in Section 1.1, then reviews the literature on astronomical PAF developments and RFI mitigation in Section 1.2 discussing concurrent efforts by other groups internationally in PAF work. An outline of the rest of the dissertation is given in Section 1.3 and dissertation contributions are highlighted in Section 1.4.

1.1 Background on Astronomical Radio Telescopes

1.1.1 Current Radio Telescopes

Two major types of telescopes are currently used in radio astronomy: single dish antennas and synthesis imaging arrays of multiple dish antennas. In single dish astronomy the telescope steers to one spot on the sky, forming one pixel in an image. To map out a region the telescope is steered over a grid of pointings producing a grid of pixels that form the image. The size of

the image is dictated by the size of the grid the telescope is scanned across. Image resolution is dictated by the pixel size – the larger the antenna the narrower the beamwidth, so the antenna sees a smaller spot on the sky giving finer image resolution. A detailed history and introduction to phenomenology and detection techniques for radio astronomy can be found in the excellent textbook by Kraus [2], and modern techniques for single-dish astronomy can be found in [3].

Synthesis imaging uses a large array of dishes, each looking at the same region of the sky. Each receives roughly the same signal with a phase delay and magnitude scaling. Since the sum of a collection of magnitudes and phases looks like a Fourier Transform, an Inverse Fourier Transform may be used to reconstruct the original image that gave rise to the apparent Fourier Transform that was received, as discussed in [4]. The technique is based on interferometry where each pair of antennas acts as an interferometer sensitive to a single spatial frequency, and Earth rotation provides different effective antenna spacings (projected onto the image plane) throughout an observation. The size of the image is dictated by the smallest main beam among all the antennas and the resolution of the image is dictated by the largest separation between antennas. Synthesis imaging arrays like the Very Large Array (VLA), Very Long Baseline Array (VLBA) and Extended Very Long Baseline Array (EVBLA) form much higher resolution images than a single dish can because their baselines are so much larger than a single dish could be, but they rely on good sampling of the image plane which is difficult because of the large aperture they seek to synthesize. Good references for synthesis imaging are [5], [6], and [4].

Instrumentation has taken a major step forward every few decades with the introduction of new telescopes with significantly better sensitivity, beginning with Karl Jansky's original discoveries in 1933 and the first sky survey at 160 MHz published in 1940 by Grote Reber using the parabolic reflector he built in 1937. The 1950's and 1960's were an era of large reflector telescope construction. Beginning with the establishment of the National Radio Astronomy Observatory (NRAO) in 1956 several large reflectors were built in Green Bank, WV (NRAO-GB), including the 85ft, 300ft, and 140ft, which began observations in 1959, 1962, and 1965, respectively [7]. The 305-meter Arecibo Radio Telescope began operations in 1963 [8]. Preliminary astronomical interferometry work at the same time led to proposals in the late 1960's for the Very Large Array in Socorro, NM which was constructed through the 1970's and officially opened for proposals in 1978 [7]. The following two decades saw advancements centered on millimeter wave astronomy,

synthesis imaging and calibration techniques, and very long baseline interferometry. The next major leap forward in both sensitivity and sky coverage is the Square Kilometre Array (SKA) [9], an internationally sponsored telescope with a total collecting area of 1km^2 conceived during the early 1990's and currently in the process of design specification and funding acquisition.

1.1.2 Next Generation Reflector Telescopes: Multi-Beam Feeds

Large reflector telescopes are desirable for their high sensitivity and narrow main beam that improves image resolution, but the narrow beam also means increased time required for sky surveys. To speed up sky surveys, multi-beam feeds allow the telescope to see multiple places simultaneously while maintaining the image pixel size dictated by the large reflector. There are two main approaches to multi-beam feeds. The first, called a focal plane array (FPA), places an array of single pixel horn antenna feeds near the focus of a large reflector, forming independent beams on the sky. In contrast, a phased array feed (PAF) uses a more densely packed array (on the order of $\lambda/2$ spacing) of smaller antennas with overlapping radiation patterns, whose signals are electronically combined to form beams. Since a PAF is an antenna array near the focus of a reflector it could also be considered a focal plane array, but to avoid confusion in this dissertation FPA will always refer to an array of horn antenna feeds with independent beams as opposed to a PAF where the signals are electronically combined to form beams.

Focal Plane Arrays

The FPA acts like several collocated large reflectors, each seeing an independent part of the sky, illustrated in Figure 1.1. The feeds are in fixed locations, with each element radiation pattern created by appropriate design of the physical shape of each feed to give a proper tradeoff between dish illumination and spillover, and radiation patterns are ideally non-overlapping to place an independent pixel on the sky. Received data labeled $x[n]$ in the figure are treated as P independent signals $x_i[n]$ for $i = 1 \dots P$. The Parkes Multibeam receiver [10], with 13 independent horns placed at the focus of the Parkes Radio Telescope, conducted the HI Parkes All-Sky Survey (HIPASS) from 1997 to 2000 [11]. The Arecibo L-band Feed Array (ALFA) at the focus of the Arecibo Radio Telescope [12, 13] is another example of an astronomical FPA.




Figure 1.1: The primary advantage of FPA telescopes is increased field of view provided by multiple, simultaneously formed beams. Spatial cancellation of interfering signals is also possible, but very deep nulls are required [14].

Phased Array Feeds

In a PAF the electrically small array elements have broad, overlapping radiation patterns, allowing the received signals, $\mathbf{x}[n]$, to be combined with a vector of beamformer coefficients, \mathbf{w} , into a single narrow composite beam, $y[n]$:

$$y[n] = \mathbf{w}^H \mathbf{x}[n].$$

An arbitrary number of beams can be formed from the same received data, $\mathbf{x}[n]$, using different beamformer coefficients, \mathbf{w} . Forming beams that overlap at the 3dB points keeps the computational burden low while sufficiently covering the field of view. The shape of each beam is controlled by the beamformer coefficients, designed according to criteria such as maximizing the signal to noise ratio (SNR), controlling main beam width and dish illumination, or adaptively steering beam pattern

nulls to track interferers. Multiple beams can be formed simultaneously in hardware, or in post-processing with recorded voltage samples.

Tradeoffs

The major tradeoffs between FPAs and PAFs are stability versus adaptability, sparse versus filled field of view coverage, and system and computational complexity. PAFs offer the ability to tailor the beampattern to fit various optimality constraints, adapting to changing spillover noise environments and interference. In an FPA the beampattern is not adaptable but it is usually well known, with the advantage of being able to subtract out its effect more easily. Beamformer calibration is much harder for a PAF because the beam shape depends on receiver phases and gains at each element which fluctuate over time. Low error estimates of the complex array response are needed for wavefronts arriving from every direction that a beam will be steered or where a response pattern control is needed. This requires a sequence of calibration observations of a bright reference sky source. Sidelobe responses can be specified but require a large number of calibration vectors that must be updated periodically and require high-SNR pointings at a calibration source to determine the response beyond levels at the first or second sidelobe.

Field of view coverage is sparse for an FPA, requiring multiple pointings to fill gaps between pixels. PAF elements see the full field of view, so all image pixels are available using a grid of beamformer calibrators to form narrow beams in each desired direction. System complexity is lower for an FPA than a PAF, since the multiple horn feeds are treated independently. FPAs use fewer array elements since the elements are electrically large and patterns do not overlap. Cryo-cooling is difficult for the large number of elements in a PAF because of space limitations at the reflector focus. Radiometric calibration can be done independently on each horn using calibrated noise sources at each feed in an FPA, but in a PAF the phase relationships among elements are important so this type of gain calibration is insufficient. Computational complexity per beam is the same for FPAs and PAFs, but PAFs have the added computational burden of actually forming the beams. As the number of beams formed by a PAF becomes large, the overall computational burden for a PAF becomes significant.

Despite the challenges in calibration and computational burden, the flexibility of PAFs makes them enough of an attractive new technology in astronomy to warrant continuing develop-

ment by groups around the world. This dissertation focuses on PAFs for single dish astronomy but the techniques presented here are being applied to synthesis imaging, including the interference cancelation discussed in Chapter 5. The SKA will likely involve phased arrays mounted on single dishes, and synthesis imaging arrays made up of collections of phased-array-fed reflectors. The Australian ASKAP array currently under development, with six installed dishes to date, is based entirely on a PAF design approach [15]. In May 2011 BYU hosted the first PAFSKA Workshop with attendees being major participants in the SKA PAF development community from CSIRO (Australia), DRAO (Canada), ASTRON (the Netherlands), and others.

1.2 Literature Review

The work presented in this dissertation is part of a large international effort in astronomical PAF development. This section begins with the motivation for PAF development from an astronomer's perspective and refers to international projects to develop astronomical PAFs. Advances in PAF calibration and experimental radio camera imaging results provide the state of the art in beamforming for experimental PAFs. This dissertation presents some of the earliest and highest-sensitivity experimental astronomical detections with a PAF. Finally, RFI mitigation is discussed, beginning with common sources of interference and available mitigation techniques, then focusing on subspace projection interference mitigation which is extended by this dissertation to cancel interferers in difficult environments by applying a motion model to the interference subspace.

1.2.1 Scientific Motivation for PAFs in Radio Astronomy

Many design parameters have been explored in radio telescopes including instantaneous bandwidth, system noise temperature, angular resolution, and time and frequency resolution. Two areas yet to be explored are greater collecting area and wider fields of view. PAFs provide wider fields of view by forming multiple beams without loss of sensitivity in each beam. Observations of extended HI regions with a single pixel feed are time consuming, so the following scientific studies will benefit from the wide field of view offered by a PAF: observations of HI clouds interacting with the disk of the Milky Way, surveys of gas clouds outside other galaxies covering large areas, studies of kinematics of extended neutral hydrogen (HI) clouds contained by galactic groups, and

additional studies of the galactic center. A PAF on the GBT would aid in analyzing the interstellar medium using 18 cm emission from hydroxyl (OH) to understand how to minimize the effect of foreground matter on cosmological investigations [14].

1.2.2 Previous Developments for PAFs in Radio Astronomy

Arrays in Radio Astronomy

Early advances in radio astronomy imaging were due to improved receiver electronics and larger antennas [2], [3]. The introduction of array antennas, particularly the VLA with synthesis imaging brought a huge improvement in image resolution [4]. The recent introduction of centimeter-wave focal plane arrays, notably the Parkes multibeam receiver [10], and the ALFA feed [13], and higher-frequency FPAs for the Green Bank Telescope (GBT) like the 7-pixel K-band feed [16] and the 64-pixel MUSTANG bolometer array at 90 GHz [17], give the small pixel size and high sensitivity of a large reflector while forming multiple beams on the sky.

Previous PAF Developments

NRAO developed an early 19-element array of sinuous antennas [18–20]. Since 2003, there has been an array feed collaboration between BYU and NRAO, demonstrating beamforming and RFI mitigation with a seven element hexagonal array on a 3-m reflector in 2006 [21] and more recent results on the Green Bank 20 Meter Telescope presented in this dissertation (see also [14, 22–26]). Groups with The Netherlands Foundation for Research in Astronomy (ASTRON) and in particular the Westerbork Synthesis Radio Telescope (WSRT) and Low-frequency Array (LOFAR) have demonstrated good illumination efficiency for a wideband, dual polarized Vivaldi array (APERTIF) [27–29]. The Canadian National Research Council (DRAO) has developed a Phased Array Feed Demonstrator (PHAD), also of Vivaldi elements [30]. The Australian Commonwealth Scientific and Industrial Research Organization (CSIRO) is developing a wideband connected checkerboard array for the (Australian Square Kilometre Array Pathfinder (ASKAP) synthesis imaging array) [31]. The Karoo Array Telescope (meerKAT) being developed in South Africa will be upgraded to a PAF implementation. ASKAP and meerKAT are Square Kilometre

Array (SKA) pathfinder instruments [9]. Currently all active PAF development worldwide is being carried out for L-band observations.

1.2.3 Calibration and Beamforming

A good introduction to general beamformer design is given in [32, 33], including coverage of fixed, deterministic beamformer design based on array geometry, as well as statistically optimal beamformer design. This dissertation uses two of the several well-known statistically optimal beamforming algorithms: the linearly constrained minimum variance (LCMV) and maximum sensitivity (max SNR) beamformers [32, 33]. LCMV minimizes total output power while satisfying a set of linear constraints. A good literature review for calibration and beamforming for radio astronomy is given by [34]. Willis [35] demonstrates that well-behaved PAF beam shapes can be designed in simulation using deterministic beamforming techniques with little coma, distortion, or polarization rotation as beams are steered off axis. This approach requires many beamformer calibrators that are hard to obtain experimentally with sufficient signal to noise ratio. Additionally, due to gain instabilities and difficulties in matching cable lengths it can be hard to transfer beamformers designed in simulation to an experimental system. For these reasons the PAF community has favored statistically optimal techniques in practice. The maximum sensitivity beamformer has been used for most experimental PAF observations reported to date and the only demonstrated experimental radio camera images formed to date use statistically optimal, adaptive, data-dependent beamforming algorithms [24–26, 28, 36]. Radio frequency interference (RFI) mitigation is an important goal of PAF beamformer design. Recent results demonstrating feasibility of PAF-based adaptive nulling algorithms for interference mitigation in astronomical applications include [14, 21–23]. A survey of RFI mitigation techniques is the focus of Section 1.2.5.

1.2.4 PAF Experiments

Array Characterization, Antenna/Front End Optimization, Mutual Coupling

Several technical challenges arise in astronomical PAF work because of the extreme requirements for sensitivity and antenna characterization. Noise temperatures for individual elements in isolation do not transfer directly to active beamforming arrays; instead the system temperature

is affected by element mutual coupling and must be measured using the formed beam. New definitions unifying IEEE standard definitions for single feed and array antennas, including definitions for aperture efficiency and beam sensitivity have been proposed in [37, 38] to consider the system equivalent noise temperature. Measurements of the array isotropic noise response can be used to determine the aperture efficiency and system noise temperature associated with formed beams [37].

A major concern with PAF development has been optimizing the interface between the array and front end receiver electronics [39, 40]. Due to mutual coupling between array antenna elements, standard impedance matching techniques used for single-port antennas must be extended to multi-port systems using the theory of active impedances [41–43]. Optimum source impedances of front end amplifiers must be matched to a beam-dependent active impedance. For communications applications, decoupling networks have been explored [39] but are likely too lossy and narrow band for astronomical instruments. More promising approaches include a noise matching condition that is optimal in an average sense over the array field of view [42, 44]), and the design of array antennas with active impedances that remain as close as possible to the LNA optimum source impedance over the array field of view and operating bandwidth, as demonstrated in [45].

The system described in Chapter 4 of this dissertation has provided a platform for experimentally verifying the theoretical developments in optimized antenna design to compensate for mutual coupling and reduce system noise temperature.

PAF Results and First Radio Camera Images

Efforts at BYU have focused on the best sensitivity for narrowband observations. Other groups are exploring more broadband array elements and system bandwidths. Early PAF sensitivity results and radio camera images obtained at the WSRT were reported in [36], and at NRAO/BYU in [24]. These early experiments have been proof-of-concept demonstrations to characterize embedded element beampatterns, system noise temperatures, and multi-pixel radio camera images of known sources with comparison to images obtained from other telescopes. Radio camera images created using PAFs have been reported in [14, 24–26, 28, 36]. This dissertation describes the first isotropic noise response measurements using a new type of facility to perform a Y-factor type noise measurement using the entire array. These results were first reported in [14].

1.2.5 RFI Mitigation and Subspace Projection

RFI Sources

An excellent introduction to RFI sources is given in [46], where it is noted that even locating instruments in undeveloped areas with regulatory protection does not avoid many man-made sources such as satellite downlinks. Satellites are a problem because there is nowhere in the world that a radio telescope can hide from them, and consumer electronics – both intentional radiators like cell phones and unintentional radiators like digital cameras – are a problem because of their ubiquity. Astronomers are interested in a variety of molecules, each with its own spectral characteristics and at a variety of redshifts so RFI is not limited to protected radio astronomy bands. Since astronomers avoid frequencies that have a lot of interference, it is hard to say what frequencies they might use if there were no interference there [47]. At NRAO-GB, some of the sources seen consistently are airphones in the 900 MHz band, expected to expand in the near future with broadband connectivity and 3G and 4G mobile services available on airplanes. Satellites are very strong and can get into the main beam and overload receivers. These include satellites for navigation (GPS/GLONASS), communication (Iridium out-of-band emissions), and broadcast (Sirius/XM), and often the bandwidth is more the issue than signal strength. Spread-spectrum communication systems look like broadband noise sources and are very difficult to work around. Even power line noise is strong, broadband, and recurrent [48]. Published RFI sources include: satellite downlink transmissions [49–51], radar systems [52–56], air navigation aids [57, 58], wireless communications [59, 60], and digital television broadcasts.

Various RFI Mitigation Techniques and Results

Publications related to astronomical RFI mitigation techniques dating back to 1978 can be found in [61, 62]. Other good introductions to RFI mitigation are included in [34, 63, 64]. The main approaches include blanking the corrupted data (including a Kalman tracker to track radar reflections from airplanes [52] and an FPGA implementation of radar blanking [65]), time-domain cancelation such as the LMS tracker for temporal interference cancelation described in [50], and spatial cancelation [63, 64] including auxiliary antenna-assisted spatial cancelation [66].

Various analyses have been done on the effects of interference mitigation on astronomical imaging including TDMA blanking [60], extending the CLEAN deconvolution algorithm [67], overcoming beampattern-distortion introduced PSD bias [68], and PSD bias due to cross-correlation between interference and signal plus noise [69].

Subspace Projection in Radio Astronomy

Due to its zero-forcing deep nulling behavior, the recent approach taken at BYU has focused on spatial cancelation using the subspace projection (SP) algorithm [34, 63, 64, 66, 70–72]. The first and best work applying subspace projection in an astronomical context was from Leshem et al. [63, 67, 73]. Recent results demonstrating feasibility of PAF-based interference mitigation in astronomical applications include [21–23, 68, 69].

When interference motion is relatively fast, interference to noise power ratio (INR) is low, or when very deep nulling below the noise floor is required, conventional adaptive cancelers may perform poorly. Ellingson applies the PASTd algorithm of [74] to form a subspace-tracking algorithm [64]. The work presented in Chapter 5 focuses on tracking a single interferer to produce extremely deep beampattern nulls uses a polynomial model to track a rank 1 subspace estimate.

1.3 Dissertation Outline

This dissertation presents an experimental PAF from RF front end to digital back end with results and algorithm development in the middle.

Chapter 2 presents key terminology used in this dissertation, introducing radio sources (neutral hydrogen and hydroxyl), telescope parameters (gain, sensitivity, and aperture efficiency), and statistically optimal beamformers (MaxSNR and LCMV). The signal model is given here including assumptions of stationarity for the observed signals leading to expressions for the array covariance matrix and contributions from signal, noise, and interference. The site in Green Bank, WV is described where the experimental platform of Chapter 3 was developed.

Chapter 3 gives a system description for the PAF experimental platform. The PAF consists of a 19-element array of linear dipoles mounted at the focus of the Green Bank 20 Meter Telescope, a 3.6m reference antenna provides a high-gain copy of interference for cancelation work, and the

entire system is controlled remotely over fiber optic cables from the Jansky Lab 1 km away. Major system components and the system layout are described in this chapter.

Chapter 4 presents results of two major experimental runs in 2007 and 2008. The signal model and approach to beamforming are discussed in this chapter. Bare array element radiation pattern measurements and the impact of mutual coupling are discussed. The array is further characterized using an array Y-factor method to measure the isotropic noise response. On-reflector results include total power and spectral line measurements of astronomical sources, radio camera images, and RFI mitigation demonstrations.

Chapter 5 presents algorithm developments augmenting subspace projection interference mitigation with a polynomial covariance model to enable deep interference nulling in the particularly difficult cases of low INR and rapid interferer motion. Performance limiting effects are investigated including sample estimation error for low INR, subspace smearing for interferer motion, noise bias for correlated noise due to mutual coupling and reflector obscuration, and spectral bias (scooping) due to a coupling between the temporal and spatial cancellation. A polynomial model is introduced for covariance tracking giving the dual benefit of reducing subspace smearing and sample estimation error reduction through increased data averaging. Cramer-Rao Bound analysis is performed, and experimental data verification is shown using the array and receivers described in Chapter 3 in an anechoic chamber.

Chapter 6 highlights progress on an FPGA-based digital back end system to transition the system from a narrow 425 kHz RF bandwidth prototype system to a science-ready instrument. This chapter describes international progress with CASPER hardware toward real-time correlator development and includes recommendations made to the Arecibo Observatory on future digital back end work for the AO40, a large scale PAF with a 40-beam real-time correlator/beamformer. The chapter concludes with a road map for future hardware beamformer development here at BYU.

1.4 Contributions and Publications

This section presents a list of the main contributions of this dissertation, highlighting three main areas: 1) experimental PAF results, 2) deep interference nulling, and 3) digital back end development, followed by a list of publications that have resulted from the work.

1.4.1 Contributions

- Design of a phased array feed system and experiments to detect astronomical sources and characterize the array.
 - Large group project involving over 30 people to build one of the first astronomical PAFs in the world and with highest sensitivity of its time. Major contributors include:
 - * BYU: Dr. Brian Jeffs, Dr. Karl Warnick, Jonathan Landon, David Jones, Jacob Waldron, Jiyoung Son, Alan Stemmons, Michael Elmer, Micah Lillrose, Michael Mendoza (prior work: Chad Hansen, James Nagel, Chris Ashworth)
 - * NRAO: Roger Norrod, Dr. Rick Fisher, Bob Simon, Ken Ward, Dave Woody, Bill Shank, Mike Hedrick, Rob Taggart, Pete Chestnut, Bob Anderson, Gary Anderson, Jonah Bauserman, Carla Beaudet, Greg Monk, Charlie Myers, Chris Clark, Wolfgang Baudler, Nathan Sharp, Rusty Taylor, Daryl Shinaberry
 - * This section includes significant results from the group effort, such as:
 - The highest sensitivity and lowest noise temperature of their time, with system stability on the order of days, permitting calibration vector re-use.
 - Array Y-factor noise measurement technique for characterizing astronomical PAFs that is becoming a standard in the community.
 - Radio camera images to show improved sky survey speed with multiple simultaneously formed beams.
 - RFI mitigation in experimental data.
 - My contributions include:
 - * Major leadership role coordinating teams at BYU and NRAO, training, delegating, overseeing experiment logistics/manpower/duration/travel through 5 experimental runs and nearly 5 months on site.
 - * Primary design and layout of the system spread over 2km connected via fiber optic network for central control. I oversaw all aspects of design, and no one knew more about the system than I did.

- * Responsibility for design and construction of the isotropic noise measurement facility and experiments. Following our lead other groups internationally built similar facilities.
 - * Conducting RFI mitigation experiments and demonstrating cancelation in real data.
 - * Running antenna range tests to characterize mutual coupling.
- Improved RFI mitigation is achieved in the challenging cases of low interference-to-noise ratio (INR) and moving interference by combining subspace projection (SP) beamforming with a polynomial model to track a rank 1 subspace.
 - Factors limiting covariance estimation in SP on a per-STI basis are explored, including sample estimation error, subspace smearing, and correlated noise bias.
 - Overall null depth improvement of between 6 and 30dB without requiring prior knowledge of optimal STI length.
 - Phase normalization that better preserves the continuous variations than simply normalizing to zero the phase on the first array element.
 - Expressions for the gradient, Hessian, and Jacobian of the algorithm’s objective function for use in numerical optimization.
 - Cramer-Rao Bound analysis on polynomial coefficient estimation.
 - Real data experiments provide experimental verification of the principles described in this dissertation.
 - Created preliminary framework for increasing the PAF system bandwidth from 425 kHz to hundreds of MHz bandwidth required by astronomers for fully science ready system.
 - Inaugurated BYU contribution to international CASPER development during 6 months at the University of California, Berkeley.
 - Began FPGA-based digital back end development using a 64-input, 50 Msamp/sec ADC and ROACH processing board. Working with another student at BYU, created a

real-time spectrometer using building blocks developed at Oxford as part of the larger CASPER collaboration.

- Made recommendations on feasibility and timelines for Arecibo Observatory and future BYU proposals, synthesizing information gathered from CASPER contacts.
- Created a road map toward a 64-input 20 MHz F/X/B correlator/beamformer laying out incremental steps to a full hardware back end.
- Gathered and synthesized information on progress made by other groups building correlators with CASPER hardware and alternative approaches to digital back end development for large scale PAFs.

1.4.2 Publications

Several publications have resulted from this work, including:

- Journal Articles

- **J. Landon**, B.D. Jeffs, K.F. Warnick, “Model-Based Subspace Projection Beamforming for Deep Interference Nulling”, submitted, IEEE Transactions on Signal Processing
- **J. Landon**, M. Elmer, J. Waldron, D. Jones, A. Stemmons, B.D. Jeffs, K.F. Warnick, J.R. Fisher, and R.D. Norrod, “Array Feed Calibration, Beamforming and Imaging”, The Astronomical Journal 139 (2010) 1154–1167.
- 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, p. 635–646
- 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.

- Full Manuscript Reviewed Conference Proceedings

- D.A. Jones, K.F. Warnick, B.D. Jeffs, **J. Landon**, J. Waldron, J.R. Fisher, R. Norrod, “Modeled and measured mutual impedances, element patterns, and sensitivity for a 19

element focal plane array”, Antennas and Propagation Society International Symposium, AP-S 2008. IEEE, 5-11 July 2008 Page(s):1–4.

- Abstract Reviewed and Other Conference Articles, Presentations, and Reports

- K. F. Warnick, D. Carter, T. Webb, **J. Landon**, V. Asthana, M. Elmer, R. Norrod, and J. R. Fisher, “Towards a high sensitivity, cryogenic phased array feed antenna for the Green Bank Telescope”, in URSI General Assembly and Scientific Symposium of International Union of Radio Science, Istanbul, Turkey, Aug 2011.
- **J. Landon**, B. D. Jeffs, K. F. Warnick, D. Carter, T. Webb, M. Elmer, J. R. Fisher, and R. Norrod, “An interference canceling phased array feed for the Green Bank Telescope”, URSI National Radio Science Meeting, Boulder, CO, Jan. 5–8, 2011.
- B. Jeffs, K. Warnick, **J. Landon**, M. Elmer, J. Fisher, and R. Norrod, “Recent progress with the BYU/NRAO phased array feed”, in Calibration and Imaging Workshop, CALIM2009, Socorro, New Mexico, Mar. 2009, https://safe.nrao.edu/wiki/pub/Software/CalIm09Program/brian_jeffs.pdf.
- K. Warnick, B. Jeffs, **J. Landon**, J. Waldron, D. Jones, J. Fisher, and R. Norrod, “Phased array antenna design and characterization for next-generation radio telescopes”, in Antenna Technology, 2009. iWAT 2009. IEEE International Workshop on, Mar. 2009, pp. 1–4.
- K. F. Warnick, B. D. Jeffs, **J. Landon**, J. Waldron, D. A. Jones, R. Fisher, and R. Norrod, “Beamforming and imaging with the BYU/NRAO L-band phased array feed”, Proceedings of the 13th International Symposium on Antenna Technology and Applied Electromagnetics and the Canadian Radio Sciences Meeting (ANTEM/URSI), Banff, AB, Canada, Feb. 15–18, 2009. Invited.
- **J. C. Landon**, B. D. Jeffs, K. F. Warnick, J. R. Fisher, and R. D. Norrod, “Improved Low INR Interference Cancellation Using Phased Array Feeds”, URSI General Assembly, Chicago, IL, Aug. 9–16, 2008.

- M. J. Elmer, **J. Landon**, B. D. Jeffs, K. F. Warnick, J. Waldron, J. R. Fisher, and R. D. Norrod, “Optimizing Phased Array Feeds to Reduce Elevation-Dependent Spillover Noise in Radio Telescopes”, URSI General Assembly, Chicago, IL, Aug. 9–16, 2008.
- K. F. Warnick, B. D. Jeffs, **J. Landon**, J. Waldron, D. Jones, R. Fisher, and R. Norrod, “BYU/NRAO 19-Element Phased Array Feed Modeling and Experimental Results”, URSI General Assembly, Chicago, IL, Aug. 9–16, 2008.
- 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”, Calibration and Imaging Workshop, CALIM2008, Perth Australia, 7–9 Apr. 2008.
- **J. Landon**, K. F. Warnick, B.D. Jeffs, D. Jones, R. Norrod, R. Fisher, “A 19-Element L-band Focal Plane Array for Interference Mitigation with Auxiliary Reference Antenna on the Green Bank 20 meter Telescope”, AAS 211th Meeting, Austin, TX, Jan. 7–11, 2008.
- **J. Landon**, D. Jones, B. D. Jeffs, K. F. Warnick, R. Fisher, and R. Norrod, “Interference cancellation and sensitivity optimization using an L-band focal plane array on the Green Bank 20m Telescope”, URSI National Radio Science Meeting, Boulder, CO, Jan. 3–6, 2008.
- K. F. Warnick, B. D. Jeffs, **J. Landon**, J. Waldron, D. Jones, A. Stemmons, “BYU/NRAO 2007 Green Bank 20 Meter Telescope focal plane array–modeling and experimental results”, invited presentation, SKADS MCCT Technical Workshop on “Design of Wide-band Receiving Array Systems”, Dwingeloo, The Netherlands, Nov. 26–30, 2007.
- K. F. Warnick, J. Waldron, **J. Landon**, M. Lilrose, and B. D. Jeffs, “Experimental results on interference mitigation with a 19 element array feed”, Proceedings of 2nd European Conference on Antennas and Propagation, Edinburgh, UK, Nov. 11–16, 2007.

Chapter 2

Background

This chapter presents key terminology used in this dissertation, both from radio astronomy and signal processing. The material presented here is limited to the radio sources, telescope parameters, and beamforming concepts used in this dissertation. Thorough introductions to astronomical sources and detection techniques for radio astronomy can be found in [2, 75], and to array signal processing and beamforming in [33, 76, 77]. The chapter proceeds as follows: Section 2.1 presents the notation used throughout the dissertation. Section 2.2 describes the two radio sources (HI and OH) observed experimentally in Chapter 4 and the telescope parameters (gain, sensitivity, aperture efficiency) characterized in that chapter. Section 2.3.1 presents the signal model and assumptions of stationarity for the observed signals leading to expressions for the array covariance matrix and contributions from signal, noise, and interference. Section 2.3.2 presents the statistically optimal beamformers (MaxSNR and MVDR) used in Chapters 4 and 5 for defining the telescope beam pattern. The zero-forcing subspace projection beamformer used for RFI mitigation demonstrations in Chapter 4 and extended with a polynomial model in Chapter 5 is presented in Section 2.3.3. The site in Green Bank, WV where the experimental platform of Chapter 3 was developed is described in Section 2.4.

2.1 Notation

1. $E\{\mathbf{A}\}$: Expected value of random \mathbf{A} .
2. $\mathbf{A}^T, \mathbf{A}^H$: Transpose and conjugate transpose of \mathbf{A} .
3. $\hat{\mathbf{A}}$: Sample estimate of some parameter, \mathbf{A} .
4. $\check{\mathbf{A}}$: Calibration data set estimate of \mathbf{A} .
5. $\tilde{\mathbf{A}}$: Polynomial model estimate of some parameter, \mathbf{A} .

6. $\{\mathbf{A}\}_{i,j}$: The i, j th element of \mathbf{A} .
7. $\|\mathbf{a}\|, \|\mathbf{A}\|_F$: l_2 vector norm and Frobenius matrix norm.
8. $\text{Re}\{\alpha\}, \text{Im}\{\alpha\}$: Real and imaginary parts of complex α .
9. $\angle\alpha, \bar{\alpha}$: Phase angle and complex conjugate of α .
10. $\mathbf{A} \circ \mathbf{B}$: Khatri-Rao, or column-wise Kronecker product.
11. $\mathbf{1}_P, \mathbf{I}_P$: $P \times 1$ vector of ones and $P \times P$ identity matrix.

2.2 Astronomy Background

Throughout this dissertation, various units and system characterizations are used that are unique to radio astronomy. This section clarifies the use of some of these terms that may be unfamiliar to engineers outside of astronomy.

2.2.1 Fundamentals

Brightness Temperature – Kelvin (K)

The most basic source of radio emission is due to blackbody radiation from a source at a particular temperature T in Kelvin (K). A perfect blackbody is an idealized material that absorbs all incident energy at all frequencies and reflects none; it is also a perfect emitter. According to Planck's radiation law, a blackbody radiates uniformly in all directions with a spectral brightness B_f that depends only on the temperature and the observation frequency:

$$B_f = \frac{2hf^3}{c^2} \left(\frac{1}{\exp(hf/k_B T) - 1} \right),$$

where B_f is the spectral brightness in $\text{W}/\text{m}^2/\text{sr}/\text{Hz}$, h is Planck's constant (6.63×10^{-34} J), f is frequency in Hz, k_B is Boltzmann's constant (1.38×10^{-23} J/K or W/K/Hz), T is the absolute temperature of the source in Kelvin, and c is the speed of light (3×10^8 m/s) [2, 78]. Radiation observed at a particular brightness corresponds to a source at a particular absolute temperature, so

radio sources are often described in terms of brightness temperature, the absolute temperature of a source that would give rise to the observed brightness.

The brightness temperature makes a nice analogy to the power available from a resistor at a particular temperature connected to a matched load. Nyquist showed that for such a resistor at temperature T in Kelvin, the available power is

$$P_a = k_B T B,$$

where P_a is the available power in Watts and B is the observation bandwidth in Hz [78]. The noise power added by a receiver can thus be converted into an equivalent noise temperature, T_{eq} or T_{sys} – the absolute temperature of an equivalent resistor placed at the input to a noiseless receiver. This conversion allows the added system noise temperature and the observed astronomical source power to be interpreted as sources external to the telescope. Then the total received power, converted to equivalent temperature in Kelvin is

$$T = T_A + T_{\text{sys}}, \tag{2.1}$$

where T is the total received power, T_A is the antenna temperature due to the astronomical source, and T_{sys} is the temperature of an equivalent source representing the noise added by the system.

Flux Density – Jansky (Jy)

The Jansky (Jy), named in honor of Karl Jansky, the father of radio astronomy, is a measure of the incident power arriving at a radio telescope:

$$1 \text{ Jy} = 10^{-26} \text{ W/m}^2/\text{Hz}.$$

The brightest radio sources are on the order of a few thousand Jy, and the weakest detectable sources are on the order of micro-Janskys (μJy) for a highly sensitive telescope. The concept of flux density can be understood as a sheet of energy arriving at the telescope with power distributed evenly per unit area so that a larger aperture telescope collects more total power than a smaller aperture telescope but the incident energy density is unchanged. The term “flux” connotes the flow of power radially outward from the source or coming into the telescope. The energy is also

distributed in frequency so that a wider bandwidth receiver captures more incident power than a narrower bandwidth receiver, thus reducing time required for successful source detection.

The following example illustrates how to make the Jy more familiar to engineers accustomed to dealing with spectrum analyzers whose units are in dBm/Hz. The noise added by the receiver can be represented as a resistor at the same ambient temperature as the equivalent noise temperature of the receiver. The noise power supplied by a resistor at room temperature (290 K) to a matched load has a power spectral density of -174dBm/Hz, using $P = k_B T B$, where P is power in Watts, k_B is Boltzmann's constant (1.38×10^{-23} W/K/Hz), T is room temperature (290 K), and B is the bandwidth in Hz of the receiving system. An ideal 20m diameter radio telescope has collecting area of $100\pi\text{m}^2$, so a 1 Jy source would produce $100\pi \times 10^{-26}/1\text{mW} \approx -205\text{dBm/Hz}$, giving an instantaneous SNR of -31dB for a room temperature receiver. In contrast, communication systems often require SNRs on the order of a few dB or tens of dB for minimal errors.

2.2.2 Spectral Line Sources

Spectral line sources are an important class of astronomical sources in addition to the black-body radiators described earlier. A molecule changing energy states can emit or absorb electromagnetic radiation at a particular frequency. The two sources discussed in this dissertation are neutral hydrogen (HI) and hydroxyl (OH).

Neutral Hydrogen (HI)

Neutral hydrogen, abbreviated HI or H I, has a spectral line associated with the transition between two closely spaced energy states, called a hyperfine transition, and is one of the most important spectral lines observed in radio astronomy. The radiation measured in the laboratory at 1420.405751786(30) MHz is one of the most precisely measured physical quantities [75] and is commonly referred to by its wavelength as the 21cm line. The two closely spaced energy levels are associated with electron spin in the neutral hydrogen atom, giving rise to a magnetic dipole moment. When the spin of the electron leads to a magnetic dipole moment that is aligned with the dipole moment of the nucleus, the energy is slightly higher than when the dipole moments are opposed. The probability of spontaneous transition between states is such that an atom remains in

the higher state for 11 million years, but collisions reduce that expected time to 400 years [2, 75]. These probabilities are small, but considering the vast quantity of HI available for study, these transitions can be observed constantly. Doppler shifts to the HI line due to motion toward or away from the Earth has been an important part of studying the evolution of the universe and identifying the spiral structure of our galaxy.

Hydroxyl (OH)

There are four main spectral lines associated with the hydroxyl (OH) molecule: 1612.231 MHz, 1665.400 MHz, 1667.358 MHz, and 1720.529 MHz. The OH line involves an electric dipole moment which is much more intense than the magnetic dipole transition of the hydrogen line. Typical electric dipole moments are a factor of 46 times larger than that of an equivalent magnetic dipole moment [75]. The transition probability of OH is also much higher than that of HI, so it can be detected despite its much smaller distribution (smaller by a factor of 20 million over the path to CasA) [2].

If the population of higher energy levels in OH exceeds that of the lower energy levels, a radio wave passing through OH molecules may induce more transitions from high to low than from low to high. Since the transition from high to low is associated with emission and from low to high with absorption, this means more energy will be released than is absorbed as the wave passes through the medium, leading to amplification of the radio wave. This behavior is known as microwave amplification by stimulated emission of radiation, called a MASER. It relies on having background material providing continuum radiation and large optical depth of the OH medium with the population inversion (higher energy states more highly populated) [2, 75].

2.2.3 Telescope Parameters

This section gives basic definitions and intuition on radio telescope figures of merit. Many additional characterizations exist that give insight into the design and performance of a radio telescope but this section is limited to those appearing in Chapter 4.

Gain

The general definition of gain for an arbitrary antenna is the directivity when reduced by the radiation efficiency. Directivity is a measure of the directionality as compared to an isotropic radiator. The ratio of power radiated in the peak direction relative to the power radiated by an isotropic radiator in that direction leads to the expression:

$$D = A_{\text{eff}} \frac{4\pi}{\lambda^2},$$

where D is the directivity (dimensionless), A_{eff} is the effective collecting area of the aperture in m^2 , and λ is the wavelength of the radiated energy [78]. Multiplying D by the radiation efficiency gives the gain, G .

For a radio telescope an additional definition of gain is useful, defined as the ability of the telescope to convert incident flux density in Jy to antenna temperature in Kelvin. Using this definition, gain is found using:

$$G = \frac{T_A}{F^{\text{s}}} \quad (\text{K/Jy}), \quad (2.2)$$

where G is the gain with units of K/Jy , T_A is the antenna temperature in Kelvin, and F^{s} is the source flux density in Jy. This is a measured quantity that takes into account all of the losses due to a non-ideal antenna aperture and receiving system.

Sensitivity

Sensitivity can be expressed as:

$$\frac{A_{\text{eff}}}{T_{\text{sys}}} = \frac{k_B B}{S^{\text{inc}}} \text{SNR} \quad (\text{m}^2/\text{K}), \quad (2.3)$$

where A_{eff} is the effective antenna collecting area in m^2 , T_{sys} is the system noise temperature in Kelvin, and S^{inc} is the incident signal power density in one polarization in W/m^2 . Incident signal power S^{inc} is found using $1/2 \times 10^{-26} \times B \times F^{\text{s}}$, where the $1/2$ comes from assuming the source power is spread equally across both polarizations but these dipoles only receive one polarization, and $10^{-26} \times B$ converts from Jy ($10^{-26} \text{W}/\text{m}^2/\text{Hz}$) to power density in W/m^2 .

This definition of sensitivity gives a measure of telescope performance that depends on the quality of the receiver and the telescope illumination. A telescope with large physical diameter that is poorly illuminated by the feed has reduced effective collecting area, A_{eff} . A feed that over-illuminates the reflector gives increased A_{eff} but adds spillover noise as it sees beyond the edges of the reflector. This added spillover noise increases T_{sys} , reducing sensitivity. A lower noise receiver placed on a smaller telescope may perform better than a noisier receiver on a larger telescope, an idea captured by this definition of sensitivity.

Aperture Efficiency

An ideal aperture antenna uses its entire physical area to collect a section of the incident power. In reality, the aperture is usually nonuniformly illuminated, so the effective area measures the actual power-collecting area of the antenna. Aperture efficiency is the ratio of the effective collecting area to the physical collecting area:

$$\eta_{\text{ap}} = \frac{A_{\text{eff}}}{A_{\text{phys}}}, \quad (2.4)$$

where A_{phys} is the physical area of the antenna ($100\pi\text{m}^2$, in the case of the Green Bank 20 Meter Telescope), and A_{eff} is the effective area of the telescope, which can be estimated using the sensitivity expressed in m^2/K from $A_{\text{eff}}/T_{\text{sys}}$:

$$A_{\text{eff}} = T_{\text{sys}} \frac{A_{\text{eff}}}{T_{\text{sys}}}. \quad (2.5)$$

Due to reciprocity, the receiving pattern of the antenna is the same as the radiation pattern of the antenna as a transmitter. Thinking of the antenna as a transmitter, the feed cannot provide a step-function illumination of the reflector where all transmitted energy is distributed evenly across the reflector and none spills over the reflector edges onto the ground. As a receiver, this means the feed is seeing the ground as a bright thermal radiator beyond the edges of the reflector in close proximity competing with the astronomical signals coming from much farther away. To reduce this ground spillover, the feed illuminates a smaller portion of the reflector. The downside to under-illuminating the reflector is that it means the telescope is effectively much smaller. As a rule of

thumb, typical radio telescope feeds taper the illumination pattern by 10 dB at the edges of the reflector to trade off aperture efficiency and spillover efficiency.

2.3 Beamforming Background

This section presents the signal model and beamformers used throughout the dissertation. The signals are all assumed to be wide-sense stationary processes, and the array covariance is computed only at a single frequency because the system is narrowband. Equations are given for the statistically optimal Maximum SNR and LCMV/MVDR beamformers as well as the zero-forcing Subspace Projection beamformer that is used in interference mitigation tests.

2.3.1 Signal Model

The system bandwidth is assumed to be sufficiently narrow so that $B \ll D/c$ where B is the bandwidth, D is the PAF diameter and c is the speed of light. For wider band systems the beamforming can be done in a similar way where the total bandwidth is first broken into frequency subbands where the above condition holds. The array output is a $P \times 1$ vector of complex basebanded voltages

$$\mathbf{x}[n] = \mathbf{a}_s s[n] + \boldsymbol{\eta}[n],$$

where \mathbf{a}_s is the array response due to the signal of interest, $s[n]$, and $\boldsymbol{\eta}[n]$ is a vector of noise and possibly interference. Assuming zero-mean, wide-sense stationary signals (except for moving interference) and statistical independence between signal and noise sources, the array covariance and correlation matrices can be used interchangeably and are written as

$$\begin{aligned} \mathbf{R} &= E\{\mathbf{x}[n]\mathbf{x}^H[n]\} = \mathbf{R}_s + \mathbf{R}_\eta \\ \mathbf{R}_s &= \sigma_s^2 \mathbf{a}_s \mathbf{a}_s^H \\ \mathbf{R}_\eta &= \mathbf{R}_{\text{rec}} + \mathbf{R}_{\text{sp}} + \mathbf{R}_{\text{sky}} + \mathbf{R}_{\text{loss}} + \mathbf{R}_{\text{int}}[n], \end{aligned}$$

where σ_s^2 is the signal power, \mathbf{R}_{rec} , \mathbf{R}_{sp} , \mathbf{R}_{sky} , \mathbf{R}_{loss} , and $\mathbf{R}_{\text{int}}[n]$ are the covariance contributions due to the receiver, spillover region, sky [non-signal sources including atmosphere, cosmic microwave background (CMB) and galactic background (GB)] antenna resistive losses, and interfer-

ers, respectively. Due to the stationarity assumption, all covariances except $\mathbf{R}_{\text{int}}[n]$ do not depend on n , but because of moving interference \mathbf{R}_η can vary slowly with n . Unless specific mention is made of interference cancelation, $\mathbf{R}_{\text{int}}[n]$ is assumed zero.

\mathbf{R} is estimated for the j^{th} short-term integration (STI) window using the sample autocorrelation matrix:

$$\hat{\mathbf{R}}^{(j)} = \frac{1}{L} \sum_{n=jL}^{(j+1)L-1} \mathbf{x}[n]\mathbf{x}^H[n] = \frac{1}{L} \mathbf{X}_j \mathbf{X}_j^H \quad (2.6)$$

$$\mathbf{X}_j = [\mathbf{x}[jL], \mathbf{x}[jL+1], \dots, \mathbf{x}[(j+1)L-1]],$$

where \mathbf{X}_j is the j^{th} STI data window of length L samples. Estimating \mathbf{R} requires that L is chosen such that wide-sense stationarity holds over the data window. In the presence of moving interference L must be short enough to provide approximate stationarity, but in interference-free cases L should be as long as possible and the superscript (j) will be dropped where this does not introduce confusion.

2.3.2 Adaptive PAF Beamforming

The astronomical PAF community has favored adaptive beamformers over deterministic beamformers in experimental results reported to date. Statistically optimal adaptive beamformers are based on the array covariance matrix which can be estimated from sample data prior to an observation, accounting for the current state of the array and noise environment. Deterministic beamformers require knowledge of the dish Airy pattern including feed support scattering and blockage as well as antenna and receiver gains and phases which are difficult to obtain with a high degree of confidence due to their time-varying, temperature-dependent nature. This section describes the two adaptive beamformers used in this dissertation: the maximum signal-to-noise ratio (MaxSNR) beamformer [32, 77] and the linearly-constrained minimum variance (LCMV) beamformer [33, 79]. For the case of a single constraint in LCMV (all uses of LCMV in this dissertation use only a single constraint), the beamformer becomes the minimum-variance distortionless response (MVDR) beamformer [33, 76, 80].

Maximum Signal-to-Noise Ratio (MaxSNR): The SNR, $\mathbf{w}^H \mathbf{R}_s \mathbf{w} / \mathbf{w}^H \mathbf{R}_\eta \mathbf{w}$ is maximized using [77]

$$\mathbf{w}_{\text{MaxSNR}} = \arg \max_{\mathbf{w}} \frac{\mathbf{w}^H \mathbf{R}_s \mathbf{w}}{\mathbf{w}^H \mathbf{R}_\eta \mathbf{w}} \quad (2.7)$$

that leads to the generalized eigenvalue problem

$$\mathbf{R}_s \mathbf{w}_{\text{MaxSNR}} = \lambda_{\max} \mathbf{R}_\eta \mathbf{w}_{\text{MaxSNR}}. \quad (2.8)$$

This is an example of a fixed-adaptive beamformer since it adapts to the signal and noise environments at calibration but remains fixed throughout the observation. To make this an active-adaptive beamformer, \mathbf{R}_η would be updated periodically with calibrations using off-pointings $\hat{\mathbf{R}}_{\text{off}}$ in place of \mathbf{R}_η .

Linearly Constrained Minimum Variance (LCMV): The total beamformer output power is minimized subject to a linear constraint [79]:

$$\arg \min_{\mathbf{w}} \mathbf{w}^H \mathbf{R} \mathbf{w} \quad \text{subject to} \quad \mathbf{C}^H \mathbf{w} = \mathbf{f}, \quad (2.9)$$

where \mathbf{C} is a list of steering vectors and \mathbf{f} is a list of constraints. Minimizing (2.9) yields:

$$\mathbf{w}_{i,\text{LCMV}} = \mathbf{R}^{-1} \mathbf{C} [\mathbf{C}^H \mathbf{R}^{-1} \mathbf{C}]^{-1} \mathbf{f}.$$

For the case of the single constraint, LCMV is also called the Minimum Variance Distortionless Response (MVDR) beamformer or Capon beamforming [33, 76, 80]. Specifying unity gain in the desired direction, $\mathbf{C} = \mathbf{a}_i$ and $\mathbf{f} = 1$, so the minimization becomes

$$\arg \min_{\mathbf{w}} \mathbf{w}^H \mathbf{R} \mathbf{w} \quad \text{subject to} \quad \mathbf{a}_i^H \mathbf{w} = 1$$

leading to

$$\mathbf{w}_{i,\text{MVDR}} = \mathbf{R}^{-1} \mathbf{a}_i / (\mathbf{a}_i^H \mathbf{R}^{-1} \mathbf{a}_i) = \alpha \mathbf{R}^{-1} \mathbf{a}_i, \quad (2.10)$$

where α is a scalar that does not affect the beamformer performance. Essentially LCMV/MVDR keeps the signal response in the desired direction constant while minimizing the response to all

other sources contributing to \mathbf{R} . The form of (2.10) is very similar to the MaxSNR beamformer, as observed in [81] and [82]. Both MaxSNR and LCMV/MVDR have a form that includes a variation of \mathbf{R}_n^{-1} , which illustrates that their optimality constraint relies on whitening the noise and then forming a spatial matched filter in this whitened space to the SOI. For a rank 1 source, the MaxSNR and LCMV beamformers are analytically identical.

2.3.3 Active Interference-Canceling Beamformers

The noise power inversion property of MVDR mentioned in the previous section is good at reducing the overall noise contribution but if deep nulls in a particular direction are required a different approach is needed. A strong interferer coming in through the sidelobes will be canceled by MVDR until its total power contribution is on the same order as the residual noise floor, but the optimality constraint will then focus on further driving down noise power to reduce the total power coming from the sidelobes. Residual interferers at the level of the noise floor will still overwhelm astronomical signals of interest, since these are normally well below the noise floor. Ideally the interferer would be completely canceled with a true zero-response beam pattern null in the direction of the interferer. Subspace Projection (SP), described in Section 4.7.2 and Chapter 5, projects the total received signal onto a space orthogonal to the interferer attempting to null it out completely. The SP weights are given by:

$$\mathbf{P}_{\text{int}}^{\perp} = \mathbf{I} - \mathbf{U}_{\text{int}} \mathbf{U}_{\text{int}}^H \quad (2.11)$$

$$\mathbf{w}_{\text{SP}} = \mathbf{P}_{\text{int}}^{\perp} \mathbf{w}, \quad (2.12)$$

where \mathbf{U}_{int} is a matrix formed by partitioning the SVD of \mathbf{R} such that its columns span the interference subspace, and \mathbf{w} is the beamformer that would have been used in the absence of interference, designed using MaxSNR or MVDR.

2.4 Location: NRAO – Green Bank, WV

Figure 2.1 shows the location of the various telescopes at the Green Bank observatory. The observatory is home to the 100-meter diameter Robert C. Byrd Green Bank Telescope (GBT), one

of the largest movable structures in the world. Green Bank was chosen for its remote location to reduce radio frequency interference (RFI) because it sits in a rural valley surrounded by mountains to further reduce interference. Because the location is within the National Radio Quiet Zone (NRQZ), the observatory is unique in its ability to request that permanent radio transmitters are directed away from the observatory. Special care is taken at the observatory to reduce RFI, including only permitting vehicles with diesel engines in Zone 1 (the area near the telescopes) since they lack spark plugs that cause RFI, and locating the control rooms in a Faraday cage at the Jansky Lab which is itself located roughly 0.5km from the nearest telescope.

Telescopes are controlled from the Jansky Lab using fiber optic communication, and the received signals are in most cases modulated onto fiber and sent back to the Jansky Lab for processing. The BYU PAF system is unusual in that the signals are recorded at the telescope in Zone 1 so the digital back end had to be inside a shielded rack to reduce RFI produced by the digital electronics.




Figure 2.1: Overhead view of the site. The Green Bank 20 Meter Telescope is about 1km from the Jansky Lab and the 3.6m antenna is about 0.5km from the telescope and the lab. Image courtesy of NRAO/AUI/NSF.

Chapter 3

Development of an Experimental Phased Array Feed for Radio Astronomy

3.1 Introduction

This chapter describes an experimental PAF system developed for the Green Bank 20 Meter Telescope. It represents the first major PAF installed on a telescope at NRAO-GB and at the time was one of only two such experimental systems in the world. The foundation of this work was laid with preliminary tests using an array of 19 sinuous crossed dipoles on the Green Bank 43 Meter Telescope [18–20]. Designing and installing this PAF system has been a joint learning experience at BYU and NRAO, requiring creativity to meet the system engineering challenges because there was little institutional knowledge about designing a large scale PAF system like this one. Other observatories around the world are developing PAF systems as well, notably in the Netherlands at ASTRON’s WSRT [28] and Australia’s CSIRO [31] (see Section 1.2.2).

PAFs are part of the worldwide SKA development plans but there are still many challenges to be solved involving calibration stability, polarization purity, dish illumination, radiation pattern shape, and mutual coupling effects. This is a new application for adaptive array systems, requiring lower noise and higher sensitivity than other applications, motivating further study in beamformer design. Characterizing the noise response and sensitivity of the array are challenging because the array is part of a large telescope and because the noise temperature is low enough that it is difficult to measure. The development of this experimental system is an important part of answering these theoretical questions.

One early design decision was to develop a narrowband (425 kHz) system focusing on beamformer design and highest possible sensitivity, rather than a full broadband science-ready instrument. In this way the BYU PAF development was able to make contributions with a much smaller budget and development team than those available at international observatories who started with narrowband systems but set a higher priority on eventual broadband designs. Hav-

ing a narrowband system simplified the digital back end because voltages could be sampled at a relatively low rate and streamed to disk for post-processing without any prior frequency channelization or other real-time processing. Beamformers could be designed based on the recorded data and reprocessed as necessary. Analog beamforming elements are not required, simplifying the front end electronics. Frequency dependence in cables and components throughout the system could be ignored, simplifying system calibration. Eventually the system bandwidth must be expanded to several hundred MHz as required by the scientific needs of astronomers. Chapter 6 describes ongoing work to create a digital back end capable of handling tens and eventually hundreds of MHz of RF system bandwidth.

This experimental system has served an important role in verifying the theoretical models for sensitivity and multiple beam formation proposed by the astronomical PAF community. Chapter 4 discusses the need for in situ calibrations for computing beamformer coefficients, and presents some of the first astronomical observations made using a PAF. Many of the simulation codes developed at BYU for analyzing PAF behavior assume a PAF mounted on a radio telescope with the particular physical characteristics of this system.

One aspect of this system that is unique among astronomical PAFs currently in development internationally is its mobility. Lacking a permanent installation at an observatory, the system was designed to be installed and removed with minimal effort. The mobility of the system led to the system later being installed on the Arecibo Telescope for focal plane mapping experiments in Puerto Rico. It also allowed the system to be installed at the anechoic chamber at the University of Utah for interference cancelation tests discussed in Chapter 5 of this dissertation.

Section 3.2 presents an overview of the system in terms of four main subsystems. The rest of this chapter is organized according to those subsystems, describing constituent components in greater detail. Chapter 4 describes experiments to observe astronomical sources and characterize the system performance.

3.2 System Layout

The system consists of 4 main functional blocks: 1) RF front end in the front end box (FEB), 2) digital back end in the telescope pedestal, 3) antenna position control and data acquisition control in the Jansky Lab, 4) reference antenna for RFI tests. The entire system was designed

to provide remote operation from the Jansky Lab control room including pointing for the telescope and reference antenna, RF observation frequency, and digital recording and processing. Optical fiber connects all components to the control room 1 km away. Figure 3.1 shows the connections among each of these subsystems. The initial concept of the system was envisioned with an earlier version of this figure and later versions such as this one have shown up in many conference presentations and proposals over the course of the design, construction, and deployment of the system.

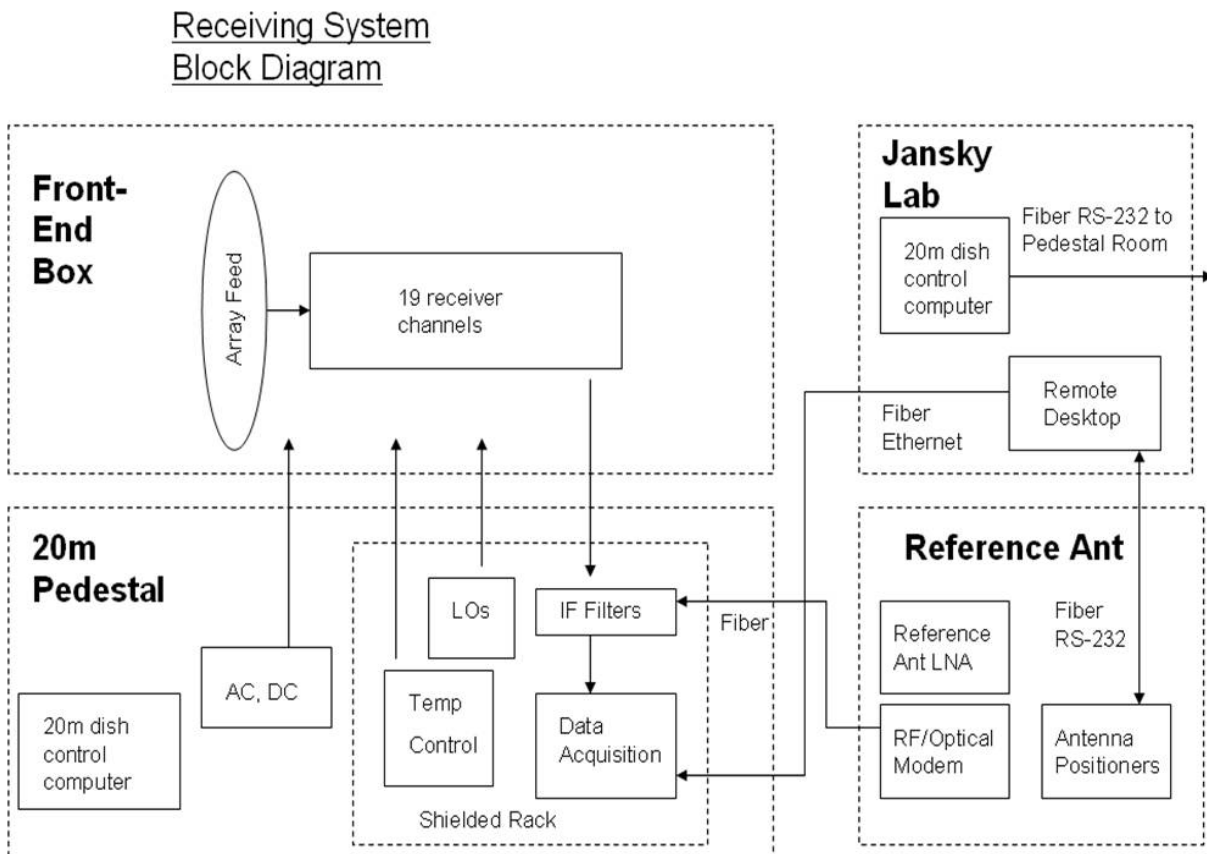


Figure 3.1: Receiving system block diagram, showing communication among the various subsystems. The system was controlled from the Jansky Lab by logging into the computer in the telescope pedestal to control LO frequencies and begin data acquisition.

The RF front end consists of a 19-element array of L-band dipoles at the focus of the telescope, mounted on a standard front end box (FEB) that houses the low-noise amplifiers (LNAs) and frequency downconversion system. Receivers in the FEB perform a two-stage downconversion to a

final IF of 2.8125 MHz and cables running down the feed support arms into the telescope pedestal carry the IF signals to the digital back end. Telescope pointing was controlled by a dedicated computer in the Jansky Lab.

The core of the data acquisition system was the computer sitting inside the shielded rack in the telescope pedestal. For a single feed the received signals would be modulated onto optical fiber and demodulated at the Jansky Lab for processing. In the absence of enough fibers from the telescope to the Jansky Lab, the digital back end had to be placed in the telescope pedestal, inside a military-grade shielded rack to reduce interference generated by the digital equipment that would corrupt observations at this and other telescopes on site. A fiber modem inside the rack permitted Internet access and remote control of the data acquisition system. 20 inputs are synchronously sampled at 1.25 Msamp/sec with 12-bit quantization and streamed to disk. Beams are formed in post processing with the recorded voltage samples. Logging in remotely to the data acquisition computer over Remote Desktop allows the user to set LO frequencies through a Labview/HPIB interface, begin data acquisition, and transfer files back to the lab. This remote operation is a primary reason to locate the LO signal generators in the rack despite the cable loss over the 170ft lengths of cable up to the FEB.

Roughly 0.5km away, a 3.6m antenna was used as a high-gain reference for interference mitigation tests. Throughout this chapter, “reference antenna” refers to this 3.6m antenna, and “telescope” refers to the Green Bank 20 Meter Telescope. The signal from the reference antenna was directly modulated onto fiber and fed into the shielded rack in the telescope pedestal where it was downconverted and sampled synchronously with the PAF signals. Pointing for the reference antenna was controlled from the Jansky Lab using an RS-232/optical link out to the M² Antenna Systems, Inc. (Fresno, CA) antenna rotator control boxes sitting in their own shielded enclosure.

3.3 RF Front End

3.3.1 Green Bank 20 Meter Telescope

Completed in 1994, the Green Bank 20 Meter Telescope was built by RSI and funded by the US Naval Observatory as part of a very long baseline interferometer for earth orientation measurements. It was shut down in 2000 due to budget constraints and re-commissioned in 2007

for these BYU PAF experiments. The reflector has a diameter of 20m and $f/d = 0.43$, with rms surface accuracy of 0.8mm. The dish has an elevation-over-azimuth mount design with a fast slew rate of $2^\circ/\text{second}$ on each axis and 34 arcsec rms pointing residuals [83]. The 20m diameter gives a beamwidth of roughly $\theta_{3\text{dB}} \approx \lambda/D \approx 0.5^\circ$ at 1600 MHz.



(a) The Green Bank 20 Meter Telescope.

(b) The reflector surface.

Figure 3.2: The Green Bank 20 Meter Telescope. (a): The Green Bank 20 Meter Telescope where the 19-element BYU PAF was mounted for experiments. The reflector is on an elevation-over-azimuth mount and the telescope pedestal room houses the PAF digital back end. (b): The reflector surface shown here puts into perspective the huge collecting area, emphasizing the viability of the telescope for astronomical observations. The reflector has surface accuracy of 8mm rms and can slew at the very fast rate of $2^\circ/\text{sec}$ on each axis.


Figure 3.2(a) shows a view of the telescope with the 20m reflector sitting on top of an air-conditioned concrete pedestal room that houses the pointing control electronics, temperature control, and BYU digital back end. Figure 3.2(b) gives a sense of the size of the reflector surface.

The far-field for this antenna begins at roughly $D^2/\lambda \approx 2\text{km}$. This distance is important because the RFI cancelation algorithm described in Chapter 5 was developed for a source in the far field. The possibility of placing a transmitter on a nearby mountain top was considered but deemed impractical. Instead, interference was generated on site using signal generators and horn antennas in fixed locations or with motion generated by placing the antenna in the bed of a moving truck or carried by hand while walking back and forth in front of the Jansky Lab. The algorithm of Chapter 5 assumed a source in the far field, so the proximity of the interference to the telescope

may be one reason the algorithm did not perform as expected for on-reflector interference cases. In that chapter, the real data experimental verification was done with the array in an anechoic chamber.

3.3.2 Phased Array Feed

Figure 3.3(a) shows the PAF mounted at the focus of the telescope, and Figure 3.3(b) shows a close-up of the array mounted on the FEB. The PAF consists of 19 dipoles arranged in



(a) BYU PAF mounted at the focus of the telescope.

(b) Close-up view of the of the array.

Figure 3.3: The BYU PAF mounted at the focus of the Green Bank 20 Meter Telescope. (a): BYU PAF mounted at the focus of the telescope. (b): Close-up view of the of the array clamped to the front of an NRAO standard format front end box (FEB) for the Green Bank 20 Meter Telescope.

a hexagonal pattern over a ground plane 26 inches (0.66 meters) in diameter. The array antenna elements are all single-polarization balun-fed half-wave dipoles tuned to 1600 MHz, placed 0.25λ from the ground plane. The dipoles use a simple, robust design with slightly thickened arms for increased bandwidth, made of 1/4" diameter copper tubing. The arms attach to a section of semi-rigid coaxial cable that feeds through the ground plane directly into SMA-connectorized low noise amplifiers. The other length of copper making the balun shown in Figure 3.4(a) was hollowed semi-rigid coax. The ground plane was made of standard copper-clad FR-4 laminate printed circuit board material. For increased rigidity the ground plane was reinforced with additional rib strips of FR-4 soldered on the back of the ground plane.

Array elements achieve a 570 MHz impedance bandwidth (-10 dB reflection coefficient) or a 35% bandwidth relative to the center frequency, which is modest but adequate for initial tests. The spectral line of neutral hydrogen (HI) is at 1420 MHz, and two main spectral lines for hydroxyl (OH) are at 1612 MHz and 1665 MHz. Detections of the bright HI source CygA and OH masers W49N and W30H at these frequencies are shown in Chapter 4. The array elements are arranged in two concentric hexagonal rings around a center dipole, shown in Figure 3.4(b). The hexagonal grid permits 0.6λ spacing between dipoles, slightly farther apart than the typical $\lambda/2$ used for a rectangular array. The increased inter-element spacing reduces mutual coupling without undersampling or producing grating lobes. Details of the array design and construction can be found in [82, 84].




Figure 3.4: Dipole dimensions and numbering. (a): Dimensions of the dipoles used in the PAF. The dipoles are balun-fed half-wave dipoles tuned to 1600 MHz [14]. (b): Layout and numbering of the array antenna elements in a hexagonal pattern over the ground plane. Elements are spaced 0.6λ (at 1600 MHz) from each other to slightly reduce mutual coupling without introducing grating lobes in the beam pattern.

The dipoles connect through the array ground plane to uncooled low-noise amplifiers (LNAs) surrounded by Styrofoam for better temperature stability. Initial experiments on the Green Bank 20 Meter Telescope in 2007 used Mini-Circuits, Inc. (Brooklyn, NY) ZEL-1217 amplifiers with 29dB gain and 1.5dB noise figure or 120 Kelvin noise temperature. Later experiments in 2008

used custom amplifiers from Ciao Wireless, Inc. (Camarillo, CA) with 33 K noise temperature and 41 dB gain at 1600 MHz. These LNAs were measured by L. Belostotski (U. Calgary) to have the following noise parameters at 1600 MHz: $T_{\min} = 33K$, $R_n = 3.4\Omega$, and $\Gamma_{\text{opt}} = 0.07\angle 90^\circ$.

3.3.3 Analog Front End Electronics




Figure 3.5: RF receiver electronics inside the FEB. At right, the frequency downconverter boxes feed IF signals to the rear bulkhead panel at left. The LO amplifier/splitter network is in the lower center of the image and DC power supplies are at left.

Signals from the LNAs feed through 24' lengths of minibend® (Astrolab, Inc.) flexible coaxial cable into receiver electronics that perform a two-stage downconversion to the final intermediate frequency (IF) centered at 2.8125 MHz. Figure 3.6 shows a photo inside one of the 10 room-temperature dual-channel frequency downconverter boxes. The electronics use primarily SMA-connectorized Mini-Circuits components because of their low cost and ease of assembly. The RF signal was first passed through an image-rejection bandpass filter at 1200-1700 MHz. The first local oscillator (LO1) does a high-side mix tunable around 2000 MHz to mix RF down to the first IF at 400 MHz. A 10 MHz wide bandpass filter centered at 400 MHz was followed by LO2 fixed at 396 MHz to mix down to the final IF of 2.8125 MHz. The final IF filter bandwidth was 425 kHz. An example tuning scenario to receive RF = 1420.8 MHz is: LO1 = 1819.6125 MHz,

and $LO2 = 396$ MHz, yielding $IF = 2.8125$ MHz. Additional design details for the frequency downconverter electronics are given in [82].




Figure 3.6: Frequency downconverter box [82]. Electronics perform a 2-stage mix from 1600 MHz RF (tunable from 1400-1800 MHz) to 2.8125 MHz IF. Components are primarily SMA-connectorized commercial components from Mini-Circuits connected with semi-rigid coaxial cable.

The two LOs sit inside the shielded rack in the telescope pedestal room and the two LO signals are carried up to the FEB using high quality low-loss LMR-400 coaxial cable from Times Microwave Systems, Inc. (Wallingford, CT). At these LO frequencies there is still some loss through the cables, so the LO sources are set to lower power and then amplified inside the FEB to get the power high enough for each signal to be split 10 ways to feed each downconverter box and then split once more inside the box to feed both receiver channels with the +7dBm required by the mixers. Details on the LO amplification and splitting network can be found in [84].

3.3.4 Front End Box (FEB)

The PAF was mounted on an NRAO standard format front end box (FEB) for the Green Bank 20 Meter Telescope that houses the receiver electronics and measures 28in \times 28in \times 60in. The array ground plane was clamped to a 28in \times 28in square aluminum plate using an aluminum ring 1/2 inch wide, 25.5 inches in diameter, and 1/16 inch thick that screws into threaded holes in

the square plate, with seams waterproofed using RTV, seen in Figure 3.3(b). Analog front end electronics are mounted inside the FEB, shown in Figure 3.5.

The FEB was made of steel with 1/2 inch plywood sheets attached inside the top and bottom, allowing electronics to be mounted using drywall screws. These plywood sheets also create a channel for circulating air past Peltier units on the top and bottom of the FEB to stabilize and maintain the inside temperature at 25C. An additional fan blows directly onto the LO amplifier/splitter network since these amplifiers went unstable when they became too hot. A thermistor sits inside the FEB with sensor line running down the telescope cable wrap to a standard control unit that stabilizes FEB temperature by variously setting the Peltier units to cool or heat as needed.

The entire assembly was sealed and waterproofed, with IF signals feeding out of the FEB through SMA bulkhead feedthroughs on the back panel and LO signals feeding in through N-type bulkhead feedthroughs sealed with RTV. Figure 3.7 shows the rear of the FEB with the bulkhead feedthrough panel.




Figure 3.7: The FEB viewed from behind with IF cables feeding out the rear bulkhead panel and down the telescope feed support arm to the pedestal room housing the digital back end. Peltier units on the FEB maintain 25°C temperature stability.

3.4 Telescope Pedestal

3.4.1 IF and LO

In Figure 3.7 the cables running between the FEB and pedestal run along a feed support arm and through a cable wrap system inside the pedestal. IF signals are carried down along 170ft lengths of high quality coaxial cable LMR195 from Times Microwave. LMR195 is double-shielded including a foil and braid outer conductor specified to give 90dB shielding. At 10 MHz the loss is less than 1dB/100ft, so the loss introduced by the long cable run is negligible at the IF of 2.8125 MHz.

Inside the shielded rack the IF signals pass through a tray of final IF filters and amplifiers to a final bandwidth of 425 kHz before arriving at the ADC. The IF filter tray has three choices of filter bandwidths, 425 kHz, 1 MHz, and 5 MHz, so the ADC can run at 1.25, 2.5, and 10 Msamp/sec although in practice the system has been configured exclusively with the 425 kHz bandwidth filters and sampled at 1.25 Msamp/sec because of the disk write speed and limited number of filters available at the higher bandwidths.

The 2.8125 MHz IF was chosen to permit baseband subsampling techniques with the below-Nyquist sample rate of 1.25 Msamp/sec producing an aliased copy of the 425 kHz IF band, yielding the desired image for the IF between 0 and 625 kHz at the ADC inputs. The primary RF frequencies used in Chapter 4 are mostly within the range of 1400-1650 MHz although the RF was tunable from 1200-2000 MHz using LO1 without modifying the final IF.

Figure 3.8 shows the rack that houses the digital back end. The rack was a military-grade shielded rack with 90dB shielding effectiveness that was tested in the observatory's anechoic chamber to be sure it met the RFI leakage standards for equipment installed on site. IF cables feed into the rack through SMA connectors on a bulkhead panel, LO signals feed up to the FEB through N-type connectors on the bulkhead panel. Internet and reference antenna signals feed in over fiber through a length of waveguide attached to the bulkhead.

Two Agilent 8648D signal generators sit in the rack to provide LO signals. Power levels and frequency tuning are controlled remotely from the Jansky Lab by logging into the data acquisition computer and running a LabVIEW interface that communicates with the signal generators over an GPIB connection. The GPIB control was enough of a convenience to motivate having



Figure 3.8: Shielded rack holding the digital back end and signal generators for LO signals. A single server-class PC holds 5 ADC cards with 4 inputs per card for a total of 20 synchronously sampled inputs at 1.25 Msamp/sec. All data are streamed to disk for post-processing.

the signal generators in the rack instead of the FEB but LO cable loss was a source of RFI and required approval from the site RFI management team. The amplifier/splitter network described in Section 3.3.4 reduced the required power level at the signal generators to reduce RFI from LO cable loss.

3.4.2 Digital Back End

A single network server class PC houses all data acquisition cards and a high speed array of hard drives for real-time data streaming to disk. 20 inputs (19 PAF signals and one reference antenna) are synchronously sampled by the data acquisition computer at 1.25 Msamp/sec with 12-bit quantization for high dynamic range. The computer hosts five ADC cards from National

Instruments with four inputs each for a total of 20 inputs. The server was set up with a 4-drive, 15000 rpm SCSI RAID 0 array configured for maximum speed, striping across all four drives without redundancy. Disk access was limited by the PCI bus, not disk writing. At 1.25 Msamp/sec on all 20 channels, data can be recorded continuously for 20 minutes at a time and then immediately restarted every 20 minutes for very long acquisitions. The maximum acquisition length in this mode is nearly two hours, when the 247 GB disk array fills up (20 inputs x 12-bit samples x 1.25 Msamp/sec = 37.5 MB/sec).

Data acquisition as well as LO control can be done at the base of the telescope during maintenance time when the door can be open on the shielded rack, or remotely by logging into this computer over the Internet. One LabVIEW VI (a program in LabVIEW is called a virtual instrument, or VI) communicates with the LOs over GPIB, one VI displays power spectral density (PSD) plots for each of the 20 inputs, and one VI runs the data acquisitions. Real-time PSD displays provide quick verification that the various amplifier stages are working properly based on a quick look at the noise floors levels and flatness before running a data acquisition. The VI that runs the data acquisition communicates through TCP/IP with the telescope control computer to ensure the telescope is pointed correctly before acquiring data.

In post-processing 6144-sample blocks of 12-bit data are unpacked into 16-bit integers and then complex-basebanded in a digital receiver. The final IF center frequency of 2.8125 MHz with 425 kHz bandwidth was chosen so that sampling sub-Nyquist at 1.25 Msamp/sec places aliased copies of the data at $\pi/2$ radians/sec discrete frequency. This means the digital mix to complex baseband uses $\cos[n\pi/2]$ and $\sin[n\pi/2]$, which conveniently evaluate to ± 1 and $\pm j$ (where $j = \sqrt{-1}$) for all values of n , meaning that the mixing operation is done efficiently by simply negating the real/imaginary parts of the signal at appropriate times. Finally the samples are decimated by 2 and a 70-tap digital Parks-McClellan low-pass filter is applied. Codes in Matlab and C perform complex basebanding, correlating, beamforming, imaging, and RFI mitigation, all in post-processing.

The array is spatially narrowband because of its small physical extent. Using $BW \ll c/d \approx (3 \times 10^8 \text{ m/s}) / (0.66 \text{ m}) \approx 450 \text{ MHz}$, it can be concluded that since the digital back end has only 425 kHz bandwidth, no additional frequency channelization is required before beamforming. As the system incrementally moves toward a science-ready digital back end with hundreds of

MHz bandwidth the array will no longer be narrowband and frequency channelization will be incorporated in the data processing.

3.5 Reference Antenna

RFI mitigation algorithms exist that benefit from a high-gain reference copy of the interferer [50, 85]. In these experiments a 3.6 meter antenna was used as the auxiliary, located about 0.5km from the telescope and steered toward the RFI source. Figure 3.9 shows the 3.6m dish with the Green Bank 20 Meter Telescope in the background. Following the LNA at the antenna, the RF signal was modulated onto fiber which fed into the Jansky Lab, then back out to the shielded rack in the telescope pedestal. In the shielded rack the signal feeds through a frequency downconverter box like those in the FEB, described in Section 3.3.3, and then is sampled synchronously with the signals from the PAF.

The antenna uses a custom commercial positioner design from M² Antenna Systems, controlled remotely from the Jansky Lab over RS-232 modulated onto fiber. The control boxes produced enough RFI that a special enclosure was built for them so they could sit in Zone 1 at the antenna. Using the commercial software package NOVA this auxiliary antenna can track satellites like GPS, GLONASS, Iridium, or weather satellites like GOES-11.

3.6 Jansky Lab

The control rooms at the Jansky Lab are housed inside an RF enclosure to reduce radiation onto the site from the telescope control computers and racks of digital back end electronics. Control room windows have wire mesh and outer doors seal with a locking mechanism to ensure the room acts as a good Faraday cage.

Control for the BYU PAF system was located next to the control computer for the Green Bank 20 Meter Telescope. Scripts lay out a series of telescope pointings during each experiment, and the telescope control computer was configured to communicate over TCP/IP with the data acquisition computer in the telescope pedestal to be sure the telescope had arrived at the desired pointing and stopped moving before acquiring data.



Figure 3.9: 3.6m reference antenna in front of RFI enclosure for positioner control boxes, with the Green Bank 20 Meter Telescope in the background.

Next to the telescope control computer, a separate computer was used to log into the data acquisition computer in the shielded rack to configure LO frequencies and begin data acquisitions. This computer also had the satellite tracking program NOVA installed, and using its serial port connected to a fiber modem it controlled the reference antenna pointing.

3.7 Conclusions

This chapter has presented details of one of the first astronomical PAF systems in existence, built by BYU in collaboration with NRAO. It is an important proof-of-concept platform for studying beamformer design in an astronomical context where high sensitivity and good calibration are required. The system design was greatly simplified by the choice of a narrow system bandwidth allowing sampled voltages to be streamed to disk for analysis in post-processing. The next chapter describes experiments done using this system to characterize the system, detect astronomical radio sources, and demonstrate RFI cancellation.

Building a large scale system like this one presents many logistical challenges, especially since it is one of the first of its kind at NRAO-GB and the design has been a learning experience for all involved. The primary system improvements and extensions required for this system to become

a fully science-ready instrument include: increased system bandwidth, dual-polarization antennas, cryo-cooled LNAs, and a simplified user interface. Each of these tasks is a large undertaking because of the large number of antenna inputs for the array.

The system bandwidth must be increased by a factor of 1000, from a few hundred kHz to a few hundred MHz. Correlations and beamformers will need to be formed for each frequency channel. Data storage requirements are problematic at the high sample rates required for wideband operation, so real-time signal processing in hardware will be required. Analog beamforming at the front end may reduce the data rate through the system if a few beams are formed, but if the number of formed beams is on the order of the number of array elements, this solution may not offer any advantage over forming beams at the back end. Work is already underway to build an FPGA-based digital back end, described in Chapter 6.

A dual-polarization array is required for studying polarized radio sources. Several designs have already been created and tested, including one in development for the GBT [86]. That design also includes cryo-cooling for LNAs – a challenge because there are so many LNAs for the array. Developing a convenient user interface for the system is a step that will happen once the system design has transitioned over to an observatory and the key research questions concerning PAF design and characterization have been answered.

Chapter 4

First PAF Results on the Green Bank 20 Meter Telescope

4.1 Introduction

The experimental results reported in this chapter represent a significant contribution to the effort to make PAFs a viable technology for radio astronomy, presenting some of the first astronomical observations in the world made using a PAF. The experiments here focus on two main goals: first, proof-of-concept demonstrations that astronomical sources can be detected using this PAF; second, characterizations of the array in terms of its sensitivity and noise response. The approach taken at BYU has been to focus on narrowband, high-sensitivity results, where other groups internationally have focused on broader bandwidth antennas and receivers. The isotropic noise response characterization experiments described in Section 4.4.3 are the first of their kind in the radio astronomy world, and following this example, other groups around the world have constructed measurement facilities like the one described here.

Forming beams with the array requires steering vectors, or beamformer calibrators, that specify the appropriate gains and phases to be applied to each array element. This chapter begins in Section 4.2 with the argument that these beamformer calibrators must be found using experimental data from the array mounted on the telescope rather than pre-computed calibrators based on array geometry. Building a geometry-based model that takes into account all effects of scattering from feed supports, antenna element radiation patterns, and receiver gains and phases is unrealistic. All experimental PAF results reported so far have used beamformers based on in situ calibrations. The calibration procedure described in Section 4.2 produces a lower variance estimate of the calibration vector than simply using the dominant eigenvector of the signal covariance matrix.

Section 4.3 describes proof-of-concept experiments showing astronomical detections using the PAF. Total power measurements are shown using continuum sources (Cygnus A) and OH sources (W49N and W3OH), as well as a spectral detection of the OH line at 1665 MHz. The

experimental platform built from commercial off-the-shelf components has receiver electronics stable enough in gain and phase to permit re-use of calibration vectors obtained on different days. The beam sensitivity, gain and aperture efficiency are reported, based on experiments using off-the-shelf Mini-Circuits low-noise amplifiers with 120K noise temperature. These numbers are a good starting place for improvements to the array, and later experiments discussed in Sections 4.4-4.6 use custom-built amplifiers with 33K noise temperature to achieve higher sensitivity.

One of the most important results coming from this work is the development of a Y-factor technique using room temperature microwave absorber and cold empty sky to characterize the isotropic noise response of an astronomical array. The isotropic noise response of an array is used in tying the beamformer output to physically meaningful results, discussed in Section 4.4.1, and in characterizing the PAF system temperature and aperture efficiency, discussed in Section 4.6.1. The novel measurement technique described in Section 4.4.3 has been adopted by other groups in the astronomical PAF community, and is becoming a standard technique for characterizing the PAF.

Mutual coupling perturbs the embedded dipole radiation patterns and increases the overall system noise temperature. Section 4.5.1 shows antenna range measurements of the bare array, with ripples in the dipole radiation patterns due to mutual coupling. These perturbed patterns give additional evidence that beamformer calibrators are best found from experimental data because of the difficulty to account for all array response factors in simulation. Section 4.5.2 reports the equivalent noise temperature for the dipoles and receivers, showing the roughly 20K increase in noise temperature compared with measurements of the LNAs and receivers alone.

The isotropic noise response is used in Section 4.6 to compute the sensitivity and noise temperature for the PAF. Comparing these results with those reported by Oosterloo [36] for the APERTIF array mounted at WSRT in April 2008 ($2.2\text{m}^2/\text{K}$ sensitivity and $125\text{K}-135\text{K}$ T_{sys}), the proof-of-concept sensitivity measurement reported in Section 4.3.2 (November 2007, using 120K amplifiers) are slightly lower but the results from Section 4.6.2 (August 2008, using 33K amplifiers) are 50% better. These results are $1.33\text{m}^2/\text{K}$ and $3.3\text{m}^2/\text{K}$ sensitivity and 112K and 66K T_{sys} for 2007 and 2008, respectively.

Increasing sky survey speed by forming multiple beams on the sky is one of the primary motivations for astronomical PAF development. Section 4.7 presents several radio camera images made from grids of MaxSNR beams using calibration vectors obtained experimentally. Another

promising aspect of PAFs is the ability to spatially cancel RFI using adaptive beamforming. RFI mitigation is demonstrated in Section 4.3.6, showing the frequency spectrum before and after cancellation. Section 4.7.2 shows an RFI-corrupted radio camera image before and after cancellation.

4.2 Calibration

The basic approach to deterministic beamformer design is to apply temporal FIR filter design techniques to the spatial filtering problem where the antenna elements are spatial samples of a continuous aperture and the beamformer coefficients are chosen using, e.g., windowed filter or equiripple filter design. This approach assumes a fully calibrated array, i.e., the receiver gains and phases are known exactly as well as the array response to sources from arbitrary directions, including all antenna effects of the dish structure and antenna elements. Willis [35] demonstrates that well-behaved PAF beam shapes can be designed in simulation using deterministic beamforming techniques with little coma, distortion, or polarization rotation as beams are steered off axis.

Deterministic beamformer design requires many beamformer calibrators that are hard to obtain experimentally with sufficient signal to noise ratio. Additionally, due to gain instabilities and difficulties in matching cable lengths it can be hard to transfer beamformers designed in simulation to an experimental system. For these reasons the PAF community has favored statistically optimal techniques in practice. The maximum sensitivity beamformer has been used for most experimental PAF observations reported to date, and the only demonstrated experimental radio camera images formed to date use statistically optimal, adaptive, data-dependent beamforming algorithms [24–26, 28, 36]. Section 2.3.2 describes the adaptive beamformers used in this chapter.

The adaptive beamformer design approach requires knowledge of the array covariance structure, which can be estimated for a stationary signal environment and periodically updated with calibrations. Adaptive beamforming is used for the experimental results reported here because it accounts for the spillover noise contribution, interference sources, and mutual coupling. Adaptive beamforming can be thought of in two senses: fixed-adaptive, where the beamformer is chosen using the array covariance estimated once from a calibration data set and then held constant while the dish is steered across the sky, and active-adaptive, where the beamformer is updated repeatedly to actively cancel interference or adapt to the covariance changes with new dish pointings.

For each beam mainlobe peak formed by the array and any other points on the beam where constraints are desired, a calibration vector \mathbf{a}_i is required to steer the beam in direction Ω_i . Periodic recalibration is necessary as receiver gains and phases vary in time. The calibration procedure is as follows [14]:

1. *Noise-only covariance, $\check{\mathbf{R}}_\eta$* : Steer the array to a (relatively) empty patch of sky so $\mathbf{x}[n] \approx \boldsymbol{\eta}[n]$ for a long period (several minutes) and compute $\check{\mathbf{R}}_\eta = \frac{1}{L} \mathbf{X} \mathbf{X}^H$ as in (2.6).
2. *Signal-plus-noise covariance, $\check{\mathbf{R}}_i$* : While tracking the brightest available source, steer the array in directions Ω_i around the source, giving $\mathbf{x}[n] = \mathbf{a}_i s[n] + \boldsymbol{\eta}$, where \mathbf{a}_i is the desired calibration vector at direction Ω_i . Calculate $\check{\mathbf{R}}_i$ over the same period as $\check{\mathbf{R}}_\eta$ in the previous step.
3. *Array response, \mathbf{a}_i* : Compute $\mathbf{a}_i = \check{\mathbf{R}}_\eta \mathbf{u}_{\max}$ where \mathbf{u}_{\max} is the maximum solution to the generalized eigen equation $\check{\mathbf{R}}_i \mathbf{u}_{\max} = \lambda_{\max} \check{\mathbf{R}}_\eta \mathbf{u}_{\max}$. This noise-whitening method produces a lower variance estimate with less bias arising from the correlated noise in \mathbf{R}_η than does the simpler approach of using the dominant eigenvector of $(\check{\mathbf{R}}_i - \check{\mathbf{R}}_\eta)$. The vector \mathbf{a}_i is assumed unit-norm so that $\mathbf{R}_s \approx \lambda_{\max} \mathbf{a}_i \mathbf{a}_i^H$.

4.3 Proof-of-Concept PAF Demonstration

This section presents the first results obtained with the BYU PAF during the Summer and Fall of 2007. The main purpose of these experiments is to demonstrate that the PAF can actually detect astronomical sources and begin to characterize the PAF in terms of sensitivity and system noise temperature. The beamwidth of a formed beam is compared with that for the center element of the array. A spectral line detection is presented first, followed by measurements of the sensitivity, gain, and aperture efficiency based on total power measurements of CygA. Initial tests were run to determine if the PAF would benefit from sitting off-focus to spread the received energy across more array elements, explore adapting the spillover noise response as the telescope tips down in elevation, and demonstrate RFI cancelation for a moving RFI source. In later sections these results are improved upon to get better sensitivity, characterize the array more fully, and create radio camera images.

4.3.1 OH Spectral Line Detection

As a test of astronomical source detection using the PAF, Figures 4.1(a) and 4.1(b) demonstrate spectral line detection of the OH maser W49N, using a MaxSNR beamformer. The source is weak enough that the spectral line detection required differencing the PSDs for an on-source pointing and an off-source (noise-only) pointing. This experiment demonstrates a couple of important characteristics: first, the PAF can detect astronomical sources; second, the system is stable enough to observe sources that require several seconds of integration for the on and off pointings. This second point is even more remarkable because the calibration vector that was used in forming the beam was actually found on a different day from when these on/off pointings at the OH source were taken, meaning that the system is stable in gain and phase over multiple days.

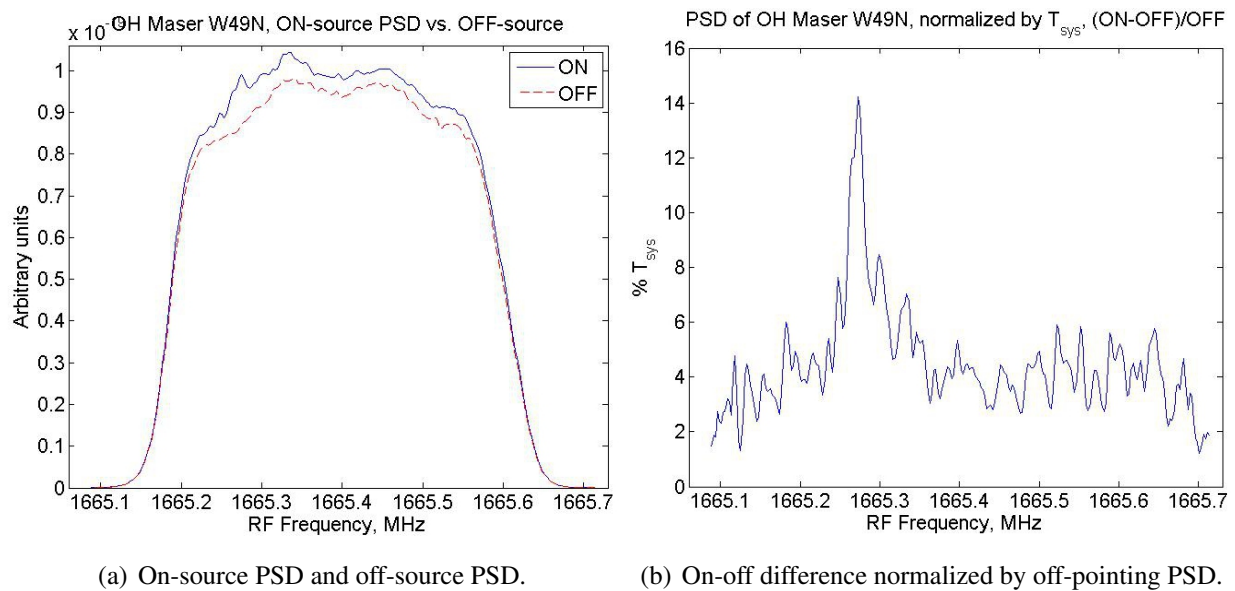


Figure 4.1: Spectral line detection of OH source W49N. (a): PSDs for OH maser W49N on-pointing and off-pointing, using MaxSNR formed beam in each case. (b): W49N OH maser on-off difference normalized by off-pointing PSD. The OH line at 1665.401MHz shows up at 1665.27MHz. The beam was formed using a MaxSNR beamformer with calibration vector taken on a different day, demonstrating system gain and phase stability.

Figure 4.1(a) shows two PSD plots in linear units, for an on-source pointing (signal plus noise) and an off-source pointing (noise only). The PSD was computed using Welch's periodogram with 256 frequency bins, over a 7 second data record for each pointing. The difference between the

two plots is the PSD due to the signal only, and is shown in Figure 4.1(b), in units normalized by the noise-only PSD, i.e., $(\text{PSD}_{\text{on}} - \text{PSD}_{\text{off}})/\text{PSD}_{\text{off}}$. The main OH line at 1665.401 MHz detected here can be seen Doppler shifted to 1665.27MHz due to its motion away from the earth. The line is greater than 6x larger than the residual rms deviations in the noise floor, indicating successful detection.

4.3.2 Beam Sensitivity


The radio telescope gain, sensitivity, and aperture efficiency were determined using CygA, a 1380 Jy source, by pointing the telescope over a series of steps to the left, directly at, and to the right of the source to measure on-source and off-source power. Using the total power from the on-source and off-source pointings, the antenna temperature, T_A due to the source can be estimated by applying Equation (2.1):

$$T_{\text{on}} = T_A + T_{\text{sys}}, \quad T_{\text{off}} = T_{\text{sys}}.$$

With T_A from this measurement and $F^s = 1380$ Jy from the known source flux density, the gain, sensitivity, and aperture efficiency are then found using (2.2)-(2.5).

Figure 4.2(a) shows the received power as a function of time as the telescope stepped in cross-elevation (perpendicular to an elevation slice) across the source in 24 arcmin (half beamwidth) steps, with 20 second dwell times per pointing. The slice across the source has discrete steps instead of a smooth curve because fewer pointings were used to give clear boundaries between on-source and off-source pointings and to give plenty of integration time at each pointing. Vertical axis units are in Kelvin, found using $T = P/(k_B B)$, where P is received power in Watts, T is equivalent noise temperature in Kelvin, k_B is Boltzmann's constant, and B is system bandwidth in Hz. ADC voltages were calibrated using known input sources in the lab, then voltages were referred to a 50Ω source to get power in Watts. Here 390kHz is used for the bandwidth to make the numbers more physically sensible. The figure compares the antenna temperature for the center dipole and for a formed beam using a MaxSNR beamformer.

Subtracting the T_{off} from T_{on} to get T_A gives $209\text{K} - 140\text{K} = 69\text{K}$ for the center element and $187\text{K} - 112\text{K} = 75\text{K}$ for the formed beam. This leads to $G = T_A/F^s = 0.050$ K/Jy for the center element and 0.055 K/Jy for the formed beam.



(a) Center element and MaxSNR antenna temperature. (b) Offset plot to illustrate narrower beam and higher gain for MaxSNR.

Figure 4.2: Beam sensitivity comparison for center dipole and MaxSNR formed beam, in a slice across CygA (1380 Jy) in 24 arcmin steps dwelling 20 seconds/pointing. Long dwell times and distinct steps between pointings make it easier to distinguish between on-source and off-source pointings. (a): Formed beam Gain and Sensitivity are 0.055 K/Jy and 1.33 m²/K, respectively, compared with 0.05 K/Jy and 0.99 m²/K for the center element. (b): Offset plot created by adding 28K to the MaxSNR beam to align T_{sys} , demonstrating narrower beamwidth and higher gain of MaxSNR beam. The center element sees higher power during the off-source pointing between 1.4 and 1.6 minutes indicating a wider main beam. The maximum overall power is higher for the MaxSNR beam due to its higher gain.

From (2.3), sensitivity is found using

$$\frac{A_{\text{eff}}}{T_{\text{sys}}} = \frac{k_B B}{S^{\text{inc}}} \text{SNR} \quad (\text{m}^2/\text{K}).$$

where A_{eff} is the effective antenna collecting area in m², T_{sys} is the system noise temperature in Kelvin, and S^{inc} is the incident signal power density in one polarization in W/m². Incident signal power S^{inc} is found using $1/2 \times 10^{-26} \times B \times F^s$, where the 1/2 comes from assuming the source power is spread equally across both polarizations but these dipoles only receive one polarization, and $10^{-26} \times B$ converts from Jy (10^{-26} W/m²/Hz) to power density in W/m². For this observation using CygA,

$$\begin{aligned} \frac{A_{\text{eff}}}{T_{\text{sys}}} &= \frac{k_B B}{S^{\text{inc}}} \text{SNR} \\ &= \frac{1.38 \times 10^{-23} \times 390 \times 10^3}{1/2 \times 10^{-26} \times 390 \times 10^3 \times 1380} \text{SNR}. \end{aligned}$$

SNR is found using $T_A/T_{\text{sys}} = 69\text{K}/140\text{K}$ for the center element and $75\text{K}/112\text{K}$ for the formed beam. Thus $A_{\text{eff}}/T_{\text{sys}} = 0.99 \text{ m}^2/\text{K}$ for the center element and $1.33 \text{ m}^2/\text{K}$ for the formed beam, meaning a 33% signal processing improvement.

Aperture efficiency is found using $\eta_{\text{ap}} = A_{\text{eff}}/A_{\text{phys}}$. For the 20m telescope diameter, $A_{\text{phys}} = \pi(20/2)^2$ and A_{eff} is found by multiplying the sensitivity ($A_{\text{eff}}/T_{\text{sys}}$) by T_{sys} . For the center element and formed beam, these are respectively:

$$\frac{0.99 \text{ m}^2/\text{K} \times 140\text{K}}{100\pi\text{m}^2} = 44\% \quad \text{and} \quad \frac{1.33 \text{ m}^2/\text{K} \times 112\text{K}}{100\pi\text{m}^2} = 47\%.$$

These numbers are lower than those expected for a good single feed, but they are still reasonable results, especially considering the initial proof-of-concept nature of this experiment. Compared to the only other astronomical PAF in existence at the time, these results from November 2007 are slightly lower than those reported the following April by Oosterloo [36] for the APERTIF array of Vivaldi elements mounted at the WSRT, $2.2\text{m}^2/\text{K}$ sensitivity and $125\text{K}-135\text{K}$ T_{sys} . The noise temperature of the LNAs on the BYU PAF during this experiment is 120K . Later experiments using 33K LNAs give better sensitivity results, presented in Section 4.6. That section gives improved results from later experiments with the BYU PAF, and includes sensitivity values for several large radio telescopes in Table 4.5 for comparison.

4.3.3 Dish Illumination Control

The previous section calculated a 33% sensitivity improvement using a MaxSNR beamformer over a center element beam. Figure 4.3 illustrates the MaxSNR spillover noise minimization to improve SNR and sensitivity. The azimuth scan through CygA compares power output from the center element with the output of a fixed MaxSNR beamformer that minimizes total power subject to the constraint of unity gain in the boresight direction. Power is normalized with respect to the maximum power seen through the slice. The difference between power on-source and off-source is greater for the formed beam because of the spillover noise minimization achieved with MaxSNR.

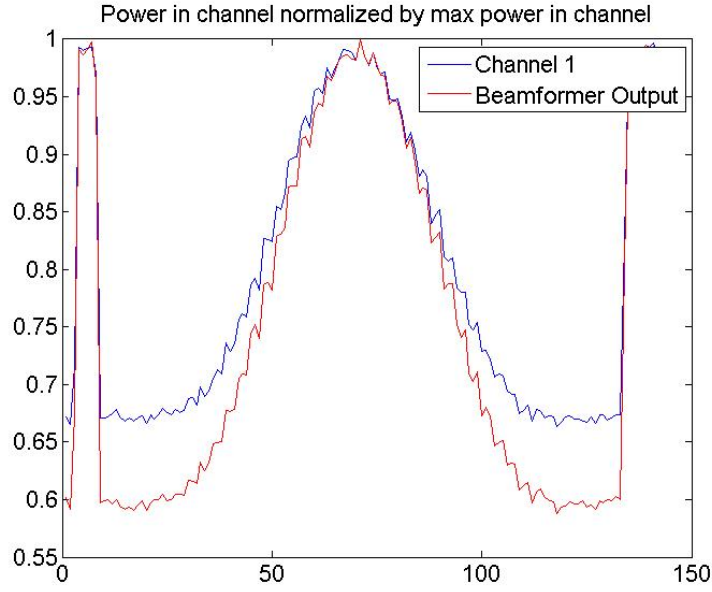


Figure 4.3: Slice across CygA comparing center element with fixed MaxSNR formed beam. The difference between on-source and off-source power is greater for MaxSNR because it optimizes the beam pattern to reduce spillover noise so the power from CygA is a greater percentage of the total power.

4.3.4 PAF De-focusing

PAF beamforming relies on the imperfect ability of the parabolic reflector to focus all of the available energy to the center of the focal plane, allowing neighboring antenna elements to collect some additional power that can be combined to form a beam. Since the majority of the available energy is focused on the center element, the beamformer has reduced flexibility in controlling the dish illumination beam pattern because the outer ring elements are de-emphasized. From an FIR filter perspective, the reflector forces a tapered window weighting on the array elements. Another set of experiments involved moving the PAF in front and behind the focal plane to intentionally sample energy before it is fully focused, spreading more energy to the outer ring PAF elements. Some improvement is gained by moving the PAF closer to the dish for telescopes with short focal length because of the resulting reduction in spillover efficiency. For the Green Bank 20 Meter Telescope, however, the best sensitivity is actually when the PAF is right at the focal plane. This is an important result discussed in Jacob Waldron’s thesis [84].

4.3.5 Adaptive Spillover Noise Control

An advantage to active PAFs over offset horn feed FPAs is the ability to adaptively optimize the aperture efficiency and spillover efficiency. As the dish steers down in elevation the spillover region is not solely 250K ground, but part of the spillover now includes some of the sky above the horizon. Figure 4.4 explores the idea of computing a new MVDR beamformer at each elevation pointing in an attempt to tradeoff spillover from hot ground in favor of illuminating the spillover region including cold sky, reducing the overall spillover noise level. The response in the boresight direction is always constrained to be unity gain, while the power everywhere else is minimized. The output power from each beamformer is assumed to be only due to noise since there is no signal of interest in these elevation slices, so the lower power curve is more desirable since it represents less noise power. There is a spike at 40° elevation, due to the appearance of an unexpected source entering through a sidelobe. Had it entered through the mainlobe the response should be the same since both beamformers used the same unity mainlobe gain constraint, and the SNR also appears lower for the active beamformer which is consistent with the beamformer's goal of minimizing total power outside of the mainlobe.

Assuming the noise reduction is only due to trading sidelobes looking at the ground for sidelobes looking at the sky, it was surprising that the noise power increased significantly at low elevations. The plot here served as the basis for much more extensive analysis presented in [14], showing that some improvement can be gained through updating the beamformer to optimize the spillover at lower elevation, but in fact the dominant contribution to the changing antenna temperature as the dish steers in elevation is the increased atmosphere seen through the high gain main beam. Here in Figure 4.4 that same effect is seen where at low elevation the total power increases significantly as the mainlobe sees more atmosphere, thus increasing the apparent sky temperature.

4.3.6 RFI Mitigation Demonstration

Another advantage of PAFs for radio astronomy is the ability to adaptively place beam-pattern nulls for canceling interference. Figure 4.5 shows the PSD before and after interference cancelation. While the dish tracked CygA at 1600 MHz, a CW tone in the middle of the RF band was broadcast from the bed of a moving truck. The fixed beamformer clearly shows a strong

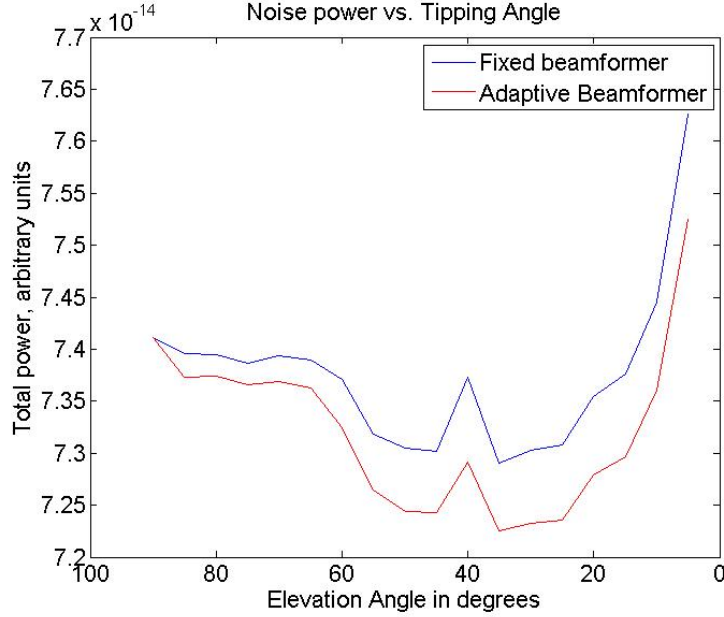


Figure 4.4: Elevation slice through (ideally) empty sky, comparing actively updated MVDR beamformer and fixed MVDR beamformer to demonstrate actively optimizing the spillover noise reduction. With no signal of interest, the beamformer output is noise-only, so the lower output power of the active beamformer is desirable [23].

source in the middle of the band that disappears with cancellation. Here RFI mitigation is done using subspace projection (SP), described in Section 2.3.3. In this example, 100 short-term integrations (STIs) of length 6144 samples (5 milliseconds) are used to estimate the j^{th} sample array covariance matrix $\hat{\mathbf{R}}^{(j)}$ as in (2.6). For the j^{th} STI, the SP beamformer is found by applying (2.11) and (2.12) on a per-STI basis:

$$\begin{aligned} \mathbf{P}_{\text{int}}^{\perp(j)} &= \mathbf{I} - \mathbf{U}_{\text{int}} \mathbf{U}_{\text{int}}^H \\ \mathbf{w}_{\text{SP}}^{(j)} &= \mathbf{P}_{\text{int}}^{\perp(j)} \mathbf{w}, \end{aligned}$$

where \mathbf{U}_{int} is estimated as the dominant eigenvector of $\hat{\mathbf{R}}^{(j)}$ and \mathbf{w} is a MaxSNR beamformer calibrated for a boresight beam. The beamformer $\mathbf{w}^{(j)}$ is then applied to all data within the j^{th} STI.

Knowledge of temporal signal characteristics was not required; this interference mitigation is done spatially with beampattern nulling. Because the INR is high in this data set, the contribution to the array covariance matrix due to interference is easily identified. In Section 4.7.2 this approach

is applied for multiple beam pointing directions. Chapter 5 more fully addresses SP interference mitigation, extending this approach to cases where the interference contribution to $\hat{\mathbf{R}}$ is difficult to identify because of low INR and rapid motion.

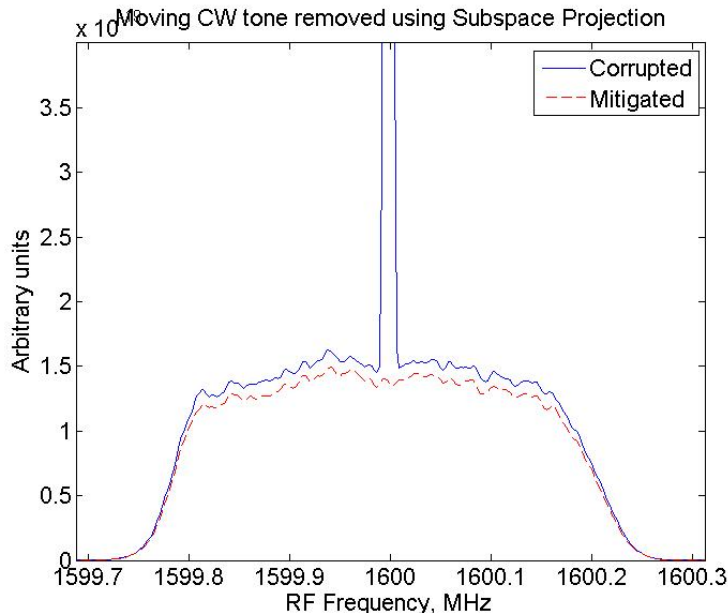


Figure 4.5: Demonstration of RFI mitigation using SP beamforming to remove a CW tone from the middle of the band. The interference was broadcast from the bed of a moving truck while the telescope tracked CygA [23].

4.4 Array Noise Response Characterization

Defining antenna figures of merit for a beamforming array is more complicated than for a single antenna because of the dependence on the beamforming coefficients. Several definitions and formulas that connect passive single feed antennas to active beamforming arrays are included here to emphasize the importance of the quantity $\mathbf{R}_{\text{ext,iso}}$ in characterizing an active beamforming array. Section 4.4.2 describes a generalized Y-factor approach to measure $\mathbf{R}_{\text{ext,iso}}$ experimentally and then Section 4.4.3 discusses experiments conducted at NRAO-GB to obtain these data for the BYU PAF. Section 4.5.2 reports on a test to measure the dipole noise temperature reduction with greater spacing between elements and compares noise temperatures measured at the LNA with the dipoles in various configurations. All results in this section use the 33K LNAs from Ciao Wireless.

4.4.1 Significance of $\mathbf{R}_{\text{ext,iso}}$ in Connecting Active Arrays and Passive Antennas

Denoting the output power for a beamforming array due to some incident signal as

$$P_{\text{sig}} = \mathbf{w}^H \mathbf{R}_{\text{sig}} \mathbf{w},$$

it is apparent that the output power depends on the beamforming coefficients in two ways: 1) the beamformer weights steer the beam toward or away from the signal, and 2) an arbitrary scaling on the coefficients appears to scale the power, i.e., $\mathbf{w}' = 10\mathbf{w}$ leads to $P_{\text{sig}} = \mathbf{w}'^H \mathbf{R}_{\text{sig}} \mathbf{w}' = 100\mathbf{w}^H \mathbf{R}_{\text{sig}} \mathbf{w}$. The necessary antenna definitions and formulas to tie the output of a beamforming array to physically meaningful values are found in [38], including reference plane definitions. Consider a single antenna surrounded by an isotropic noise field at temperature T_{iso} (290K at room temperature). The available power external to this single antenna (prior to antenna losses due to radiation efficiency and mismatch) is

$$P_{\text{ext,iso}} = k_B T_{\text{iso}} B,$$

where k_B is Boltzmann's constant, B is the system bandwidth, and T_{iso} is the brightness temperature of the isotropic noise field. For a beamforming array the external power available is the same, $P_{\text{ext,iso}} = k_B T_{\text{iso}} B$, which can be rewritten as

$$P_{\text{ext,iso}} = \mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}.$$

Shifting the reference point to immediately after the antenna, the *beam isotropic noise response* [38] is

$$P_{\text{t,iso}} = \mathbf{w}^H \mathbf{R}_{\text{t,iso}} \mathbf{w},$$

where $\mathbf{R}_{\text{t,iso}} = \mathbf{R}_{\text{ext,iso}} + \mathbf{R}_{\text{loss}}$, under the condition that the physical temperature of the array is also T_{iso} . The difference between the external available power and the beam isotropic noise response is the radiation efficiency [38],

$$\eta_{\text{rad}} = \frac{P_{\text{ext,iso}}}{P_{\text{t,iso}}}.$$

Comparing the beam isotropic noise response with the power available from the isotropic noise environment gives the necessary scaling to convert the formed beam power back to the physical available power. This scaling is called the *available receiver gain* [38],

$$\begin{aligned}
G_{\text{rec}}^{\text{av}} &= \frac{P_{\text{t,iso}}}{k_{\text{B}}T_{\text{iso}}B} \\
&= \frac{\mathbf{w}^H \mathbf{R}_{\text{t,iso}} \mathbf{w}}{k_{\text{B}}T_{\text{iso}}B} \\
&= \frac{\mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}}{\eta_{\text{rad}}k_{\text{B}}T_{\text{iso}}B}, \tag{4.1}
\end{aligned}$$

which serves to undo the arbitrary scaling introduced by the beamformer coefficients. The *beam equivalent available power* is the output of the beamformer scaled by the available receiver gain $G_{\text{rec}}^{\text{av}}$. For an arbitrary signal, the *available signal power* [37, 38] is

$$P_{\text{sig}}^{\text{av}} = \frac{P_{\text{sig}}}{G_{\text{rec}}^{\text{av}}} = \left[\frac{\eta_{\text{rad}}k_{\text{B}}T_{\text{iso}}B}{\mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}} \right] \mathbf{w}^H \mathbf{R}_{\text{sig}} \mathbf{w},$$

which ties the output of the beamformer to a physical quantity that corresponds to the power seen by an equivalent passive antenna. In addition to making a connection between active arrays and passive antennas, the two preceding equations are used in finding the system noise temperature, T_{sys} , and the effective area, A_{eff} [38]:

$$T_{\text{sys}} = \frac{T_{\text{noise}}}{G_{\text{rec}}^{\text{av}}} \quad A_{\text{eff}} = \frac{P_{\text{sig}}^{\text{av}}}{S_{\text{sig}}},$$

where the signal of interest is a plane wave of incident flux density S_{sig} and polarization matched to the beam such that A_{eff} is at a maximum. This further illustrates the usefulness of $\mathbf{R}_{\text{ext,iso}}$.

4.4.2 A Novel Y-factor Approach to Obtain $\mathbf{R}_{\text{ext,iso}}$

Calculation of the available receiver gain requires knowledge of $\mathbf{R}_{\text{ext,iso}}$, which can be estimated through a matrix generalization of the Y-factor approach. In the Y-factor technique the equivalent noise temperature of the receiving device is found by using a hot load and cold load of known temperatures T_{hot} and T_{cold} to produce two noise power outputs when connected at the

input of the device:

$$\begin{aligned} N_{\text{hot}} = P_{\text{hot}} + P_{\text{noise}} &= k_{\text{B}}T_{\text{hot}}B + k_{\text{B}}T_{\text{noise}}B \\ N_{\text{cold}} &= k_{\text{B}}T_{\text{cold}}B + k_{\text{B}}T_{\text{noise}}B. \end{aligned} \quad (4.2)$$

Finding the ratio

$$\frac{N_{\text{hot}}}{N_{\text{cold}}} = \frac{T_{\text{hot}} + T_{\text{noise}}}{T_{\text{cold}} + T_{\text{noise}}} = Y$$

and solving for T_{hot} in terms of Y gives

$$T_{\text{eq}} = \frac{T_{\text{hot}} - YT_{\text{cold}}}{Y - 1}, \quad (4.3)$$

where T_{eq} represents the device noise temperature as an equivalent resistor placed at the input to the device prior to any gain through the system that would introduce the same amount of noise power $k_{\text{B}}T_{\text{noise}}B$ added by the device.

To solve for $\mathbf{R}_{\text{ext,iso}}$, the receiver is surrounded with isotropic noise environments at temperatures T_{hot} and T_{cold} . For a noiseless receiver the structure of the covariance stays constant and is just scaled by the appropriate temperature so

$$\frac{\mathbf{R}_{\text{hot}}}{T_{\text{hot}}} = \frac{\mathbf{R}_{\text{ext,iso}}}{T_{\text{iso}}} = \frac{\mathbf{R}_{\text{cold}}}{T_{\text{cold}}}$$

or

$$\mathbf{R}_{\text{hot}} = \frac{T_{\text{hot}}}{T_{\text{iso}}} \mathbf{R}_{\text{ext,iso}}$$

for a noiseless receiver.

Experimentally, \mathbf{R}_{hot} and \mathbf{R}_{cold} are estimated as sample covariance matrices obtained with absorbing foam covering the full field of view or the array looking at cold empty sky, respectively, as described in the next section. In that case, the noise terms added by the receiver and cable loss

are included:

$$\hat{\mathbf{R}}_{\text{hot}} = \frac{T_{\text{hot}}}{T_{\text{iso}}} \hat{\mathbf{R}}_{\text{ext,iso}} + \hat{\mathbf{R}}_{\text{rec}} + \hat{\mathbf{R}}_{\text{loss}} \quad (4.4)$$

$$\hat{\mathbf{R}}_{\text{cold}} = \frac{T_{\text{cold}}}{T_{\text{iso}}} \hat{\mathbf{R}}_{\text{ext,iso}} + \hat{\mathbf{R}}_{\text{rec}} + \hat{\mathbf{R}}_{\text{loss}}. \quad (4.5)$$

Solving for $\hat{\mathbf{R}}_{\text{ext,iso}}$ by subtracting (4.5) from (4.4):

$$\begin{aligned} \hat{\mathbf{R}}_{\text{hot}} - \hat{\mathbf{R}}_{\text{cold}} &= \frac{T_{\text{hot}} - T_{\text{cold}}}{T_{\text{iso}}} \hat{\mathbf{R}}_{\text{ext,iso}} \\ \frac{T_{\text{iso}}}{T_{\text{hot}} - T_{\text{cold}}} (\hat{\mathbf{R}}_{\text{hot}} - \hat{\mathbf{R}}_{\text{cold}}) &= \hat{\mathbf{R}}_{\text{ext,iso}}. \end{aligned} \quad (4.6)$$

Plugging $\hat{\mathbf{R}}_{\text{ext,iso}}$ from (4.6) back into (4.4) to solve for $\hat{\mathbf{R}}_{\text{noise}}$:

$$\begin{aligned} \hat{\mathbf{R}}_{\text{hot}} &= \frac{T_{\text{hot}}}{T_{\text{iso}}} \frac{T_{\text{iso}}}{T_{\text{hot}} - T_{\text{cold}}} (\hat{\mathbf{R}}_{\text{hot}} - \hat{\mathbf{R}}_{\text{cold}}) + \hat{\mathbf{R}}_{\text{rec}} + \hat{\mathbf{R}}_{\text{loss}} \\ \hat{\mathbf{R}}_{\text{hot}} - \frac{T_{\text{hot}}}{T_{\text{hot}} - T_{\text{cold}}} (\hat{\mathbf{R}}_{\text{hot}} - \hat{\mathbf{R}}_{\text{cold}}) &= \hat{\mathbf{R}}_{\text{rec}} + \hat{\mathbf{R}}_{\text{loss}} \\ \frac{T_{\text{hot}} \hat{\mathbf{R}}_{\text{cold}} - T_{\text{cold}} \hat{\mathbf{R}}_{\text{hot}}}{T_{\text{hot}} - T_{\text{cold}}} &= \hat{\mathbf{R}}_{\text{rec}} + \hat{\mathbf{R}}_{\text{loss}}. \end{aligned} \quad (4.7)$$

Equation (4.6) shows how to find the available receiver gain $G_{\text{rec}}^{\text{av}}$ in (4.1), and (4.7) gives insight into the noise added by the receiving system, \mathbf{R}_{rec} and \mathbf{R}_{loss} .

4.4.3 Experimental $\mathbf{R}_{\text{ext,iso}}$ Measurements

Measuring $\mathbf{R}_{\text{ext,iso}}$ as in (4.6) requires two isotropic noise environments to produce \mathbf{R}_{hot} and \mathbf{R}_{cold} . As a lower cost alternative to heating or cooling an anechoic chamber, copper screen was used to block out ground spillover from the array backlobes and the array was presented with noise fields at two temperatures. Figure 4.6 shows the pyramidal structure used for $\mathbf{R}_{\text{ext,iso}}$ measurements. The hot isotropic noise source was microwave absorber that filled the opening of the copper ground screen, and for the cold isotropic noise source the array was pointed toward (ideally) empty sky. T_{hot} was ambient temperature, 17°C or 280K, and T_{cold} was 7K, an estimate provided by NRAO slightly higher than the cosmic microwave background temperature to account for some scattered weak sources in the sky. Other groups around the world involved in astronomical



Figure 4.6: Overhead view of ground screen. The array is mounted on the FEB containing the front end receiver electronics, and the wood frame of the ground screen sits on the structure that holds the FEB in the outdoor test facility shown here. The copper screen is rounded around the edges of the wood frame to help reduce diffraction that would introduce interference from sources around the screen.

PAF work have followed our lead and built similar ground screen enclosures for their own isotropic noise response measurements.

Ground Screen Construction

The ground shield was made from copper screen stapled to a wood frame. The smaller opening of the pyramid was made to fit snugly down around the FEB, with inner dimensions just larger than the 28in×28in FEB. The larger opening widens to 10ft×10ft. The FEB sat in a positioning fixture that also served to support the ground screen. To allow the ground screen to make electrical contact and help seal the bottom opening of the pyramid, copper screen was clamped to the array ground plane and draped down the sides of the FEB, shown in Figure 4.7.




Figure 4.7: Bare array with copper mesh clamped down around it

A wooden frame was built to fit around the FEB that sat on the positioning fixture supporting the weight of the ground screen and located the base of the ground screen level with the array ground plane. The copper screen was stapled to the wood frame and soldered every few inches along each seam to reduce RF leakage. The copper extended beyond the top opening of the frame and was rounded to reduce diffraction at the edges of the ground screen. The ground screen was lifted into place using a wood frame with ropes tied to the corners that lifted up from below with ropes tied to a cherry picker.

Hot Absorber/Cold Sky Setup

The hot isotropic noise field was created using an 8ft×8ft sheet of absorber made from standard 2ft×2ft squares of 8 inch microwave absorber absorber glued to a plywood backing. The absorber was suspended above the copper pyramid from the cherry picker that would lower it into the large opening of the ground screen and set it so the tips of the absorber nearly touched the copper screen to completely fill the aperture of the array. To be sure the presence of the absorber so close to the array would not change the electrical properties of the array, the dipole S-parameters were compared with the absorber at various distances from the array. Little change was observed, indicating the presence of the absorber does not affect the array electrically.



Figure 4.8: Absorber being lowered into the ground screen

Computing $\mathbf{R}_{\text{ext,iso}}$

For each experimental run, the absorber was held over the array while recording data for 30 seconds, then removed and another 30 seconds of data were recorded looking at the sky. In each experimental run the hot/cold noise fields were presented in quick succession to keep the array characteristics as consistent as possible for the Y-factor test. The sample covariance matrices $\hat{\mathbf{R}}_{\text{hot}}$ and $\hat{\mathbf{R}}_{\text{cold}}$ were then computed using the absorber data set for $\hat{\mathbf{R}}_{\text{hot}}$ and the sky data set for $\hat{\mathbf{R}}_{\text{cold}}$, with Eq. (2.6). Plugging $\hat{\mathbf{R}}_{\text{hot}}$ and $\hat{\mathbf{R}}_{\text{cold}}$ into (4.6) with $T_{\text{hot}} = 280\text{K}$ and $T_{\text{cold}} = 7\text{K}$ gives $\mathbf{R}_{\text{ext,iso}}$. Note the presence in (4.6) of T_{iso} . This value is arbitrary and cancels out wherever $\mathbf{R}_{\text{ext,iso}}$ is used.

4.5 Mutual Coupling Measurements

Mutual coupling across the PAF results in non-ideal signal and LNA noise correlation that must be taken into account in beamformer design, noise level calculation, and sensitivity optimization. The pattern distortion (relative to the simple dipole model) seen in Figure 4.9(b) results from re-radiation to neighboring antennas. Section 4.5.2 demonstrates the increased system temperature for this array, comparing measured noise temperatures for an isolated antenna with the noise temperature for antennas embedded in the array. Simple windowed filter beamformer design is less effective because of the signal coupling across the array, and the changes induced in the array overlap matrix (a scaled version of the array covariance assuming an isotropic noise field) motivate the statistically optimal beamforming design discussed in Section 4.2.

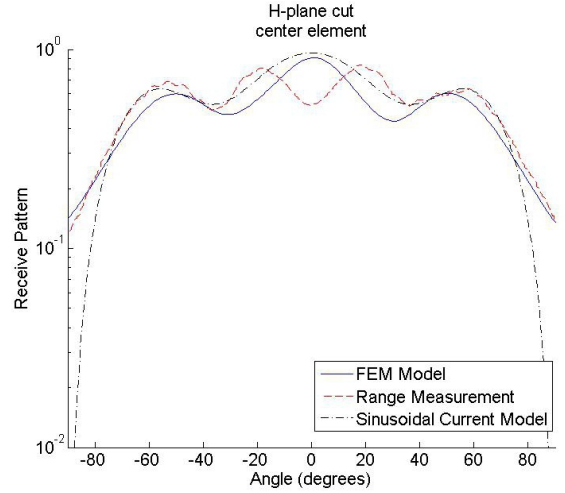
4.5.1 Element Radiation Patterns with Mutual Coupling

To measure the array element radiation patterns, the PAF was characterized on the outdoor antenna test range at NRAO-GB, shown in Figure 4.9(a). The FEB was mounted on a rotating positioner on the left tower and a fixed transmitter was on the right tower. To get cuts through orthogonal antenna polarizations the front plate of the FEB was rotated 90° with the center element kept on the axis of rotation. The E-plane and H-plane (as in Figure 4.9(b)) are defined as follows: the E-plane is any plane containing the electric field vector and the direction of maximum radiation and the H-plane is any plane containing the magnetic field vector and the direction of maximum radiation [87, 88]. For a \hat{z} -directed dipole, the E-plane is any plane containing the z-axis – i.e., the y-z plane or x-z plane; varying the polar angle (the angle with respect to the z-axis) in one of these planes gives a half-sine shaped radiation pattern. For the same \hat{z} -directed dipole, the H-plane is the x-y plane; varying the azimuthal angle (the angle with respect to the x-axis in the x-y plane) gives a constant radiation pattern.

In the H-plane cut in Figure 4.9(b), the receiving dipole and transmitting dipole are both vertical throughout the azimuth scan so the radiation pattern would ideally be constant but the measured radiation patterns shows significant variations due to mutual coupling. Figure 4.9(b) plots the modeled radiation pattern using a model that includes mutual coupling, compared with the measured radiation pattern for the center dipole. In the figure, experimentally measured results are



(a) Outdoor antenna range at NRAO-GB.



(b) Embedded element radiation pattern.

Figure 4.9: Antenna range measurements. (a): Outdoor antenna test range at NRAO-GB used to measure array element radiation patterns. The FEB is mounted on the tower at left and a fixed transmitter is located on the tower at right. The FEB sits on a positioner that can rotate $\pm 180^\circ$ and the cables are run down to the room at the base of the telescope where the positioner control and PAF digital back end are located for data recording. (b): Center element H-plane radiation pattern, comparing measured results with simulated results, demonstrating the pattern variations induced by mutual coupling [23].

compared with two modeled curves, one using a detailed finite-element method (FEM) model for the dipole array elements, and the other using sinusoidal line currents for the array elements [89]. The significant agreement is qualitative in nature; both models predict two major dips in received power but the measured pattern shows three dips. Still, the overall shape and scale of the variations is a good match. To verify that the effects were not due to multipath interference, the distance from the transmitter to the array was changed by 0.3λ but the pattern variations were unchanged. Out toward $\pm 90^\circ$ off boresight, neighboring elements and even the FEB itself block the transmitter, so the pattern drops off as expected. The 0.6λ spacing between array elements was chosen to reduce mutual coupling without introducing grating lobes for elements spaced too far apart [81].

4.5.2 Increased System Noise Temperature Due to Mutual Coupling

Several experiments were run to measure the equivalent noise temperature, T_{eq} , of the dipoles embedded in the array at various separations. Values of T_{eq} were measured using the Y-factor method with noise sources connected at the LNA inputs. These were compared with T_{eq} measured using hot/cold noise fields in the ground screen measurement setup described earlier.

These measurements show the increase in T_{eq} due to antenna connector and cable loss, mismatch, and mutual coupling.

Embedded Element Noise Temperatures

The first column of Table 4.1 reports T_{eq} measured using the standard Y-factor technique from (4.3) using SMA connectorized loads connected at the input to the LNAs. The second and third columns of Table 4.1 give measured values of T_{eq} for the LNAs connected to the array presented with hot and cold noise fields using microwave absorber and cold sky with the array sitting in the ground screen setup using Eq. (4.8). For these measurements, microwave absorber or cold sky were presented to the array for 30 seconds at a time, recorded and correlated to estimate \mathbf{R}_{hot} and \mathbf{R}_{cold} , and the total power measured on each dipole is the corresponding diagonal element of the correlation matrix. The equivalent noise temperature for the i^{th} array element is found using the Y-factor technique on each correlation matrix diagonal element:

$$Y_i = \frac{[\mathbf{R}_{\text{hot}}]_{ii}}{[\mathbf{R}_{\text{cold}}]_{ii}} \quad T_{\text{eq},i} = \frac{T_{\text{hot}} - Y_i T_{\text{cold}}}{Y_i - 1}. \quad (4.8)$$

As a baseline comparison to remove any mutual coupling effects, a separate ground plane was constructed with a single dipole in the center, shown in the second column. The full array measurement in the third column shows the results for the full array with all LNAs and receivers connected.

There is a significant increase in noise temperature going from the individual LNA T_{eq} measurements to the single dipole measurement, even without any mutual coupling penalty. The 17K increase is likely due to cable and connector loss in the dipole and the mismatch between the 50Ω LNA and the (ideally) 73Ω dipole radiation resistance.

Another increase in T_{eq} due to mutual coupling appears on each element for the full array measurement. The increase is less for elements on the outer ring (elements 8-19) than for interior elements, consistent with the intuition that mutual coupling is higher for a dipole surrounded by neighbors than for a dipole in isolation. Overall, the center element sees a 34K increase in T_{eq} , 17K due to individual dipole design, and 17K due to mutual coupling. The first 17K can be reduced using lower-loss materials, more careful connector design, and impedance matching to the LNA.

Table 4.1: Noise temperature measurements (all values given in Kelvin)

Antenna No.	Coaxial Load LNA	Isolated Dipole	Single Element in Array
1	30.5	47.5	64.2
2	34.0		64.1
3	35.3		65.4
4	33.5		64.8
5	34.3		64.4
6	31.3		60.9
7	33.8		65.6
8	33.1		54.4
9	33.0		57.9
10	30.5		58.1
11	31.8		63.1
12	31.7		56.3
13	32.1		54.5
14	31.9		58.0
15	33.7		56.3
16	31.6		59.3
17	34.8		66.9
18	32.2		59.3
19	33.3		58.8

The second 17K can be reduced using active matching to compensate for the mutual coupling penalty. Results of work to reduce the increased noise temperature are reported in [45].

Sliding Dipole Measurements

To characterize the effect of dipole proximity on mutual coupling an experiment was run with a single dipole in the center of a ground plane and a second dipole that could be positioned at various distances along two perpendicular tracks, shown in Figure 4.10. The sliding dipole was soldered to a 2in×2in square copper base and then held in each new position with copper tape. The rest of the track was covered with copper tape during each experimental run as indicated in the figure. Table 4.2 shows the results. In the first column the sliding dipole had been terminated with a 50Ω resistor but in all other tests it was connected to an LNA and receiver chain so the mutual coupling would be symmetric (which appears to be the case). Coupling appears to disappear beyond 0.7λ.

Time constraints did not permit the repeated trials necessary to average out measurement errors, so there remains an open question in the 0.3λ and 1.2λ measurements. Since mutual coupling ought to reduce as the separation increases, it is surprising that the noise temperature appears higher on the center dipole at 0.3λ than at 0.25λ , and that the value on both dipoles is lower at 1.2λ than the value measured for a single dipole in isolation, reported in Table 4.1. Further inquiry is required to determine whether the effect is due to constructive interference at certain dipole separations or simply due to estimation error that would reduce with repeated trials.




Figure 4.10: Sliding dipole tests assembly. One dipole was fixed in the middle of the ground plane, while tracks allowed a second dipole to be placed at different distances from the center dipole to measure the effect of dipole separation on mutual coupling.

Table 4.2: Sliding dipole noise temperature measurements, in Kelvin

Separation	0.2λ	0.25λ	0.3λ	0.45λ	0.7λ	1.2λ
Center Dipole	155.5	52.7	54.1	51.2	48.1	46.5
Sliding Dipole (50Ω -terminated)		57.9	55.5	49.6	49.7	46.7

Various Terminations

Table 4.3 gives the results of experiments to see the effect of having the center element connected to its LNA and the rest of the elements unterminated (open circuit), short circuited, and 50Ω terminated. In the open circuit case, neighboring dipoles nearly disappear and T_{eq} is nearly the same as for a single isolated dipole. There is still a 3 K noise temperature increase on the center dipole because the neighboring dipoles now look like $\lambda/4$ dipoles where current can still flow but they are much less efficient radiators in the band of interest. The short circuit case gives a sense for the mutual coupling increase due only to the presence of the metal dipoles around the center element in the absence of any LNA reverse noise radiated to the center element from the LNAs on the other dipoles. Finally the 50Ω terminations turn the other dipoles into radiators with room temperature resistors connected so the noise temperature seen by the center element is much higher because it sees nearby 290K sources. The noise temperature only increases to 94.7K rather than 290K due to mismatch between the 50Ω terminators and their dipoles and the fact that the center element does not have perfect coverage of hot dipoles filling its beam pattern.

Table 4.3: Measured noise temperatures in Kelvin, for array with LNA on center dipole and various terminations on other dipoles

All other dipoles:	Open	Short	50Ω
Center dipole:	50.9	59.2	94.7

4.6 Sensitivity

With the knowledge gained from the initial proof-of-concept experiments of Section 4.3 and new lower noise LNAs (33K instead of 120K), the PAF demonstrates much better sensitivity and system noise temperature in this section. The isotropic noise response $\mathbf{R}_{\text{ext,iso}}$ described earlier gives greater insight into T_{sys} and A_{eff} , discussed here. Section 4.6.2 includes a table with sensitivity values for several single feed telescopes for comparison.

4.6.1 Using $\mathbf{R}_{\text{ext,iso}}$ in Sensitivity Calculations

Sensitivity is found from the SNR using:

$$\frac{A_{\text{eff}}}{T_{\text{sys}}} = \frac{k_B B}{S^{\text{inc}}} SNR = \frac{k_B B}{S^{\text{inc}}} \frac{\mathbf{w}^H \mathbf{R}_{\text{sig}} \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\text{noise}} \mathbf{w}}. \quad (4.9)$$

Expressing sensitivity as $A_{\text{eff}}/T_{\text{sys}}$ gives a better comparison for telescopes of different sizes since a smaller telescope could be more sensitive than a larger one if its T_{sys} is lower or its aperture efficiency and radiation efficiency, η_{ap} and η_{rad} , are higher resulting in a higher A_{eff} .

Equation (4.9) shows that the SNR and sensitivity depend on the beamformer coefficients. To get greater insight into this beam dependence, T_{sys} and A_{eff} can be broken down as follows, using [14] and [38]:

$$\begin{aligned} T_{\text{sys}} &= \eta_{\text{rad}} T_{\text{iso}} \frac{\mathbf{w}^H \mathbf{R}_{\text{noise}} \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}} \\ &\approx \eta_{\text{rad}} T_{\text{iso}} \frac{\mathbf{w}^H \hat{\mathbf{R}}_{\text{off}} \mathbf{w}}{\mathbf{w}^H \check{\mathbf{R}}_{\text{ext,iso}} \mathbf{w}}, \end{aligned} \quad (4.10)$$

where $\hat{\mathbf{R}}_{\text{off}}$ is an estimate of the covariance for an off-source pointing, ideally at empty sky, and $\check{\mathbf{R}}_{\text{ext,iso}}$ is a calibration measured in advance using (4.6). In (4.10), η_{rad} can be written using [14] and [38] as:

$$\eta_{\text{rad}} = \frac{P_{\text{ext,iso}}}{P_{\text{ext,iso}} + P_{\text{loss}}} \quad (4.11)$$

$$= \frac{\mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}}{\mathbf{w}^H (\mathbf{R}_{\text{ext,iso}} + \mathbf{R}_{\text{loss}}) \mathbf{w}}. \quad (4.12)$$

Measuring \mathbf{R}_{loss} is difficult so [14] adopts the diagonal covariance model:

$$\mathbf{R}_{\text{loss}} = \sigma_{\text{loss}}^2 \mathbf{I} = \frac{T_{\text{loss}}}{T_{\text{iso}}} \frac{\mathbf{w}^H \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}} \mathbf{I},$$

where we use $T_{\text{loss}} = 5\text{K}$ from measurements of cable loss.

Effective area A_{eff} is found using

$$A_{\text{eff}} = \eta_{\text{rad}} \eta_{\text{ap}} A_{\text{phys}},$$

where A_{phys} is the known physical area of the reflector, η_{rad} is the radiation efficiency found using (4.12), and η_{ap} is the aperture efficiency found as the ratio of the collected incident power to the total available incident power:

$$\begin{aligned}
\eta_{\text{ap}} &= \frac{P_{\text{sig}}^{\text{av}}/\eta_{\text{rad}}}{A_{\text{phys}}S^{\text{inc}}} = \frac{1}{A_{\text{phys}}S^{\text{inc}}} \frac{P_{\text{sig}}}{\eta_{\text{rad}}G_{\text{rec}}^{\text{av}}} \\
&= \frac{1}{A_{\text{phys}}S^{\text{inc}}} \frac{k_{\text{B}}T_{\text{iso}}B}{\mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}} \mathbf{w}^H \mathbf{R}_{\text{sig}} \mathbf{w} \\
&= \frac{k_{\text{B}}T_{\text{iso}}B}{A_{\text{phys}}S^{\text{inc}}} \frac{\mathbf{w}^H \mathbf{R}_{\text{sig}} \mathbf{w}}{\mathbf{w}^H \mathbf{R}_{\text{ext,iso}} \mathbf{w}}.
\end{aligned} \tag{4.13}$$

4.6.2 Measured Sensitivity Results and Comparisons to Single Feed Telescopes

Table 4.4: Measured and modeled peak beam sensitivity, system temperature, and aperture efficiency for the 19 element prototype dipole array [14]

Performance Metric	Center Element	Formed Beam	Model
Sensitivity	2 m ² /K	3.3 m ² /K	3.7 m ² /K
T_{sys}	101 K	66 K	69 K
η_{ap}	64%	69%	81%

Table 4.4 reports the sensitivity, system temperature, and aperture efficiency for the BYU PAF. For comparison to single feed telescopes, Table 4.5 gives sensitivity results for several radio telescopes in m²/K as well as T_{sys} and aperture efficiency [90]. These numbers are much higher than the PAF results on the Green Bank 20 Meter Telescope because of the much larger physical areas of the telescopes in this table. Since the GBT and Effelsberg Telescope are both single-dish parabolic telescopes like the Green Bank 20 Meter Telescope, their sensitivity results can be easily scaled down to correspond to the physical area of the Green Bank 20 Meter Telescope. The GBT and Effelsberg Telescope have 100 meter diameters, giving $2500\pi\text{m}^2$ collecting area, compared with the $100\pi\text{m}^2$ of the Green Bank 20 Meter Telescope. Scaling the sensitivity values from Table 4.5 for these telescopes by $100/2500$ gives sensitivities of $11\text{m}^2/\text{K}$ and $6.6\text{m}^2/\text{K}$ for the GBT and Effelsberg Telescope, respectively. These numbers are better by factors of 2 and

nearly 4 than those reported in Table 4.4, indicating that there is still room for improvement. Since T_{sys} for the Effelsberg Telescope is smaller by more than a factor of 2 than T_{sys} for the PAF on the Green Bank 20 Meter Telescope, the sensitivity could surpass that of the Effelsberg telescope when the T_{sys} of the PAF is reduced to a similar level through the work already underway to design a cryo-cooled feed and array elements that compensate for mutual coupling.

Table 4.5: Sensitivity, system temperature, and aperture efficiency for several single feed radio telescopes [90]

Telescope	Sensitivity (m^2/K)	$T_{\text{sys}}(K)$	η_{ap}
Arecibo	683	40	0.7
EVLA	464	20	0.7
GBT	275	20	0.7
Effelsberg	165	30	0.63
WSRT	124	30	0.54

4.6.3 Sensitivity Degradation for Steered Beams

Steering beams off boresight to form a multi-pixel image improves sky survey speed but off-steered beams have degraded sensitivity because the array begins to look at the reflector at a skewed angle, de-focusing the focal spot Airy pattern and some of the de-focused energy is lost off the edge of the finite array aperture. Figure 4.11 plots sensitivity as a function of beam steering angle in half-power beamwidths (HPBW), roughly 0.7° for a 20m reflector. Beams should be within 1dB of the maximum sensitivity, so the figure suggests they should be steered no more than $\pm 1.25^\circ$ from boresight. One particularly interesting aspect of Figure 4.11 is the dips in sensitivity. Since the array is made of dipoles at discrete points it is not surprising that there are dips, but the location of the dips is surprising. There is a local minimum for the beam steered 0° off boresight, meaning that the energy is focused directly on the center element. Similarly for other steering directions, the sensitivity maxima occur for beams focused halfway between dipoles and minima occur when beams are focused on a dipole. This is caused by the complicated interactions with the dipoles and ground plane forming an “Eleven feed” configuration of two dipoles above

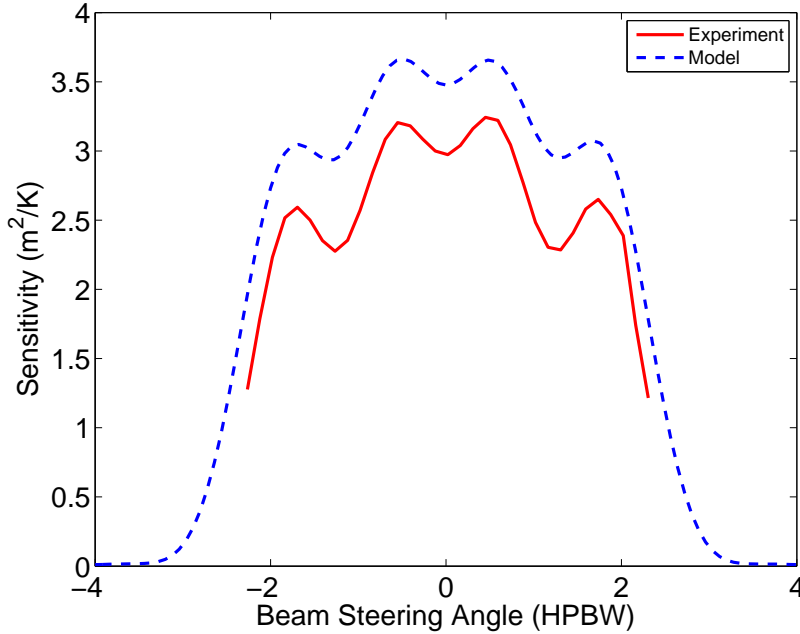


Figure 4.11: Sensitivity of the BYU PAF as a function of beam steering angle for beams steered off boresight [14].

a ground plane (“Eleven” refers to the linear dipoles next to each other looking like the number 11) that has been given recent attention by Olsson, Kildal and others for ultra-wideband antenna applications [91–93].

4.7 Radio Camera Images

The main appeal of astronomical PAFs is the survey speed improvement gained through forming multiple simultaneous beams on the sky, each with the resolution and sensitivity given by the large reflector. The images presented in this section represent some of the first multi-pixel or radio camera images created using an astronomical PAF, from [14].

For images larger than the 2.5° array field of view defined by the 1dB sensitivity limit as discussed in Section 4.6, the images are made from a mosaic of sub-images or tiles where each tile was created from a single dish pointing with a grid of beamformers applied at each pointing. In these mosaiced images the dish pointings were spaced approximately 1° . Voltage samples were streamed to disk for a single dish pointing, and then in post-processing a grid of beamformers was applied, with each beamformer steered to a different grid point to create one of the pixels in the image. Sidelobes are better than 10dB down for a beam steered 0.6° from boresight. Calibration

grids were done as described in Section 4.2, in grids of 33×33 and 65×65 pixels spaced 0.1° with a subset of the full calibration grid used in forming each mosaic tile.

4.7.1 Multiple Tile Image Mosaics




Figure 4.12: Mosaic image of W49N formed from a 3×3 grid of tiles, each made from a grid of 9×9 beams at each dish pointing [14].

Figure 4.12 shows an example of a multi-pixel image of the OH source W49N. The observation is over the 425 kHz system bandwidth centered on the 1612 MHz OH line. The 3×3 mosaic shows OH source W49N (3C398) in the top left and the nearby 3C397 at bottom right. Each of the nine mosaic tiles has 11×11 pixels from 121 simultaneous beams. The image is oversampled as compared to a more practical radio camera which would likely form beams spaced roughly at the HPBW crossover distance or less, but it illustrates the fine sampling possible. The number of beams that can be formed is limited only by the time required to get calibrators and the computational resources needed to form beams.

Figure 4.13(a) shows the Cygnus X region, observed at 1600 MHz with a 5×5 mosaic for 25 total pointings. The inset circle indicates the HPBW for a single pixel beam. As a comparison, Figure 4.13(b) is from the Canadian Galactic Plane Survey image [94], but blurred by convolution

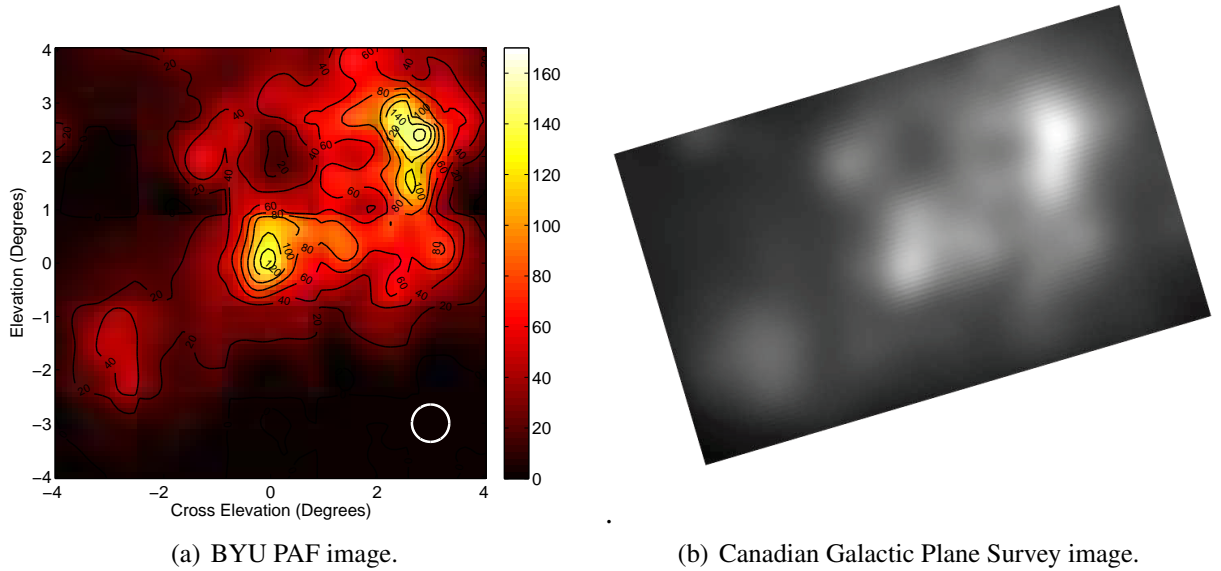


Figure 4.13: Cygnus X region. (a): Cygnus X region imaged using the BYU PAF at 1600 MHz using 5×5 mosaic. The image is oversampled with respect to the 0.7° main beam of the 20m denoted by a circle at bottom right, where a normal coarse grid would be formed with beams overlapping at the HPBW crossover points [14]. (b): Cygnus X region from the Canadian Galactic Plane Survey for comparison with (a). The image was convolved with the 0.7° beamwidth of the 20 Meter Telescope to produce comparable image resolution. The center of the map is approximately $20^h 44^m, +42^\circ$ (J2000) with north to the upper left [94].

with the equivalent beam pattern of a 20m telescope to match resolution scales. The Cygnus X radio camera image contains approximately 3000 pixels. A more practical coarse grid spacing of about half the HPBW would require 600 pixels. A single horn feed would require 600 pointings (one for each pixel) to form such an image, compared with 25 (one for each mosaic tile) for the radio camera. Assuming the single horn feed and the radio camera dwell the same length of time at each of the 600 or 25 pointings, respectively, this radio camera provides an imaging speed up of 24 times in forming the same 600-pixel image.

4.7.2 Demonstrating RFI Mitigation in a Radio Camera Image

The telescope tracked the OH source W30H with intentional interference created using a signal generator and hand-held antenna while walking back and forth about 1km from the telescope. The interference source was an FM modulated signal with 200kHz bandwidth centered at 1665 MHz obscuring the OH spectral line and half of the system passband. Figure 4.14(a) shows




Figure 4.14: RFI mitigation for a radio camera image of W3OH [14]. (a): Radio camera image before RFI source was turned on. (b): Image made from RFI-corrupted data. (c): Image made after subspace projection beamforming to cancel the interference.

the uncorrupted image using data collected before the interference source was turned on, Figure 4.14(b) shows an image created when the RFI source was turned on, and Figure 4.14(c) shows the image after RFI mitigation, with the signal of interest now clearly visible after cancellation with a subspace projection beamformer. The subspace projection beamformer [34, 63, 64, 70–72] will be discussed in greater detail in Chapter 5 where a motion model is applied to improve null depth for moving interferers with low INR. To produce Figure 4.14(c), The beamformer weights were

updated with each STI to adapt to the moving interference source using:

$$\begin{aligned}
\hat{\mathbf{R}}^{(j)}[\mathbf{U}_{\text{int}}|\mathbf{U}_{\text{s},\eta}] &= [\mathbf{U}_{\text{int}}|\mathbf{U}_{\text{s},\eta}]\mathbf{\Lambda} \\
\mathbf{P}_{\text{int}}^{\perp(j)} &= \mathbf{I} - \mathbf{U}_{\text{int}}\mathbf{U}_{\text{int}}^H \\
\mathbf{w}_{i,\text{SP}}^{(j)} &= \mathbf{P}_{\text{int}}^{\perp(j)}\mathbf{w}_{i,\text{MaxSNR}},
\end{aligned} \tag{4.14}$$

where j is the STI index, sample covariance $\hat{\mathbf{R}}^{(j)}$ is computed as in (2.6), and $\mathbf{w}_{i,\text{MaxSNR}}$ is the quiescent MaxSNR beamformer weight for the beam steered in the i^{th} direction. Equation (4.15) is an eigenvector decomposition of $\hat{\mathbf{R}}^{(j)}$ with eigenvalues in diagonal $\mathbf{\Lambda}$ ordered in descending magnitude such that the dominant eigenvectors in the partition \mathbf{U}_{int} span the interference subspace and $\mathbf{U}_{\text{s},\eta}$ spans the signal and noise subspaces. The full eigenvector matrix $\mathbf{U} = [\mathbf{U}_{\text{int}}|\mathbf{U}_{\text{s},\eta}]$ is unitary.

4.8 Conclusions and Future Work

The main contribution of these experiments has been to demonstrate detection of astronomical sources using a PAF and set in motion the work to reduce system noise temperature. The results presented in this chapter are the product of a major group effort with over 30 contributors at BYU and NRAO in which I played a major leadership role. Many of the plots and results in this chapter have come from the large group collaboration and have appeared in papers and presentations, cited where appropriate, to illustrate the significance of the experimental platform and its capabilities. The 2007 proof-of-concept experiments are some of the first astronomical detections made using a PAF, and are particularly notable because they were made using a system built by students from commercial off the shelf components. The system gain and phase are stable enough that beamformer calibrators can be reused on different days. The outdoor antenna range measurements demonstrated the radiation pattern distortion caused by mutual coupling and qualitatively verified the simulations models. PAF formed beams were compared with single antenna measurements, demonstrating noise response improvement. Statistically optimal beamformers adapting to the changing spillover noise environment, and RFI mitigation using subspace projection to remove a CW tone from the spectrum were demonstrated.

The array characterization experiments of 2008 were the first to perform array Y-factor tests to find $\mathbf{R}_{\text{ext,iso}}$ for an astronomical PAF. Theory motivating the importance of $\mathbf{R}_{\text{ext,iso}}$ was presented, along with experimental measurements with reasonable values. Increased T_{eq} was seen for a single antenna compared with the LNA-only T_{eq} , and then a further increase was seen for antennas embedded in an array. Some of the first multi-pixel PAF images were shown, including a comparison of Cygnus X with an independent image convolved with a 20m telescope beamwidth. RFI mitigation was shown for a multi-pixel image.

Future work focuses on improving the ability to calibrate the array and reduce the system noise temperature. Recent work to extend these results includes work to minimize the mutual coupling penalty through active noise matching [42] and designing antenna elements optimally matched to the noise parameters over a range of beam steering angles [45] to improve sensitivity without sacrificing aperture efficiency. NRAO is working with BYU to develop a cryo-cooled array for improved sensitivity. Work has been done to incorporate array calibrations into deterministic beamformer design [95], and continued work focuses on calibrator stability over time as well as new methods for calibration that do not require repeated time-consuming scans over bright sources to update the calibrator grids.

Improved calibration and imaging techniques will be required moving forward to combine the promising flexibility of PAFs with the reliability required for scientific use. Beam patterns must be characterized farther out into the sidelobes and system stability must be improved to enable weaker source detection. Imaging artifacts must be removed along the mosaic tile boundaries in mosaiced images seen in Figures 4.12 and 4.13(a).

Two additional pieces of future work are presented in the following chapters: better interference nulling and wider system bandwidth. Interference on the order of the noise floor is still problematic but must be canceled; this is the focus of Chapter 5 where a polynomial tracking algorithm is proposed for identifying and canceling interferers in difficult conditions. Wider system bandwidth significantly increases computational requirements. Sampling speeds must be increased from MHz to GHz and real-time processing is required to deal with the increased data rate. Chapter 6 discusses work done toward increased system bandwidth using an FPGA-based digital back end.

Chapter 5

Model-Based Subspace Projection Beamforming for Deep Interference Nulling

5.1 Introduction

For decades a variety of adaptive beamforming algorithms have been used successfully to spatially cancel interfering signals [32,33], but there are yet many applications and signal scenarios where existing methods are inadequate. When interference motion is relatively fast, interference to noise power ratio (INR) is low, or when very deep nulling below the noise floor is required, conventional adaptive cancelers may perform poorly. This chapter presents an algorithm which can improve cancellation null depth by more than 10 dB under these conditions. Due to its zero-forcing deep nulling behavior, our proposed method uses the well known subspace projection (SP) algorithm [34, 63, 64, 70–72]. The SP approach is to identify the interference subspace through an eigen-decomposition of the array covariance matrix and project the total signal onto a space orthogonal to the interference.

Though the principles and proposed algorithm in this chapter are generally applicable to adaptive array interference cancellation for a wide range of sensor arrays and application areas, the discussion presented here is motivated by and focuses on phased array feeds (PAFs) for radio astronomy. This application serves as an apt illustration of the difficulties in attempting deep nulling when interference to noise ratio (INR) is low and interference motion is rapid relative to the array covariance matrix estimation time interval, called the “short-term integration” (STI) window. Non-white noise bias also arises in PAFs, limiting null depth for some algorithms.

Several major projects are currently investigating PAFs for centimeter band radio astronomy, including the Brigham Young University/NRAO array at the Green Bank Telescope (GBT) in West Virginia [20, 21, 43, 69], the NAIC AO40 feasibility study for the Arecibo telescope in Puerto Rico [45], the ASTRON APERTIF array at Westerbork Synthesis Radio Telescope in The Netherlands [27], the CSIRO Australian Square Kilometer Array Pathfinder (ASKAP) [96], the

PHAD array at Dominion Radio Astrophysical Observatory (DRAO) Composite Applications to Radio Telescopes (CART) reflectors [30], and proposed designs for the Square Kilometer Array (SKA) [9].

Radio astronomical PAFs provide an extremely challenging environment for adaptive array interference canceling, and demand high performance solutions that do not currently exist in the literature due in part to extremely low interference-to-noise ratios (INR) encountered in operational scenarios. The power spectral density of a typical signal of interest (SOI) is 30 dB or more *below* the noise floor PSD at the resolution bandwidth of the receiving system. Detecting such low level signals requires long integrations (on the order of minutes to hours) and on-source minus off-source subtraction [34]. Very good beampattern nulls must be placed to cancel an interferer to a level where it will not obscure the SOI. Even in the absence of interference, receiver electronics must be very stable since their fluctuations limit the detection of weak sources. Interference cancelation algorithms may introduce beampattern fluctuations of their own known as “pattern rumble” [68] and spectral bias known as “spectral scooping” [69]. In addition, the interferer becomes more difficult to detect as INR decreases but still obscures the SOI since the SNR may still be lower than the residual INR after canceling. A good introduction to interference mitigation in radio astronomy can be found in [34, 63, 64], and specifically SP interference mitigation in [64], which presents a low-rank sequential update algorithm to track a changing interference subspace. A related example of low-order model fitting in radio astronomy is found in [97].

As contemporary radio astronomical science goals increasingly require observing sources outside the traditional protected spectrum bands, a critical need is developing to deal with ubiquitous man-made interfering signals such as satellite downlink transmissions [49, 51], radar systems [54, 56], air navigation aids [58], wireless communications [60], and digital television broadcasts. Even locating instruments in undeveloped areas with regulatory protection does not avoid many man-made sources such as satellite downlinks.

The remainder of the chapter is organized as follows. Section 5.2 provides mathematical preliminaries, a signal model, and a detailed definition of subspace projection beamforming. Sources of motion dependent estimation error leading to limited null depth are studied. Section 5.3 proposes a matrix polynomial model for time-variation in the array covariance due to interference motion. Two algorithms with differing computational requirements are introduced to estimate

polynomial model coefficients. These algorithms address covariance estimation error issues described in the previous section. Section 5.3.5 presents a Cramer Rao lower bound analysis for error variance in the polynomial coefficients. Simulated and real data experimental results are presented in Section 5.4. It is shown that both of the proposed algorithms produce deeper cancellation nulls than other available algorithms, including a comparison with PASTd subspace tracking. Final comments and observations are found in Section 5.5.

5.2 Beamforming Issues for Moving Interference

5.2.1 Signal Model

Let $\mathbf{x}[n]$ be the $P \times 1$ vector of complex sensor array samples. For the PAF experiments reported in Section 5.4 this corresponds to the $P = 19$ separate basebanded antenna outputs.

The array sample is modeled as

$$\mathbf{x}[n] = \mathbf{a}_s s[n] + \mathbf{a}_d[n] d[n] + \boldsymbol{\eta}[n], \quad (5.1)$$

where $s[n]$ is the SOI with array response vector \mathbf{a}_s (including effects of the reflector dish, PAF array, and mutual coupling between antenna elements), $d[n]$ is the single interfering source with corresponding array response $\mathbf{a}_d[n]$. The dish is physically steered to keep the SOI at boresight, so \mathbf{a}_s is constant. As the interferer moves or as the dish tracks the SOI, $\mathbf{a}_d[n]$ changes over time. Noise vector $\boldsymbol{\eta}[n]$ is assumed to be spatially non-white, Gaussian, with correlation arising from PAF element mutual coupling and from the non-isotropic spillover ground noise field seen by the PAF beyond the edge of the dish. Signals $d[n]$, $s[n]$, and $\boldsymbol{\eta}[n]$ are assumed to be temporally wide sense stationary Gaussian random processes.

The array covariance matrix is

$$\begin{aligned} \mathbf{R}_x[n] &= E\{\mathbf{x}[n]\mathbf{x}^H[n]\} \\ &= \sigma_s^2 \mathbf{a}_s \mathbf{a}_s^H + \sigma_d^2 \mathbf{a}_d[n] \mathbf{a}_d^H[n] + \mathbf{R}_\eta \\ &= \mathbf{R}_s + \mathbf{R}_d[n] + \mathbf{R}_\eta, \end{aligned} \quad (5.2)$$

where the time-dependence in $\mathbf{R}_x[n]$ is due to the changing array response vector of the moving interferer, so it is only approximately wide sense stationary over an STI. This assumption fails if the motion rate is fast relative to STI window length. All other signals are assumed to be wide sense stationary.

Though $\mathbf{R}_x[n]$ varies continuously with n , it must be estimated by assuming stationarity during the k^{th} short-term integration (STI) window of length L samples

$$\hat{\mathbf{R}}_{x,k} = \frac{1}{L} \sum_{n=kL}^{(k+1)L-1} \mathbf{x}[n]\mathbf{x}^H[n] = \frac{1}{L} \mathbf{X}_k \mathbf{X}_k^H,$$

where $\mathbf{X}_k = [\mathbf{x}[kL], \dots, \mathbf{x}[(k+1)L-1]]$ is the L sample long data matrix for the k th STI.

Due to reflector dish focal effects, interference power levels at the array usually vary widely across elements and depend on signal arrival angle. Therefore, in the sequel, interference to noise ratio (INR) is defined as the temporal and spatial (across array elements) average interference power over the spatial average noise power.

5.2.2 Subspace Projection Interference Canceling

In all PAF radio astronomical instruments currently under development, the array is used in conjunction with the dish to simultaneously form multiple narrow beams, each steered to a different spot on the sky (see Fig. 1.1). This system, sometimes referred to as a “radio camera” permits high gain and high resolution coverage of a wider field of view than would a single horn feed, without having to scan the look direction by physically steering the dish. Sky survey speeds can thus be dramatically improved.

Assuming narrowband operation, the sampled data stream at the beamformer output can be expressed as

$$y[n] = \mathbf{w}^H \mathbf{x}[n], \quad (5.3)$$

where \mathbf{w} is the vector of beamformer weights designed to provide the desired directional response pattern. If signals of interest are broadband relative to the inverse of propagation time across the array, then $\mathbf{x}[n]$ is decomposed into spectral subbands, each with separate beamformer weights.

There are a number of candidate classical adaptive beamforming algorithms which can be used to design \mathbf{w} to steer the mainlobe as desired, control pattern shape, and place spatial nulls in direction of interfering signals [32,33]. The proposed method is based on subspace projection (SP) because its zero forcing properties can produce very deep nulls [34, 63, 64]. When interference is strong enough to dominate both noise and SOI components in $\mathbf{R}_x[n]$, beamformer weights can be computed as

$$\begin{aligned}\mathbf{w}_{\text{SP},k} &= \mathbf{P}_{\text{d},k}^\perp \mathbf{w}, \text{ where} \\ \mathbf{P}_{\text{d},k}^\perp &= \mathbf{I} - \mathbf{U}_{\text{d},k} \mathbf{U}_{\text{d},k}^H,\end{aligned}\tag{5.4}$$

and where k is the STI index, $\mathbf{P}_{\text{d},k}^\perp$ is the perpendicular projection matrix for the estimated interference subspace, and $\mathbf{U}_{\text{d},k}$ contains the dominant eigenvector(s) corresponding to interference from the partitioned eigen problem

$$\hat{\mathbf{R}}_{\mathbf{x},k}[\mathbf{U}_{\text{d},k} \mid \mathbf{U}_{\text{s},\eta,k}] = [\mathbf{U}_{\text{d},k} \mid \mathbf{U}_{\text{s},\eta,k}] \mathbf{\Lambda}.\tag{5.5}$$

Nominal weight vector \mathbf{w} is selected (e.g. by deterministic windowed beamformer design) to establish the quiescent beampattern in the absence of interference. The projection $\mathbf{P}_{\text{d},k}^\perp \mathbf{w}$ will perturb and can distort the desired quiescent beampattern, so we require that the interference direction is in the sidelobe region of the beampattern where precise response specifications other than maximum levels are not needed. Also, the rank of $\mathbf{U}_{\text{d},k}$ (i.e. the number of columns, or independent interferers) must be small relative to P or insufficient degrees of freedom remain to form the desired quiescent pattern away from the spatial nulls. The algorithm presented in Section 5.3 addresses the single interferer case, where $\mathbf{U}_{\text{d},k}$ is a column vector.

In a moving interference environment $\mathbf{w}_{\text{SP},k}$ is recomputed for each new L sample long STI to track and keep the cancelation null centered on the interferer. For a single interferer, $\mathbf{U}_{\text{d},k}$ is a $P \times 1$ vector, which is the case that is analyzed in the sequel. For multiple interferers the number of sources, m , must be estimated, using for example the minimum description length criterion [98], and $\mathbf{U}_{\text{d},k}$ is $P \times m$.

5.2.3 Factors Limiting Cancellation Performance

When SP and other block updated adaptive beamformers are used in a moving interference environment, there are several factors limiting accurate identification of the interference subspace from sample covariance $\hat{\mathbf{R}}_{\mathbf{x},k}$. This in turn inhibits cancellation performance. This chapter considers sample estimation error, subspace smearing, noise bias and spectral scooping. The first two issues are competing estimation error terms which depend on STI length. Noise bias is due to correlated noise in the system, which is typical for PAFs. Spectral scooping [69], affects cancellation for narrowband interferers including FM-modulated interferers that appear narrowband on short timescales. The proposed algorithm presented in Section 5.3 mitigates each of these detrimental effects.

Sample Estimation Error

In a stationary signal environment (i.e. non-moving interference) $\hat{\mathbf{R}}_{\mathbf{x},k}$ is a consistent estimator for $\mathbf{R}_{\mathbf{x}}[n]|_{n=kL}$ and estimation error variance decreases as $\frac{1}{L}$ [78]. To reduce estimator variance, STIs should be as long as possible, but this must be balanced against subspace smearing error due to long STIs for a moving interferer as discussed next.

Figure 5.1 presents two complementary metrics of moving interference cancellation performance: interference rejection ratio (IRR) and residual power, as functions of STI length and interference to noise ratio (INR). The solid lines and left y-axis scale show IRR obtained with SP canceling for a detailed simulation of an interferer seen by a 19 element PAF on a 20m telescope dish. IRR is defined as the INR before beamforming divided by the INR after beamforming. Details of simulation methods and a description of the signal scenario are given in Section 5.4.1.

For a given STI length, a high-INR interferer yields higher IRR because the dominant eigenvector of $\hat{\mathbf{R}}_{\mathbf{x},k}$ is then a better estimate of the true \mathbf{a}_d . For increasing STI length, IRR increases as L until the STI becomes so long that the interferer motion begins to degrade cancellation. This produces an IRR maximum, which can be seen for example at STI length of 256 samples for the 20dB INR curve.

Subspace Smearing

$\hat{\mathbf{R}}_{x,k}$ involves a time average of time varying $\mathbf{R}_d[n]$ over L samples of the k^{th} STI. Although $\mathbf{R}_d[n] = \sigma_d^2 \mathbf{a}_d[n] \mathbf{a}_d^H[n]$ is rank-one, time averaging during motion produces a higher rank matrix, requiring more eigenvectors to span the subspace. Only when STIs are short relative to motion rate can $\mathbf{P}_{d,k}^\perp$ effectively project out this higher induced rank interference component. This finite window estimation effect dubbed “subspace smearing” becomes more pronounced as STI length increases, as seen in Fig. 5.1 where it limits the achievable IRR and causes residual power $P_{\text{int,BF}}$ to rise with increasing STI length.




Figure 5.1: IRR and residual interference power dependence on STI length and INR level for a moving interferer. Solid curves show IRR values as indicated on the left y-axis scale. Dashed curves are residual interference power as indicated on the right y-axis scale.

Noise Bias

For the case of spatially white noise and no SOI, $\mathbf{R}_{s,\eta} = \sigma_\eta^2 \mathbf{I}$, and since adding a scaled identity matrix to any eigen problem does not change the eigenvectors, $\mathbf{U}_{d,k}$ computed in (5.5) is an unbiased estimator for the interference subspace containing \mathbf{a}_s in (5.1). Unfortunately, in astronomical PAFs and other applications, noise is typically colored. Thermal noise in the first low

noise amplifier (LNA) input stages propagates backward into antenna elements and is broadcast to neighboring array elements through mutual coupling in the closely packed array. This introduces a strong noise correlation. Spillover noise from the warm ground seen by the PAF beyond the dish edges is also correlated since this warm noise field is not spatially isotropic due to dish blockage and the ground–sky horizon being off-centered from the dish axis. Sky noise due to observing through warm atmosphere is also spatially correlated, particularly since it is seen through the high gain dish beam pattern. With this correlated noise, $\mathbf{R}_{s,\eta}$ is not diagonal and $\mathbf{U}_{d,k}$ is biased away from the interference subspace toward the dominant noise eigenvector. For high-INR interferers this effect is small, but as INR decreases subspace estimation suffers and $\mathbf{P}_{d,k}^\perp$ fails to block interference.

Since \mathbf{R}_s is small, approximate prewhitening can be accomplished using a previously recorded noise-only calibration data set to separately estimate $\hat{\mathbf{R}}_\eta$, which is assumed to be full rank. The whitened array covariance estimate is formed by using the factorization $\hat{\mathbf{R}}_\eta = \hat{\mathbf{R}}_\eta^{\frac{1}{2}} \hat{\mathbf{R}}_\eta^{\frac{1}{2}}$ where $\hat{\mathbf{R}}_\eta^{\frac{1}{2}}$ is required to be Hermitian, then

$$\begin{aligned} \check{\mathbf{R}}_{x,k} &= \hat{\mathbf{R}}_\eta^{-\frac{1}{2}} \hat{\mathbf{R}}_{x,k} \hat{\mathbf{R}}_\eta^{-\frac{1}{2}} \\ &\approx \mathbf{R}_\eta^{-\frac{1}{2}} \mathbf{R}_d[kL] \mathbf{R}_\eta^{-\frac{1}{2}} + \mathbf{I}. \end{aligned} \quad (5.6)$$

The resulting projection matrix is given by

$$\begin{aligned} \check{\mathbf{P}}_{d,k} &= \hat{\mathbf{R}}_\eta^{-\frac{1}{2}} (\mathbf{I} - \check{\mathbf{U}}_d \check{\mathbf{U}}_d^H) \hat{\mathbf{R}}_\eta^{+\frac{1}{2}} \\ \check{\mathbf{R}}_{x,k}[\check{\mathbf{U}}_{d,k} | \check{\mathbf{U}}_{s,\eta,k}] &= [\check{\mathbf{U}}_{d,k} | \check{\mathbf{U}}_{s,\eta,k}] \check{\mathbf{\Lambda}}. \end{aligned} \quad (5.7)$$

When $\check{\mathbf{P}}_{d,k}$ is used directly in (5.4) in place of $\mathbf{P}_{d,k}^\perp$, it prewhitens the noise in $\mathbf{x}[n]$, performs the interference perpendicular projection in that whitened space, then transforms the data back into normal space so that we can achieve the designed quiescent beam pattern.

Figure 5.2 illustrates the separate contributions of sample estimation error, subspace smearing, and noise bias on total residual interference power. The top two curves show INR after beamforming for SP, with and without prewhitening, for a -9.8dB INR interferer moving at 3°/sec through the dish sidelobes. Dashed lines are for a motionless interferer (with and without whiten-




Figure 5.2: Limiting factors in SP. For short STIs, sample estimation error is the dominant factor hindering SP; for longer STIs, subspace smearing dominates. In the case of no interferer motion, longer averaging can reduce sample estimation error only to the noise bias limit. Prewhitening lowers that limit so that for a stationary interferer, sample estimation error can be reduced arbitrarily with longer STIs.

ing) so there is no subspace smearing and null depth is limited only by sample estimation error. Without whitening, output INR improves with increasing STI length until it reaches a limit set by the noise bias. With prewhitening interference rejection continues to improve beyond the noise bias floor as L increases. The subspace smearing curve shows the residual power for the case where there is no sample estimation error and null depth is only limited due to averaged interference subspaces. These are total power INR curves where spectral scooping is not apparent, so no compensation for it is done here.

Spectral Scooping

Spectral scooping is a recently discovered phenomenon affecting narrowband interferers that severely biases the PSD estimate for shorter STI windows. A tracking interference canceling beamformer based on successive block estimates of the sample covariance will “notch out” or attenuate the signal of interest spectrum in the vicinity of the interference frequency. This effect is described in [69], along with a technique for mitigating it.

Figures 5.3(a) and 5.3(b) illustrate the PSD bias introduced by spectral scooping. The figures compare the PSD for a moving interference source entering through the dish sidelobes using a MaxSNR beamformer and SP applied to the same MaxSNR quiescent beamformer. The SP projection matrices are updated with STIs of 64 samples and 4 samples. For a CW tone interferer the spatial nulling couples into a temporal nulling that drives a null in the PSD. This behavior becomes more pronounced for short STIs. In Figure 5.3(b), SP has produced a very wide dip in the PSD and the overall power level is low even far from the interferer frequency, while the tone can still be seen on top of the spectral scoop. The algorithm presented in Section 5.3 does not suffer from spectral scooping. To make a fair comparison between that algorithm and the conventional SP approach in the real data experiments of Section 5.4.3, spectral scooping had to be mitigated in the conventional SP curves. The technique from [69] for overcoming spectral scooping is to form two projection matrices from each STI window, each based on half of the window, apply each projection matrix to the entire STI window, and then average the two resulting PSDs.




Figure 5.3: Spectral scooping illustration using a moving CW tone interferer in the dish sidelobes for a MaxSNR beamformer and SP applied using the same MaxSNR beamformer. (a): With SP computed using 64-sample STIs there is a sharp notch placed in the PSD. (b): Spectral scooping becomes more pronounced for shorter STIs; here SP is computed using 4-sample STIs resulting in a wide dip in the PSD and lower overall power. The CW tone can also be seen on top of the spectral dip.

5.3 A Polynomial Covariance Motion Model for Subspace Projection

5.3.1 Background

The primary goal of this work is to develop an adaptive beamforming method capable of placing deeper nulls than existing algorithms to cancel low INR moving interference. This addresses important but difficult radio astronomical observation scenarios for telescopes using PAFs. These are demanding requirements, and it is expected that in weighing candidate solutions, some trade offs will need to be made, for example requiring increases in complexity or being unable to support real-time processing. The critical factor in achieving deep nulls is reducing error in the instantaneous estimates of the interference subspace obtained through a time sequence of sample covariances. The foregoing analysis of factors limiting cancelation performance motivates the following set of criteria to guide algorithm and model selection such that the tightly linked trade off between sample estimation error and subspace motion smearing can be broken:

1. Long data windows are needed for extended integration times to reduce sample estimation error. Estimates based on independent short STI integrations are inadequate, so the window must span significant motion-induced time variation.
2. An interference signal model is required which explicitly incorporates continuous, smooth variation to represent time evolution of the array covariance matrix, rather than using a block-wise-constant short term stationarity assumption.
3. The model order, or parametric degrees of freedom, must be low in order to exploit noise reduction through integration across the longer window.

Deeper beampattern nulls are obtained by incorporating all of the available data in the model instead of considering each block STI independently. Choosing the appropriate STI length a priori to achieve optimal null depth is very difficult in practice, and as will be shown below, even optimal STI length in the SP algorithm does not achieve the deepest possible null.

A number of model options have been considered, including subspace tracking [64], which as shown in Section 5.4.2 only weakly satisfies criterion 1); spline fitting, which though promising was not the best match to criterion 3); and physics-based dynamic modeling of motion, which does

not provide a good solution to criterion 2). This is because (at least in part) the dominant covariance variation arises not so much from motion dynamics, but from the modulation induced by the fine scale, random deep sidelobe beampattern structure of a large dish PAF telescope. The method presented below addresses all three criteria and has been demonstrated to produce exceptionally deep nulls.

This chapter proposes an extension of the SP beamforming approach by incorporating a matrix polynomial model for time variation of $\mathbf{R}_d[n]$. This new Polynomial-augmented Subspace Projection (PSP) method improves interference subspace tracking. Interferer dynamics must follow physical laws of motion, so changes from $\mathbf{R}_{x,k}$ to $\mathbf{R}_{x,k+1}$ due to physical motion happen in a smooth way. This suggests that a good fit to a polynomial of appropriate order should be possible. Bias induced by the SOI is assumed to be negligible. This is a common radio astronomical observing scenario, where the SOI is so weak that it can only be detected after long integration times (order of minutes or hours) using on-source minus off-source differencing. Also assume that the temporal and spectral properties of $d[n]$ (i.e. its modulation) are such that there is sufficient in-band energy in each STI window to a) be detrimental to SOI observation, and b) to enable an interference subspace estimate in that STI. This would include all conventional fixed carrier sources, and many frequency hopped, spread spectrum, or pulsed modulation as long as the chip rates and frequency spread lead to significant energy per STI for the beamformer.

Polynomial coefficients are computed by fitting to the STI-based $\hat{\mathbf{R}}_{x,k}$. But, once fit, this model can be evaluated on much finer time scales so projection matrices and SP beamformer weights can be computed virtually continuously to reduce subspace smearing. Sample estimation error is also reduced because the polynomial fit utilizes all samples in the many STIs of the fitting window while the number of parameters estimated in the fit is far fewer than for separate $\hat{\mathbf{R}}_{x,k}$ estimates per STI. For best algorithm performance, it is assumed the noise is spatially white, or has been whitened using (5.6) and (5.7), based on a noise covariance estimate $\hat{\mathbf{R}}_\eta$ obtained previously from a calibration data set.

5.3.2 Problem Formulation

Consider the underlying continuous-time covariance matrix $\mathbf{R}_x(t)$ from which the sampled version $\mathbf{R}_x[n]$ is drawn. We introduce $\tilde{\mathbf{R}}_d(t, \mathbf{C}) \approx \mathbf{R}_d(t)$ as a model for the time-varying

interference component of $\mathbf{R}_x(t)$, parameterized by polynomial coefficient matrix \mathbf{C} . For a single interferer $\mathbf{R}_d(t) = \sigma_d^2 \mathbf{a}_d(t) \mathbf{a}_d^H(t)$ is rank 1, so the time progression of $\sigma_d \mathbf{a}_d(t)$ is modeled using the r th order vector polynomial

$$\tilde{\mathbf{a}}(t, \mathbf{C}) = \mathbf{C} \mathbf{t}, \text{ where} \quad (5.8)$$

$$\mathbf{t} = [1, t, \dots, t^r]^T, \text{ so}$$

$$\tilde{\mathbf{R}}_d(t, \mathbf{C}) = \tilde{\mathbf{a}}(t, \mathbf{C}) \tilde{\mathbf{a}}^H(t, \mathbf{C}) = \mathbf{C} \mathbf{t} \mathbf{t}^H \mathbf{C}^H. \quad (5.9)$$

Estimated coefficients $\hat{\mathbf{C}}$ are found by squared error minimization with respect to the STI sample covariances

$$\hat{\mathbf{C}} = \arg \min_{\mathbf{C}} \sum_{k=0}^{K-1} \|\hat{\mathbf{R}}_{x,k} - \tilde{\mathbf{R}}_d(t_k, \mathbf{C})\|_F^2 \quad (5.10)$$

$$t_k = T \left(\frac{2k}{K-1} - 1 \right), \quad (5.11)$$

where t_k is the normalized time value corresponding to the midpoint sample within the k th STI, and (5.11) ensures that $-T \leq t_k \leq T$ for $0 \leq k \leq K-1$, spanning the K STIs in the polynomial fitting window. Timescale normalization for polynomial evaluation is arbitrary in that T does not correspond to any real-world time or window width in seconds or samples. T is chosen as a tuning parameter to limit the time range when evaluating (5.9) so high order polynomial terms do not dominate curve shape and multiple bounded variations can be represented in the fitting window. Typically the value $T = 0.45$ has been used for good fitting performance.

5.3.3 Solution by Numerical Optimization

Substituting (5.9) into (5.10) using the identity $\|\mathbf{A}\|_F^2 = \text{tr}\{\mathbf{A}^H \mathbf{A}\}$ and expanding the matrix products yields

$$\begin{aligned} \hat{\mathbf{C}} &= \arg \min_{\mathbf{C}} \sum_{k=1}^K f_k \quad (5.12) \\ f_k &= \text{tr}\{\hat{\mathbf{R}}_k^H \hat{\mathbf{R}}_k + \mathbf{C} \mathbf{B}_k \mathbf{C}^H \mathbf{C} \mathbf{B}_k \mathbf{C}^H \\ &\quad - \hat{\mathbf{R}}_k \mathbf{C} \mathbf{B}_k \mathbf{C}^H - \mathbf{C} \mathbf{B}_k \mathbf{C}^H \hat{\mathbf{R}}_k^H\}, \end{aligned}$$

where $\mathbf{B}_k = \mathbf{t}_k \mathbf{t}_k^H$. This form of (5.10) is needed for computing analytical gradients and Hessians used to reduce the significant computer memory usage and improve convergence times in a numerical approach to the optimization.

A variety of widely available optimization codes are suitable for solving (5.12). Results in Section 5.4 were found using the MATLAB® function *fminunc.m* which is a subspace trust-region method based on the interior-reflective Newton method described in [99] and [100]. The gradient and Hessian of (5.12) are used in defining the trust-region problem [101] and restricting the region to a two dimensional subspace [102, 103]. Analytical expressions for both of these terms are included below, and were used by the numerical optimizer. Other algorithms exist that make use of the Jacobian, also included here.

Gradient:

Polynomial coefficients are complex-valued but in general-purpose optimizer codes they must be treated as separate sets of real and imaginary coefficients. In this context the gradient becomes:

$$\nabla_{\mathbf{C}} f = \begin{bmatrix} \text{vec}\left\{\frac{\partial f}{\partial \text{Re}\{\mathbf{C}\}}\right\} \\ \text{vec}\left\{\frac{\partial f}{\partial \text{Im}\{\mathbf{C}\}}\right\} \end{bmatrix}$$

$$\nabla_{\mathbf{C}} \left(\sum_{k=1}^K f_k \right) = 4 \sum_{k=1}^K \begin{bmatrix} \text{vec}\left\{\text{Re}\left\{(\mathbf{C}\mathbf{B}_k\mathbf{C}^H - \hat{\mathbf{R}}_k)\mathbf{C}\mathbf{B}_k\right\}\right\} \\ \text{vec}\left\{\text{Im}\left\{(\mathbf{C}\mathbf{B}_k\mathbf{C}^H - \hat{\mathbf{R}}_k)\mathbf{C}\mathbf{B}_k\right\}\right\} \end{bmatrix}.$$

Hessian:

Using \mathbf{E}_{st} to denote an elementary matrix consisting of all zeros, with a “1” in the $(s, t)^{th}$ location, the Hessian can be expressed as:

$$\mathbf{H}_{\mathbf{C}} \left(\sum_{k=1}^K f_k \right) = 4 \sum_{k=1}^K \begin{bmatrix} \text{vec}\{\text{Re}\{\mathbf{M}_1\}\} & \text{vec}\{-\text{Im}\{\mathbf{M}_2\}\} \\ \text{vec}\{\text{Im}\{\mathbf{M}_1\}\} & \text{vec}\{\text{Re}\{\mathbf{M}_2\}\} \end{bmatrix},$$

where

$$\begin{aligned} \mathbf{M}_1 &= \mathbf{E}_{st} \mathbf{B}_k \mathbf{C}^H \mathbf{C} \mathbf{B}_k + \mathbf{C} \mathbf{B}_k \mathbf{E}_{st}^H \mathbf{C} \mathbf{B}_k \\ &\quad + \mathbf{C} \mathbf{B}_k \mathbf{C}^H \mathbf{E}_{st} \mathbf{B}_k - \hat{\mathbf{R}}_k \mathbf{E}_{st} \mathbf{B}_k \end{aligned}$$

and

$$\begin{aligned} \mathbf{M}_2 &= \mathbf{E}_{st} \mathbf{B}_k \mathbf{C}^H \mathbf{C} \mathbf{B}_k - \mathbf{C} \mathbf{B}_k \mathbf{E}_{st}^H \mathbf{C} \mathbf{B}_k \\ &\quad + \mathbf{C} \mathbf{B}_k \mathbf{C}^H \mathbf{E}_{st} \mathbf{B}_k - \hat{\mathbf{R}}_k \mathbf{E}_{st} \mathbf{B}_k. \end{aligned}$$

Jacobian:

$$\mathbf{J}_{\mathbf{C}}(\mathbf{f}) = \begin{bmatrix} \mathbf{Y}_{111} & \mathbf{Y}_{P11} & \mathbf{Y}_{1r1} & \mathbf{Y}_{Pr1} & \mathbf{W}_{111} & \mathbf{W}_{P11} & \mathbf{W}_{1r1} & \mathbf{W}_{Pr1} \\ \mathbf{V}_{111} & \mathbf{V}_{P11} & \mathbf{V}_{1r1} & \mathbf{V}_{Pr1} & \mathbf{X}_{111} & \mathbf{X}_{P11} & \mathbf{X}_{1r1} & \mathbf{X}_{Pr1} \\ \vdots & \cdots & \vdots & \cdots & \vdots & \cdots & \vdots & \cdots & \vdots \\ \mathbf{Y}_{11K} & \mathbf{Y}_{P1K} & \mathbf{Y}_{1rK} & \mathbf{Y}_{PrK} & \mathbf{W}_{11K} & \mathbf{W}_{P1K} & \mathbf{W}_{1rK} & \mathbf{W}_{PrK} \\ \mathbf{V}_{11K} & \mathbf{V}_{P1K} & \mathbf{V}_{1rK} & \mathbf{V}_{PrK} & \mathbf{X}_{11K} & \mathbf{X}_{P1K} & \mathbf{X}_{1rK} & \mathbf{X}_{PrK} \end{bmatrix}$$

where $\mathbf{Y}_{stk} = -(\mathbf{E}_{st} \mathbf{B}_k \text{Re}\{\mathbf{C}^T\} + \text{Re}\{\mathbf{C}\} \mathbf{B}_k \mathbf{E}_{st}^T)$, $\mathbf{V}_{stk} = \mathbf{E}_{st} \mathbf{B}_k \text{Im}\{\mathbf{C}^T\} - \text{Im}\{\mathbf{C}\} \mathbf{B}_k \mathbf{E}_{st}^T$, $\mathbf{W}_{stk} = -(\mathbf{E}_{st} \mathbf{B}_k \text{Im}\{\mathbf{C}^T\} + \text{Im}\{\mathbf{C}\} \mathbf{B}_k \mathbf{E}_{st}^T)$, and $\mathbf{X}_{stk} = -\mathbf{E}_{st} \mathbf{B}_k \text{Re}\{\mathbf{C}^T\} + \text{Re}\{\mathbf{C}\} \mathbf{B}_k \mathbf{E}_{st}^T$.

Note that direct optimization of (5.12) contains a single phase ambiguity which inhibits numerical convergence since $(e^{j\phi} \tilde{\mathbf{a}}(t_k, \mathbf{C})) (e^{j\phi} \tilde{\mathbf{a}}(t_k, \mathbf{C}))^H = \tilde{\mathbf{a}}(t_k, \mathbf{C}) \tilde{\mathbf{a}}^H(t_k, \mathbf{C})$ for any ϕ . This applies to only one arbitrary time instance t_k while bulk phases at all remaining t_l , $l \neq k$ are resolved due to continuity of the polynomial. This observation is confirmed by the CRB analysis of Section 5.3.5. The ambiguity is resolved by arbitrarily enforcing $\angle\{\hat{\mathbf{C}}\}_{1,1} = 0$ so $\angle\{\tilde{\mathbf{a}}(0, \hat{\mathbf{C}})\}_1 = 0$. An alternate constraint parameterization suggested by [104] is forcing $\sum_{p=0}^{P-1} \angle\{\hat{\mathbf{C}}\}_{p,1} = 0$, that is to have a zero average phase across elements of $\tilde{\mathbf{a}}(0, \hat{\mathbf{C}})$. Though this was shown to reduce absolute phase estimation error variance by a factor of two, both constraint choices produce equal error variance in systems that depend only on phase differences. This is precisely the case for how $\tilde{\mathbf{a}}(t, \hat{\mathbf{C}})$ is used in (5.9) to form $\tilde{\mathbf{R}}_d(t, \hat{\mathbf{C}})$, so due to its simpler implementation in the optimizer code, we choose to require $\angle\{\hat{\mathbf{C}}\}_{1,1} = 0$.

5.3.4 Optimizer Initialization and Approximate Solution

Since (5.12) is not in general a convex optimization problem it may have multiple local minima. Performance in both accuracy and convergence time for numerical solutions is significantly improved if optimizer iterations are initialized with a good starting point near the globally optimal solution. We introduce a computationally efficient algorithm which approximates $\hat{\mathbf{C}}$ for a single interferer using a direct least squares fit to the sequence of dominant eigenvectors \mathbf{u}_k and corresponding eigenvalues λ_k of $\hat{\mathbf{R}}_{\mathbf{x},k}$ over the K STIs. The progression of this eigenvector sequence approximates the time dependent span of the varying interference subspace, but the polynomial fit generally has somewhat more error than the optimized solution of Section 5.3.3, as evidenced by higher residual interference power seen in Fig. 5.7. Even considering these errors, the approach is also useful as a computationally simple final result in higher INR scenarios when interference rejection provided by a good approximate polynomial fit is adequate.

The Eigenvector Polynomial-augmented Subspace Projection (EPSP) approximation is


$$\hat{\mathbf{C}}_{\text{eig}} = \arg \min_{\mathbf{C}} \sum_{k=1}^K \|\sqrt{\lambda_k} \mathbf{u}_k - \mathbf{C} \mathbf{t}_k\|^2, \quad (5.13)$$

which has closed-form solution

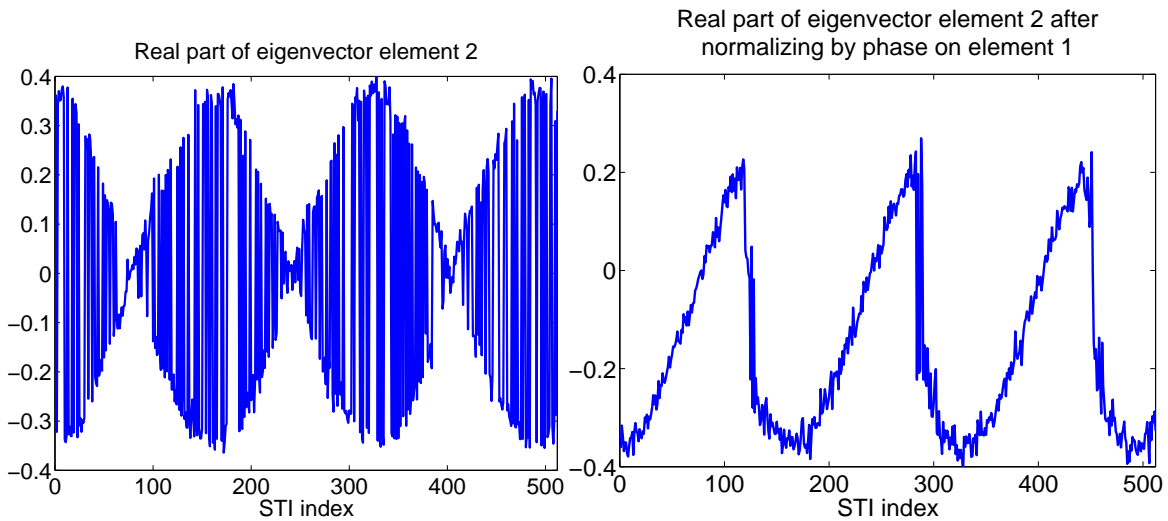
$$\hat{\mathbf{C}}_{\text{eig}} = \mathbf{U} \mathbf{\Lambda} \mathbf{T}^T (\mathbf{T} \mathbf{T}^T)^{-1}, \quad (5.14)$$

where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_K]$, each \mathbf{u}_k is normalized to unit length, $\mathbf{\Lambda} = \text{Diag}\{[\sqrt{\lambda_1}, \dots, \sqrt{\lambda_K}]\}$ is the corresponding matrix of eigenvalue square roots, and $\mathbf{T} = [\mathbf{t}_1, \dots, \mathbf{t}_K]$. When $K > r$, (5.13) is well posed for a given \mathbf{U} , but $\hat{\mathbf{C}}_{\text{eig}}$ is not unique with respect to the sequence of $\hat{\mathbf{R}}_{\mathbf{x},k}$, $1 \leq k \leq K$, because each \mathbf{u}_k is determined only to within a possibly distinct arbitrary complex constant α_k . Unit normalization of \mathbf{u}_k and scaling by $\sqrt{\lambda_k}$ resolves the magnitude ambiguity so as to match power levels with the interference subspace of $\hat{\mathbf{R}}_{\mathbf{x},k}$, but there remains an essential order K real ambiguity in the time varying eigenvector subspace matching problem. The ambiguous phase per \mathbf{u}_k must be resolved in the polynomial fit to produce a physically meaningful result.

Since for physical motion $\mathbf{a}_d[n]$ changes continuously over time, and little between STIs, the proper eigenvector phase correction should produce the “smoothest” transition from one vector



(a) True array response (smooth) and eigenvectors normalized with α_k (noisy).



(b) Un-normalized eigenvectors

(c) Eigenvectors normalized by phase on element 1

Figure 5.4: Phase normalization examples. The time progression of the real part of element two from the array signal or corresponding eigenvectors is plotted. (a): True array response (smooth) and eigenvectors normalized to α_k (noisy). Two separate axes are plotted since the curves have different scales in time and amplitude, but the close shape match suggests subspace tracking for interference cancellation using the estimate will be successful. (b): Dominant eigenvector for each of the K sample covariances, un-normalized. Rapid variation makes polynomial fitting impossible. (c): Dominant eigenvectors, after normalizing by the phase on element 1 of each eigenvector. This is relatively continuous, but requires a much higher order polynomial fit than the noisy curve in (a).

to the next. Our approach is to phase normalize each successive eigenvector for closest vector match to its predecessor. Initialize by making the first element of the first eigenvector real, so $\hat{\mathbf{u}}_1 = (e^{-j\angle \mathbf{u}_{1,1}})\mathbf{u}_1$, recall that all \mathbf{u}_k are normalized to unit length, then let

$$\begin{aligned}\hat{\mathbf{u}}_k &= e^{j\phi_k}\mathbf{u}_k, \quad k = 2, \dots, K \text{ where} \\ \phi_k &= \arg \min_{\phi} \|e^{j\phi}\mathbf{u}_k - \hat{\mathbf{u}}_{k-1}\|^2 = \angle\{\mathbf{u}_k^H \hat{\mathbf{u}}_{k-1}\}.\end{aligned}\tag{5.15}$$

These $\hat{\mathbf{u}}_k$ are used in (5.14).

Figure 5.4 illustrates the necessity of this phase normalization for polynomial fitting to the dominant eigenvector sequence. The smooth line in Figure 5.4(a) shows the true smooth variation over sample index n of the real part of the electric field seen at array element two, due to a simulated moving interferer. Figure 5.4(b) plots the corresponding $\text{Re}\{\{\mathbf{u}_k\}_2\}$ over STI index k , calculated from noisy observations of the signal. Note that without proper normalization this is very random and could not be fit to a low order polynomial. Figure 5.4(c) shows that simply normalizing so element one of each \mathbf{u}_k is real improves the continuity, but still produces a bad match to the physical variation. The noisy line in Figure 5.4(a) plots $\text{Re}\{\{\hat{\mathbf{u}}_k\}_2\}$, illustrating how (5.15) yields physically meaningful estimates of the true values seen at the array. Fitting to these values will lead to useful polynomial approximations for EPSP.

5.3.5 Cramer-Rao Lower Bound (CRB) Analysis

Following the development in [97], consider the L data vectors $\mathbf{x}[(k-1)L]$ through $\mathbf{x}[kL-1]$ in the k^{th} STI as random vectors drawn from a distribution with covariance \mathbf{R}_k . These vectors are stacked to form the $KP \times L$ matrix \mathbf{X} with sample covariance $\hat{\mathcal{R}} = \frac{1}{L}\mathbf{X}\mathbf{X}^H$ and true covariance $E\{\hat{\mathcal{R}}\}$. Data samples are independent with correlation structure arising from propagation delays across the array so \mathcal{R} is block diagonal with K blocks \mathbf{R}_k of size $P \times P$:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}[0] & \cdots & \mathbf{x}[L-1] \\ \vdots & & \vdots \\ \mathbf{x}[(K-1)L] & \cdots & \mathbf{x}[KL-1] \end{bmatrix}$$

$$\mathcal{R} = \begin{bmatrix} \mathbf{R}_1 & & \\ & \ddots & \\ & & \mathbf{R}_K \end{bmatrix}.$$

Let $\boldsymbol{\vartheta} = [\boldsymbol{\theta}_1^T, \boldsymbol{\theta}_2^T, \dots, \boldsymbol{\theta}_K^T]^T$ be the vector of all the parameters to be estimated over the entire data record, with $\boldsymbol{\theta}_k$ denoting parameter values corresponding to the midpoint sample time in the k th STI. The Cramer-Rao lower bound (CRB) on the variance for the elements of $\boldsymbol{\vartheta}$ is given by the diagonal elements of $\frac{1}{L}\mathcal{M}^{-1}$ where \mathcal{M} is the Fisher information matrix, which for Gaussian data can be expressed as [105] $\mathcal{M} = \mathcal{J}^H(\overline{\mathbf{R}}^{-1} \otimes \mathbf{R}^{-1})\mathcal{J}$ where the Jacobian \mathcal{J} is defined as $\mathcal{J} = \frac{\partial \text{vec}\{\mathcal{R}\}}{\partial \boldsymbol{\vartheta}^T}$. The sparse block diagonal structure for \mathcal{R} makes \mathcal{J} sparse, so \mathcal{M} has the same block diagonal structure as \mathcal{R} , with blocks

$$\mathbf{M}_k = \mathbf{J}_k^H(\overline{\mathbf{R}}_k^{-1} \otimes \mathbf{R}_k^{-1})\mathbf{J}_k,$$

where

$$\mathbf{J}_k = \left. \frac{\partial \text{vec}\{\mathbf{R}_k\}}{\partial \boldsymbol{\theta}_k^T} \right|_{\boldsymbol{\theta}_k = \boldsymbol{\theta}_k^{\text{true}}}.$$

A functional relationship can be employed, $\boldsymbol{\vartheta} = f(\boldsymbol{\rho})$, producing the matrices $\mathcal{F} = \frac{\partial \boldsymbol{\vartheta}}{\partial \boldsymbol{\rho}^T}$ and $\mathbf{F}_k = \frac{\partial \boldsymbol{\theta}_k}{\partial \boldsymbol{\rho}^T}$, leading to the constrained Fisher information matrix $\mathcal{M}_\rho = \mathcal{F}^H[\mathcal{J}^H(\overline{\mathbf{R}}^{-1} \otimes \mathbf{R}^{-1})\mathcal{J}]\mathcal{F} = \mathcal{F}^H\mathcal{M}\mathcal{F}$. The block diagonal structure of \mathcal{M} makes computation of \mathcal{M}_ρ a simple summation over the individual block matrices [97]: $\mathcal{M}_\rho = \sum_{k=1}^K \mathbf{F}_k^H \mathbf{M}_k \mathbf{F}_k$.

Let $\boldsymbol{\Lambda} = \mathbf{R}_\eta$ so that $\mathbf{R}_k = \mathbf{a}_k \mathbf{a}_k^H + \boldsymbol{\Lambda}$, where for correlated noise $\boldsymbol{\Lambda}$ is not diagonal. The array response vector \mathbf{a}_k changes in a step-wise constant manner from one STI to the next. Estimates of \mathbf{a}_k are all that are required for subspace projection, but nuisance parameter $\boldsymbol{\Lambda}$ must be included in the CRB calculation since its presence limits the accuracy of the \mathbf{a}_k estimate. The upper triangle of $\boldsymbol{\Lambda}$ is the conjugate of the lower triangle so the free parameter vector $\boldsymbol{\lambda}$ consists of the main diagonal and lower triangle elements of $\boldsymbol{\Lambda}$. For $\boldsymbol{\Lambda}$ of size $P \times P$ the number of free parameters is $N_\lambda = \sum_{i=1}^P i = \frac{P(P+1)}{2}$.

The parameters of interest in the estimation problem are $\boldsymbol{\theta}_k = [\mathbf{a}_k^T, \boldsymbol{\lambda}^T]^T$. Note that although $\boldsymbol{\lambda}$ appears in each $\boldsymbol{\theta}_k$, its effect is only counted once in the CRB as seen by the time-independence of \mathbf{F}^λ in (5.16).

$$\begin{aligned} \mathbf{J}_k &= \frac{\partial \text{vec}\{\mathbf{R}_k\}}{\partial \boldsymbol{\theta}_k^T} = \left[\frac{\partial \text{vec}\{\mathbf{R}_k\}}{\partial \mathbf{a}_k^T}, \frac{\partial \text{vec}\{\mathbf{R}_k\}}{\partial \boldsymbol{\lambda}^T} \right] \\ &= \left[\mathbf{J}_k^{[a]}, \mathbf{J}_k^{[\lambda]} \right] \end{aligned}$$

$$\begin{aligned} \mathbf{J}_k^{[a]} &= \left. \frac{\partial \text{vec}\{\mathbf{R}_k\}}{\partial \mathbf{a}_k^T} \right|_{\mathbf{a}_k^T = \text{true}} = (\bar{\mathbf{a}}_k \mathbf{1}_P^T) \circ \mathbf{I}_P \\ \mathbf{J}_k^{[\lambda]} &= \frac{\partial \text{vec}\{\boldsymbol{\Lambda}\}}{\partial \boldsymbol{\lambda}^T} = \mathbf{S}_{\text{LT}}, \end{aligned}$$

where \mathbf{S}_{LT} is a size $P^2 \times (P^2 + P)/2$ lower triangular part selection matrix formed from identity matrix \mathbf{I}_{P^2} by deleting columns which correspond to the upper triangle in $\boldsymbol{\Lambda}$; i.e. delete the k th columns of \mathbf{I}_{P^2} where $k = i + (j - 1)P, \forall \{(i, j) | 1 \leq j \leq P, i > j\}$.

Constraints can be added through a derivative chain-rule relationship [97, 106]. Assuming a polynomial model for \mathbf{a}_k leads to a functional constraint matrix \mathbf{F} :

$$\mathbf{a}_k = \mathbf{C} \mathbf{t}_k, \quad \mathbf{t}_k = [t_k^0, \dots, t_k^r], \quad \boldsymbol{\rho}^{(1)} = \text{vec}\{\mathbf{C}\}$$

$$\begin{aligned} \mathbf{F}_k^{[a](1)} &= \frac{\partial \mathbf{a}_k}{\partial \boldsymbol{\rho}^{(1)T}} = \mathbf{t}_k^T \otimes \mathbf{I}_P \\ \mathbf{F}^\lambda &= \mathbf{I}_P, \end{aligned}$$

where superscript (1) denotes the first in a series of constrained parameter vectors. $\mathbf{F}_k^{[a](1)}$ is further constrained by the choice used earlier to zero the phase on the first polynomial coefficient $c_{11} = \gamma_{11} e^{j\phi_{11}}$ so that

$$\boldsymbol{\rho}^{(2)} = [\gamma_{11}, c_{21}, c_{31}, \dots, c_{P(r+1)}]^T.$$

For complex z , $z = \gamma e^{j\phi}$, $\frac{\partial z}{\partial \gamma} = e^{j\phi}$ giving

$$\frac{\partial \boldsymbol{\rho}^{(1)}}{\partial \boldsymbol{\rho}^{(2)}} = \begin{bmatrix} e^{j\phi_{11}} & \mathbf{0}^T \\ \mathbf{0} & \mathbf{I}_{P(r+1)-1} \end{bmatrix} = \mathbf{F}^{[a](2)},$$

and \mathbf{F}_k is formed using $\mathbf{F}_k^{[a](1)}$, $\mathbf{F}^{[a](2)}$, and \mathbf{F}^λ :

$$\mathbf{F}_k = \begin{bmatrix} \mathbf{F}_k^{[a](1)} & \mathbf{F}^{[a](2)} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{F}^\lambda \end{bmatrix}. \quad (5.16)$$

Figure 5.5 shows simulation results for an 8^{th} order polynomial, comparing the CRB with sample error variance in the estimate $\hat{\mathbf{C}}_{\text{eig}}$ using the EPSP algorithm described in Section 5.3.4 for 2048 Monte Carlo trials. The plot has a stair step nature due to the parameter grouping, with all the 0^{th} -order coefficients showing up as the first 19 parameters plotted, then the 1^{st} -order coefficients, and so forth. For each coefficient order there is roughly the same CRB, with some variability among coefficients of the same order due to the interferer arriving at each array element in a slightly different way depending on the motion trajectory for the particular signal scenario.

Computing the CRB assumes the true underlying parameters are available. To generate a set of true polynomial coefficients, a physical motion scenario generated time varying array response vectors that were fit to a polynomial. These polynomial coefficients, \mathbf{C}_{true} were used to generate a series of array response vectors that were truly polynomial-based and of known order. The CRB was then calculated using \mathbf{C}_{true} and Monte Carlo trials were run to estimate the variance of the polynomial coefficient estimates using the EPSP algorithm for comparison with the CRB.

The CRB for this estimation problem depends strongly on the true underlying coefficients, as shown in Figure 5.6. In this figure, a very different set of polynomial coefficients were generated using the same true motion model but fitting that motion over a different time window [$T = 2.95$, in Eq. (5.11)] to produce \mathbf{C}_{true} . Running the CRB computation and Monte Carlo trials with this other set of polynomial coefficients produces the very different stair step CRB plot of Figure 5.6.

To produce a simple characterization of performance based on the CRB requires the additional step of analyzing the effect of model order mismatch, whereas the CRB analysis here

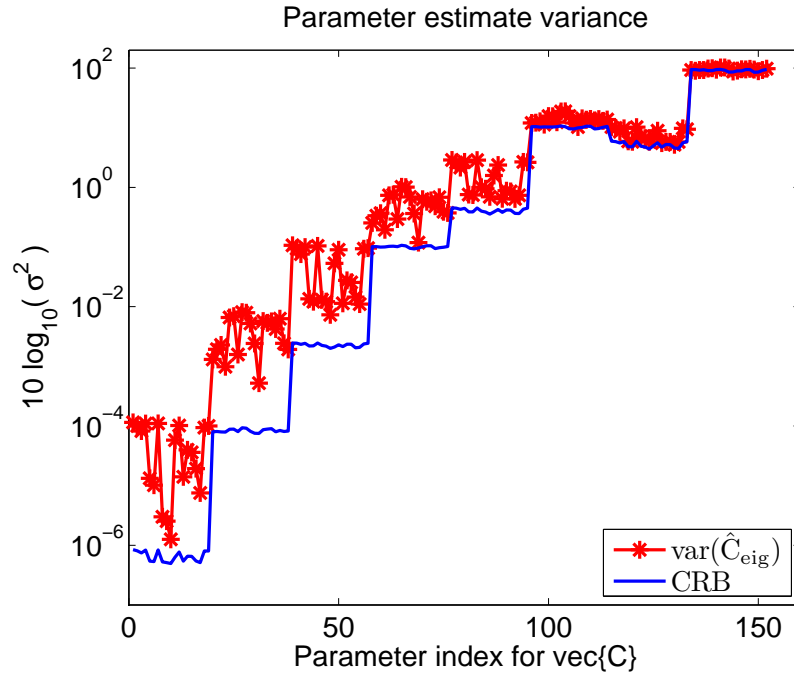


Figure 5.5: Sample error variance from Monte Carlo trials vs CRB for the EPSP algorithm using an 8th-order polynomial. Performance converges for higher-order polynomial coefficients, as expected since these parameter estimates are more strongly influenced by the noise.

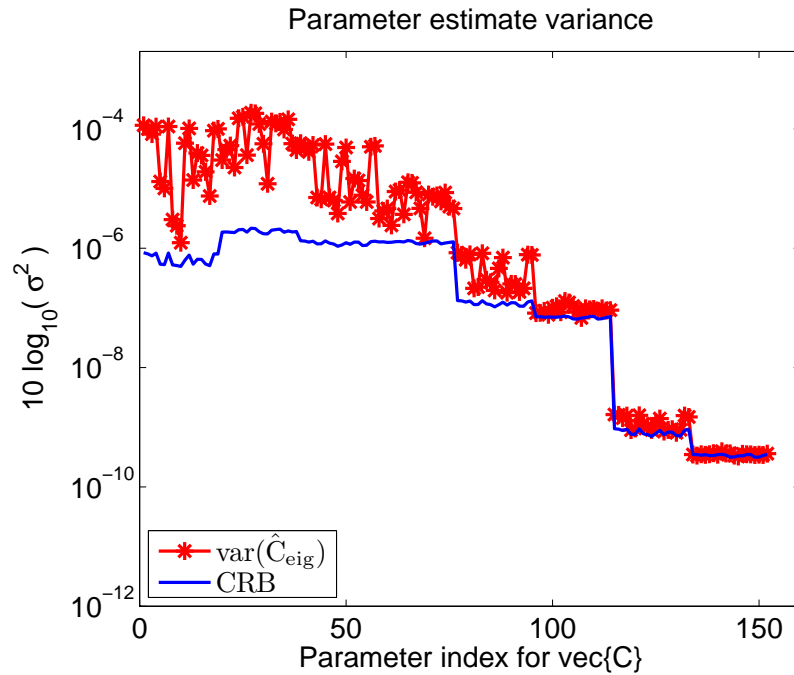


Figure 5.6: CRB and Monte Carlo trials for different set of true underlying polynomial coefficients, C_{true} . The parameter estimator variances depend strongly on the values of C_{true} , which in turn depend on the extent of the time window in Eq. (5.11).

has assumed that coefficients are estimated for a polynomial of known order. To illustrate, if the underlying motion is truly 0^{th} -order, but the EPSP algorithm is run using high polynomial order, the higher-order polynomial coefficients will all be in error and will only be nonzero due to the presence of noise. Conversely, choosing a 0^{th} -order model for motion that is truly based on a higher-order polynomial is equivalent to using conventional block-processed SP using a single STI of length matching the data record, which leads to degradation due to subspace smearing, as illustrated in Figure 5.2 for long STIs. Building a notion of model order mismatch into the CRB analysis is left as future work.

5.4 Results

5.4.1 Simulation Performance Analysis

In the following examples the simulation environment is designed to match the physical experimental platform. A detailed numerical model is used for a 19-element PAF of co-polarized linear half-wavelength broadband thickened dipoles placed over a ground plane. The 300 MHz usable antenna bandwidth is centered at 1600 MHz, which was used as the center frequency for narrowband signal simulation. Bare array element patterns were measured on an antenna test range and imported into the simulation, where a physical optics full 2D integration over the 20-meter, 0.43 f/D reflector surface gives the antenna response due to far-field sources. Blockage from the feed support structure was not included but a correlated spillover noise model and mutual coupling among antenna elements are included, using measured noise parameters of the Ciao Wireless LNAs with 33K noise temperature. The noise model is described in detail in [14]. Sample data are generated using 1.25 Msamp/s and 450kHz bandwidth.

The interference is generated using a CW tone at 0.6π discrete baseband frequency arriving with array response vectors that were calculated for a source cutting through a path covering $20 - 26^\circ$ in azimuth and $44 - 50^\circ$ in elevation. With roughly 0.7° wide sidelobes for a 20m dish, the interferer moves through three sidelobes traversing this path. Sets of STIs were calculated at windows of length 2^7 through 2^{14} samples per STI and then all STIs were prewhitened using the sample estimated calibration noise covariance matrix $\hat{\mathbf{R}}_\eta$, and polynomials were fit to each




Figure 5.7: Comparison of SP, EPSP and PSP for a 7th-order polynomial applied to the same interference scenario as in Fig. 5.2. EPSP gives 9dB improvement over best performance for SP, and over a wide range of STI lengths. In most cases PSP drives a null only 1dB deeper than EPSP, showing that the dominant eigenvector is a good model for the interference subspace in this signal scenario. For short STIs, there is enough sample estimation error in the covariance matrices that the eigenvectors are a poor estimate for the array response vectors so performance degrades when using EPSP, but PSP still produces deep nulls.

set of phase corrected eigenvectors. The quiescent (interference-free) beamformer weight vector $\mathbf{w} = \hat{\mathbf{R}}_{\eta}^{-1} \mathbf{a}_s$ used in (5.4) is a Max Sensitivity beam for an SOI kept at dish boresight.

Figure 5.7 illustrates three important aspects of using a polynomial model for the time dependence in the covariance matrix compared with the SP model that assumes independent covariance matrices for each STI: 1) SP is highly sensitive to STI length tuning but EPSP outperforms SP over a wide range of STI lengths, eliminating the need for that tuning; 2) even at the best STI length for SP, EPSP and PSP give 9dB and 10dB improvement respectively; and 3) at very short STI lengths EPSP degrades because there is too much estimation error in the covariance matrix so good set of eigenvectors cannot be found. Thus PSP is required at very short STI lengths to get the same deep nulls seen for longer STIs. Note that for STIs of about 100 samples or more, PSP gives only slight improvement over EPSP. Many of the results presented are for EPSP only, since computation time is greatly reduced (compared to PSP) while performance is nearly identical for




Figure 5.8: Performance of SP and EPSP for 8^{th} and 12^{th} order polynomials for various input INRs and motion rates. On the left, motion rate is low so the choice of polynomial order is less important and a lower polynomial order performs better since higher order polynomials fit to more of the noise. At right, motion rate is high, requiring higher-order polynomial to capture the changes in the covariance matrix over time.

modest STI lengths. The case simulated in Figure 5.7 is for an interferer at the same -9.8dB INR and $3^\circ/\text{s}$ motion rate of Fig. 5.2. In Figure 5.7, EPSP-7 denotes a 7^{th} -order polynomial was fit to the eigenvectors and the polynomial coefficients were used directly, and PSP-7 denotes that those same coefficients were used as input to the numerical optimizer.

Figure 5.8 explores four main factors that affect the performance of EPSP: 1) input INR, 2) interferer motion rate, 3) polynomial order, and 4) STI length. The first two are presented by the physical interference scenario, and the second two are chosen in the algorithm implementation. Choosing a polynomial order that is too low for the motion rate results in model-order mismatch,

as can be seen in Fig. 5.8(i), where the IRR is much improved with a 12^{th} -order polynomial compared with an 8^{th} -order polynomial because of the high motion rate. With the same input INR but much slower motion, as in Fig. 5.8(g), both polynomial orders are sufficient to produce good cancelation. In Fig. 5.8, the input INR ranges from -12dB to +12dB and motion rates are $3.2^\circ/s$, $1.4^\circ/s$, and $0.2^\circ/s$, with 2^{19} samples or 0.85 seconds of data included in the simulation, meaning that at the max motion rate of $3.2^\circ/s$, the interferer has moved a total of 2.7° or 1.9 sidelobe widths. Since the null depth for EPSP in Figures 5.7 and 5.8 does not depend strongly on the STI length of the preliminary covariance estimates, the main tuning parameter in the algorithm is the polynomial order. There is currently no method for guaranteeing the proper choice of polynomial order; at this point the approach is based on heuristic analysis of the data set. Determining a more rigorous method is left as future work.

5.4.2 Comparison with Subspace Tracking Spatial Projections

This section compares performance of the PSP and EPSP algorithms with the Subspace Tracking Spatial Projections method introduced by Ellingson for interference cancelation in radio astronomical applications [64]. This approach also uses subspace projection beamforming, but time varying estimates of the interference subspace are based on the sequential tracking algorithm PASTd developed by Yang [74]. PASTd has also been used successfully in time-varying MIMO wireless channel estimation for multi-antenna equalization and covariance tracking. There is an extensive body of related work on subspace tracking available in the literature of MIMO and other similar applications, including [107–111]. While the goal of this work is similar to that of MIMO subspace tracking, where each seeks to get a better estimate of a particular subspace of a covariance matrix, the goal is different from that discussed in this paper in two important ways. First, computational burden is a major concern for communication systems, and subspace tracking is used to reduce the number of samples required to get a good subspace estimate; here, computational burden is not the concern and the goal is instead to increase the number of samples used in order to improve the accuracy of the subspace estimate and correspondingly the null depth. Second, the subspace tracking approaches cited assume a higher-rank subspace, which is not necessary for this particular signal model. Using these approaches could be a useful step in extending this work to cancel multiple interferers, where a higher rank subspace would need to be

tracked. The PASTd-based method of Ellingson is an obvious candidate for the problem addressed in this chapter because it is designed specifically for non-stationary interference scenarios, and the zero-forcing beamforming method is expected to produce deep nulls.

The PASTd algorithm, as described in [64, 74] tracks the top M eigenvectors (assumed to correspond to interference) as follows:

For $n = 1$ to N do :

$$\mathbf{x}_r \leftarrow \mathbf{x}[n]$$

For $m = 1$ to M do :

$$\gamma_m \leftarrow (\hat{\mathbf{u}}_{m,n-1}^H) \mathbf{x}_r$$

$$\hat{d}_{m,n} \leftarrow \beta \hat{d}_{m,n-1} + |\gamma_m|^2$$

$$\hat{\mathbf{u}}_{m,n} \leftarrow \hat{\mathbf{u}}_{m,n-1} + \frac{(\mathbf{x}_r - \gamma_m \hat{\mathbf{u}}_{m,n-1}) \gamma_m^*}{\hat{d}_{m,n}}$$

$$\mathbf{x}_r \leftarrow \mathbf{x}_r - \hat{\mathbf{u}}_{m,n} \gamma_m,$$

where N is the total number of time samples to process, β is a remembering factor that weights the past eigenvalue estimate $\hat{d}_{m,n-1}$, and $\hat{\mathbf{u}}_{m,n}$ is the tracking estimate of the m th eigenvector at time n .

Figure 5.9 shows PASTd performance for various β values (along the top horizontal axis) as applied to the signal scenario used in Fig. 5.7. For comparison, SP, EPSP and PSP are plotted on the same INR out y-axis, with STI length along the bottom x-axis. The effective data window length for exponentially weighted averaging is determined by tuning parameter β which acts similarly to STI length L in SP. However, direct comparison between β and L is difficult. Values in the range $0.95 < \beta < 1$ are necessary to produce adequate noise averaging, but as $\beta \rightarrow 1$, performance suffers because of poor tracking of motion changes. With the proper choice of β , PASTd works nearly as well as choosing the optimal STI length for simple SP.

Clearly, both EPSP and PSP outperform PASTd in terms of cancellation null depth and sensitivity to tuning parameters (L vs β). This suggests that though all these methods use the same basic subspace projection approach to compute beamforming weights, and each is designed




Figure 5.9: Comparison of output INR for PASTd using different values of β , with SP, EPSP and PSP using different STI lengths, L .

to operate in a non-stationary interference environment, there are significant differences in the accuracy of the respective time-dependent covariance estimates. PASTd is a sequential process intended for dynamic subspace tracking. It excels in terms of computational efficiency, and is well suited for real-time canceling applications. Careful, perhaps manual, tuning of the β parameter may be required to match motion rates. SP is a short block processing method that has moderate computational burden, and can be used in real-time applications as long as data buffering of STI length L is possible. As shown above, this block length must be tuned manually for best null depth, but even this tuned level suffers under rapid motion. Both PASTd and SP are easily generalized to multiple interferers by increasing the number of eigenvectors assigned to the corresponding subspace estimate.

EPSP and PSP are long block processes which use all the data over an extended window, spanning significant motion, to estimate low order model parameters to track covariance changes. This enables lower error variance estimates of the time-dependent array covariance so that deeper nulls can be formed. Computational burden for EPSP is moderate, and high for PSP. Neither method at present is suitable for multiple interferer environments. Both require long data buffering

and latency that could be challenging for real-time implementation. They are well suited for post correlation processing. Most importantly, they produce the deepest cancelation nulls yet demonstrated for low INR, moving interference scenarios. This may be critical in a number of radio astronomical observing scenarios.

5.4.3 Real Data Experiments




Figure 5.10: The 19-element array mounted on a rotating table in the anechoic chamber. To get interferer motion with respect to the array, the transmitter was held fixed while the array was rotated.

Anechoic chamber tests verified EPSP experimentally using signal scenarios similar to those described in 5.4.1. An array of 19 co-polarized linear $\lambda/2$ dipoles tuned to 1600 MHz over ground plane was mounted on the rotator table in the anechoic chamber, as shown in Fig. 5.10, and the same LNAs were used whose noise parameters were used in the simulation results. Interference was generated using a CW tone at 1612 MHz (an important radio astronomical observation

frequency for OH ion emission), sampled at 1.25 Msamp/s, and the final system bandwidth was 450kHz. No SOI is included since it can be assumed the SNR is very low for an astronomical source. Interferer motion is created using a fixed transmitter and moving the array on the rotator table. Motion rates of $15^\circ/\text{sec}$, $7.5^\circ/\text{sec}$, and $3.25^\circ/\text{sec}$ were observed. Raw voltage samples were streamed to disk, then digitally filtered and complex basebanded in post processing, where the subspace projection algorithms were applied to cancel interference. The quiescent formed beam was a max sensitivity boresight beam. Results are shown with prewhitening as described in Section 5.2.3.

In simulation, computing the output INR is simple because the SP beamformer weights can be applied to the interference data and noise data separately to find the post-beamforming INR. It is more difficult to compare the results using experimental data since both SP and PSP cancel the interferer to below the noise floor. Instead of using a total power INR measurement, the PSD of the output signal is used to determine the residual PSD level in the particular bin with the interfering CW tone. To detect residual interference below the noise floor a noise-only PSD is subtracted from each curve.

Spectral scooping is clearly present with these experimental data, as is expected for narrowband interference. Thus the PSD shown for SP has been corrected using the method proposed in [69] to remove the scooping bias before a reasonable comparison could even be made. Since PSP inherently overcomes spectral scooping, no explicit bias correction was required with that algorithm.

Figures 5.11 and 5.12 show results using 0.63 seconds of data in two different interference scenarios to highlight the main tradeoff in SP between sample estimation error and subspace smearing, showing that EPSP outperforms SP in both cases. In Fig. 5.11, a slow-moving interferer ($3.25^\circ/\text{sec}$) is used so that subspace smearing is minimized, and the input INR is low (-12dB) so that sample estimation error dominates SP. For reference, a noise-only curve is included, highlighting the fact that both SP and EPSP cancel the interferer to at least the noise floor. The goal of PSP is to push interference even deeper below the noise floor for detection of extremely low-SNR signals. Figure 5.11(a) shows the PSD at a short STI window length, and Figure 5.11(b) shows the residual level in the PSD bin of the interferer as a function of STI length. The plots show the type

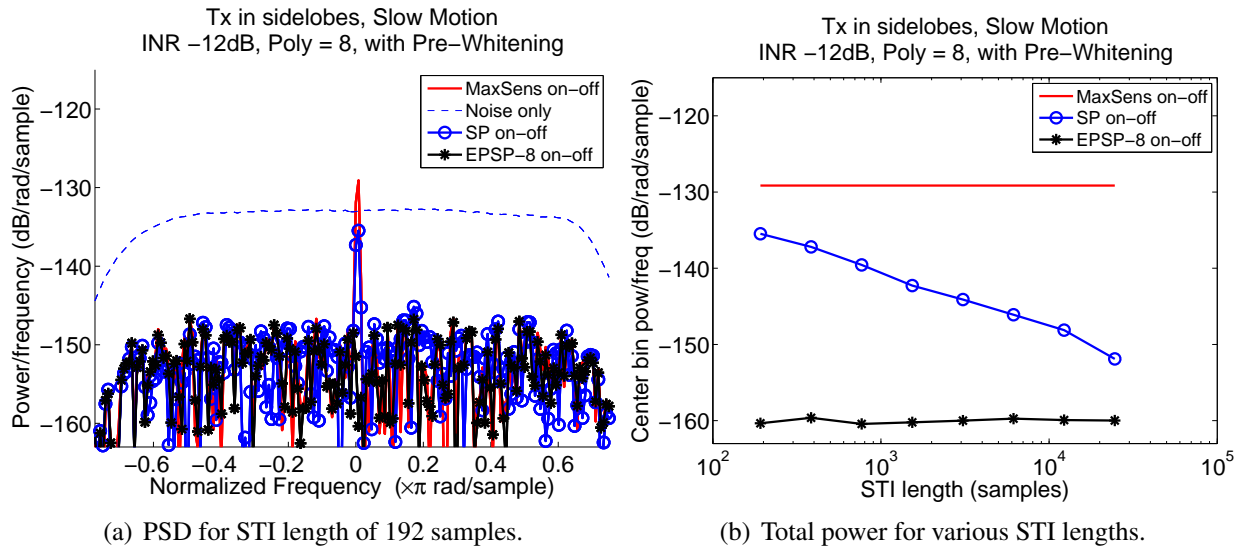


Figure 5.11: Real data experiments showing SP improvement for a slow-moving low-INR CW tone interferer in the beampattern sidelobes. In (a), the interference-free noise floor is also included, highlighting the fact that the final interference levels are below the noise floor and on-off subtraction is required to reveal the final residual interferer power. (a): PSD of beamformer output for interferer in sidelobes, using shortest STI window. (b): Residual power in center PSD bin corresponding to (a) for various STI lengths.




Figure 5.12: Real data experiments showing SP improvement for a fast-moving high-INR interferer. After on-off subtraction, (a) shows residual interference power below the noise floor. (a): PSD of beamformer output for interferer in sidelobes, using long STIs. (b): Residual power in center PSD bin corresponding to (a) for various STI lengths.

of behavior expected, where longer STIs reduce the sample estimation error allowing SP to give deeper nulls, and EPSP remains low regardless of STI length.

Figure 5.12 shows a fast-moving ($15^\circ/\text{sec}$) interferer at high INR (30dB) so subspace smearing dominates over sample estimation error. Here, longer STIs do not give significant sample estimation error improvement because the INR is already high. An 8th-order polynomial adequately captures the motion and EPSP keeps the interferer at or below the noise floor. The performance of SP begins to degrade because the longer STIs include more and more motion and the beamformer is only updated once per STI, whereas the PSP beamformer weights can be calculated at arbitrarily fine intervals.

5.5 Conclusions

The principal contributions of this chapter are: 1) characterizing the performance of SP beamforming and exploring limiting factors including sample estimation error, subspace smearing, and colored noise bias, and 2) proposing a polynomial fitting algorithm to better model interference time variation due to motion, thus improving cancelation null depth by 6 to 30dB, depending on STI length. It produces the deepest cancelation nulls yet demonstrated for low INR, moving interference scenarios, which is important in radio astronomical observing scenarios. A subtle but critical aspect of the approach is phase normalization required to give smooth eigenvector progression. It was also shown that the polynomial motion model eliminates spectral scooping that normally affects SP for narrowband interferers. Null depth improvement over SP was shown in simulation and verified with experimental data. Though an astronomical phased array feed experiment was used to motivate and demonstrate the algorithm, these methods should have general application to many other array adaptive canceling applications with moving interferers.

Future work includes extending these results to multiple interferers and apparent rank increase due to multipath scattering from supporting structures. The polynomial order is a tuning parameter like β in the PASTd algorithm; developing a more rigorous method of choosing polynomial order will add further robustness to the algorithm, and the CRB analysis should also be extended to include the notion of mismatch between the chosen polynomial order and the true polynomial order. Experimental results were shown here for observations on the order of seconds, but astronomical detections may require observations on the order of hours, so real-time

processing will need to be considered before this algorithm can be implemented on an operational astronomical instrument. Real-time concerns include rapid polynomial parameter estimation and beamforming as well as automating the parameter estimation problem so that the proper polynomial order can be chosen for the particular dataset. Obtaining good noise-only calibrations for the noise whitening step is important. In this chapter only SP beamforming was considered, but polynomial fitting can be extended to other interference mitigation algorithms since any algorithm requiring better covariance estimates for a changing subspace will also benefit.

Chapter 6

CASPER-Based Digital Back End for the BYU PAF

6.1 Introduction

This chapter details work to expand the digital PAF back end from a PC-based system doing correlation and beamforming in post-processing to an FPGA-based system using CASPER hardware for a real-time correlator/beamformer. The system consists of a hardware frequency channelizer (F engine, in keeping with CASPER parlance), correlator (X engine), and beamformer (B engine), each with its own dedicated FPGA board.

The system described in Chapter 3 samples 20 inputs at 1.25 Msamp/sec yielding 425 kHz of final system bandwidth using five PCI cards mounted in a single server-class PC. An intermediate system developed in 2010 by other students keeps the same sample rate and final bandwidth but uses five PCs, each with two ADC cards, and a main data aggregation computer. The system described in this chapter uses a single 64-input ADC card sampling at 50 Msamp/sec, giving 20 MHz of usable bandwidth per input. This is connected to an FPGA board for real-time processing and streaming to disk over 10 gigabit Ethernet (10GbE). The increases in bandwidth and number of inputs are significant steps forward in making a scientifically useful PAF system. Ultimately the goal is a 300 MHz bandwidth system with up to 80 dual-pol antennas for the Arecibo Observatory or GBT, but that development is still in the future.

This chapter covers three major topics: current digital back end progress at BYU, recommendations made to the Arecibo Observatory on hardware development for a future phased array feed beamformer, and recommendations for incremental development of a hardware correlator/beamformer here at BYU. During the summer of 2010 this research group ran the AO40 feasibility study, a project using the BYU PAF to map the focal field of the Arecibo Radio Telescope to investigate PAF feasibility. One of the project deliverables was a report recommending FPGA-based digital back end development. These recommendations provide insight into the siz-

able effort required to build a hardware back end, even in the simplified, easily reconfigurable environment offered by CASPER tools described in this chapter.

6.1.1 International CASPER Collaboration

The Center for Astronomy Signal Processing and Electronics Research (CASPER) [1] at the University of California, Berkeley was created as a joint effort between the Electrical Engineering and Computer Science Department and the Astronomy Department to develop tools for astronomical digital back end development using reconfigurable hardware. The center has become an international collaboration including major leadership from Harvard and the Karoo Array Telescope in South Africa. A principal philosophy of the CASPER collaboration is to create a scalable solution for back end development using 10GbE for communication between ADCs and computational nodes. FPGA firmware development is simplified using MathWorks® Simulink® blocks from the XILINX® System Generator Tools® library that gives a graphical environment for designing digital hardware. The graphical flow diagram is translated into VHDL and compiled to a bit file, relieving the user the challenging learning curve of mastering VHDL for digital design. The CASPER toolflow includes specialized signal processing blocks like polyphase filter banks and correlation engines, in addition to basic building blocks like multipliers and muxes provided by XILINX®. CASPER also provides ADC and FPGA boards built to interface easily with the CASPER toolflow. The iBOB and ROACH FPGA boards and the quadADC (four inputs, 200 Msamp/sec) and iADC (two inputs, 800 Msamp/sec) are commonly used throughout the CASPER community. See [1] for hardware details. The BYU Radio Astronomy Systems (RAS) group has four iADC and four ROACH boards, along with a 64-input ADC, the x64adc, discussed in Section 6.2.

6.1.2 BYU Involvement with CASPER

The Precision Array to Probe the Epoch of Reionization (PAPER) is a synthesis imaging array created by the late Professor Don Backer and Dr. Aaron Parsons at UC Berkeley. It was installed at NRAO-GB with collaboration from NRAO-Charlottesville, the University of Virginia, and the University of Pennsylvania. The PAPER correlator, built primarily by Dr. Parsons and

Jason Manley of the Karoo Array Telescope (KAT) in South Africa, is one of the largest hardware correlators built using CASPER hardware to date. From January through July 2010, Professor Backer and Dr. Parsons sponsored me at Berkeley as a visiting student researcher to help maintain the PAPER correlator and deploy it in Green Bank in its 100 MHz bandwidth, 32-input configuration. This system used four iBOBs for F engines and four ROACHs for X engines. The experience was a valuable opportunity to become familiar with CASPER development, hardware correlators, and begin preliminary tests using the x64adc before returning to BYU to begin work on a CASPER-based digital back end.

6.2 Current Beamformer/Correlator Development Efforts at BYU

A 20 MHz usable analog bandwidth, 64-input (32 dual-polarization antenna) back end has been in development at BYU since mid-2010. This instrument will support initial proof of design concept experiments and will be usable for real-time operations with sufficient bandwidth for meaningful real science observations. The current plan is to use this platform for the first science-ready PAF system installation on the GBT. The FPGA gateware developed for the 20 MHz back end will be a major stepping stone for the ultimate 300 MHz BW, 80 antenna system, but with significantly reduced cost. Insights, experience, and technology familiarization obtained with the 20 MHz system will be invaluable and reduce risk factors when the full 300 MHz system development is undertaken.

6.2.1 System Description

The BYU prototype digital back end is built using the UC Berkeley CASPER libraries and hardware. The present system shown in Figure 6.1 uses a single ROACH board with the Xilinx Virtex-5 (SXT95) FPGA. The ADC board is the x64adc [112] which supports 64 inputs with 12-bit sampling. Eight ADC chips are used, each sampling 8 independent inputs at 50 Msamp/sec.

The x64adc does anti-alias filtering at the ADC inputs with a 47pF capacitor after the balun transformer to convert single-ended input signals to differential pairs. Figure 6.2(a) shows a circuit diagram [113] where the filtering capacitor is labeled C102. This capacitor will be removed from each ADC input to eliminate the anti-alias filter and permit baseband subsampling techniques. The




Figure 6.1: Interior view of the 64-input x64adc (left) connected to a CASPER ROACH FPGA board (right). The x64adc uses 8 ADC chips, each sampling 8 independent signals for 64 total. The ROACH board hosts a Virtex-5 FPGA (mounted under the fan in the center of the board). An on-board PowerPC hosts Linux operating system that shows FPGA registers as files that can be read and written using Linux. Four separate 1Gbps Ethernet ports (far right) allow an external computer to host the ROACH and enable high-speed inter-ROACH communications with an external multiport 1 GbE switch.

PAF analog front end electronics set the final IF band between 27.5 MHz and 47.5 MHz. The ADC running at 50 Msamp/sec samples sub-Nyquist to produce an aliased copy of the IF band, yielding the desired image for the IF between 2.5 MHz and 22.5 MHz at the ADC inputs, illustrated in Figure 6.2(b). In the diagram the passband is shown with an artificial slope as a reminder that the IF band shows up in reversed frequency order because the particular Nyquist zone chosen contains an aliased copy of the negative frequencies.

Figure 6.3 shows the BYU ROACH and x64adc installation in a chassis with a front panel array of SMA connectors to accept signals from a 64 input analog receiver front end. Each SMA connector is attached to a length of RG-174 type coaxial cable purchased as an assembly from Samtec, Inc. (New Albany, IN) with 80 SMA connectors on one side and an 80-pin QSE/QTE connector on the other side for mounting to the x64adc board.

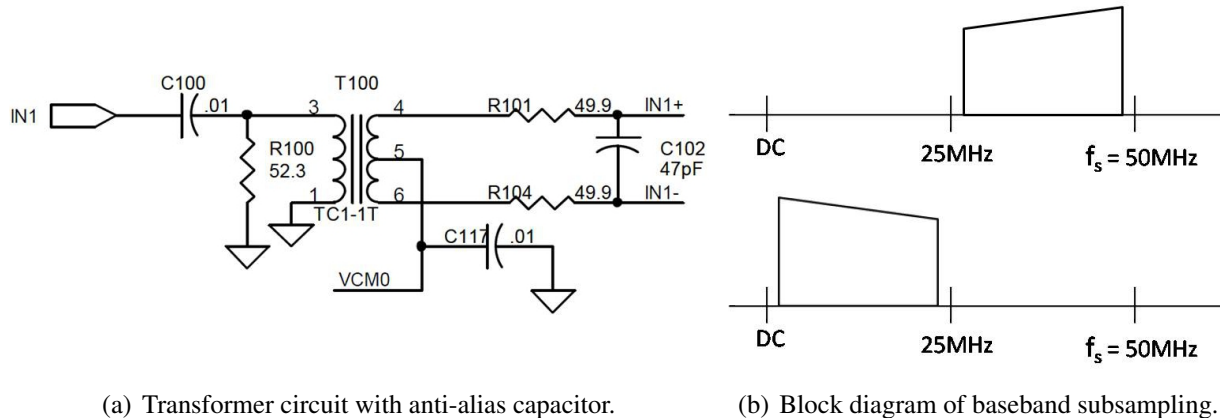


Figure 6.2: Transformer circuit with anti-alias capacitor and block diagram of baseband subsampling. (a): Circuit to transform single-ended SMA input to differential pair prior to ADC chip [113]. The capacitor labeled C102 in this diagram acts as an anti-alias filter that will be removed from each of these circuits on the x64adc to permit sub-Nyquist sampling to use a final IF higher than 25 MHz. (b): Baseband subsampling to use a final IF from 27.5 to 47.5 MHz, producing an aliased copy at baseband from 2.5 to 22.5 MHz. The artificial passband slope shown here illustrates the reversed frequency ordering introduced by using this particular Nyquist zone.

Figure 6.4 shows the BYU ROACH and x64adc installation in an instrument rack with the computers for control and gateway design. One computer is dedicated to gateway design because the CASPER toolflow requires very specific software version combinations. The toolflow computer runs 64-bit Red Hat Enterprise Linux 5.2 (RHEL5) with 64-bit MATLAB[®] R2009a and XILINX[®] ISE 11.4. Designs are created in MATLAB/Simulink and compiled using *bee_xps* into a *.bof* file – basically a bit file with wrappers used by the ROACH to program the FPGA.

One main benefit of the ROACH is that it has both an FPGA on board and a PowerPC running a modified Linux operating system called *borph*. Memory registers in the FPGA show up as files in the Linux OS. A separate computer also in the rack shown in Figure 6.4 running Ubuntu¹ Linux is used for interfacing with the ROACH and running control scripts by remotely logging into the ROACH operating system over SSH to read/write to memory registers. The *.bof* file is run as a program from the command line on the ROACH. This programs the FPGA with the bit file that

¹Although the gateway compilation computer must use RHEL5 to make the CASPER toolflow work, the control computer is not bound by the same constraints. The choice to use Ubuntu on the control computer was motivated by two considerations: 1) the PAPER correlator in Berkeley used Ubuntu and system debugging is aided by using identical configurations where possible; 2) much of the existing Python software for interfacing with the ROACHs and running tutorial scripts required a newer version of Python than was available on RHEL5 and significant work would have been required to get newer Python on RHEL5 since its package manager relied on the older version of Python.

would otherwise need to be loaded using a JTAG programmer or similar process. Interacting with the ROACHs is done using the control computer that administers a mini-DHCP server, assigning DNS names and IP addresses to each ROACH based on the MAC address of its Ethernet PHY. The ROACH filesystem is large enough to require external memory for hosting the filesystem; this can be a USB drive or MMC card that comes with the ROACH, or in our case, network mounted filesystem on the control computer running network filesystem (NFS). Using a network mounted filesystem ensures all ROACHs have the same filesystem since they all look to the same network location for it.

The control computer described here is the same one that shows up in the figures in Section 6.4. All interaction over 100 Base-T to the ROACHs is done from the control computer over SSH or telnet to program the ROACH, run control scripts and read/write memory registers. The CASPER toolflow design computer described here serves a dual purpose of CASPER design work and writing 10GbE packets to disk using its 10GbE card. This is the computer that shows up in the single PC stream-to-disk diagrams of Section 6.4, that will ultimately be replaced with a rack of computers for writing 10GbE data to disk.

The initial installation shown in these figures will support full 20 MHz spectrometer (F engine) operation, and streaming to disk of a small subset of the spectrometer frequency channels. A second ROACH board (already purchased) will be added in June, 2011 to provide real-time beamforming capability (B engine). The final expansion of the system will be to add real time correlation for calibration and continuum imaging (X engine). In the mean time, correlations required for calibration can be done in post-processing.

The current BYU 40 channel 425 kHz BW analog receiver front end downconverter system used last year in the AO40 feasibility study is currently being rebuilt to accommodate the 20 MHz BW and 64 inputs of the new ROACH-based real-time beamformer. The new design is complete, components have been purchased, and construction is scheduled for June, 2011.

Figure 6.5 provides an illustrative example of the software development environment for CASPER ROACH FPGA gateware. Flow diagrams are constructed graphically using MathWorks[®] Simulink[®] blocks from the XILINX[®] System Generator Tools[®] library. The signal flow block diagrams are compiled into gateware that is downloaded onto the CASPER FPGAs. The subsystem




Figure 6.3: Close-up front view of SMA bulkhead for the ROACH x64adc system.

shown performs the polyphase filter bank digital receiver/frequency channelization operations for the F engine functions in the back end processor.

6.2.2 BYU Real-Time Beamformer Capabilities

This section summarizes the current state of progress and near-term performance goals for the ongoing PAF back end processor development effort with the 64 input, 20 MHz system. Three principal functional capabilities are considered, including real-time spectrometer operation, data acquisition and streaming to disk for post processing and analysis, and real-time beamforming.

- Spectrometer:
 - Current System:
 - * Small number of frequency channels and antenna inputs for testing end-to-end functionality.



Figure 6.4: Equipment rack configuration with ROACH and 64 RF input connections for the x64adc (top), a second ROACH which will perform beamforming operations (middle) and host computers and data file server (bottom).

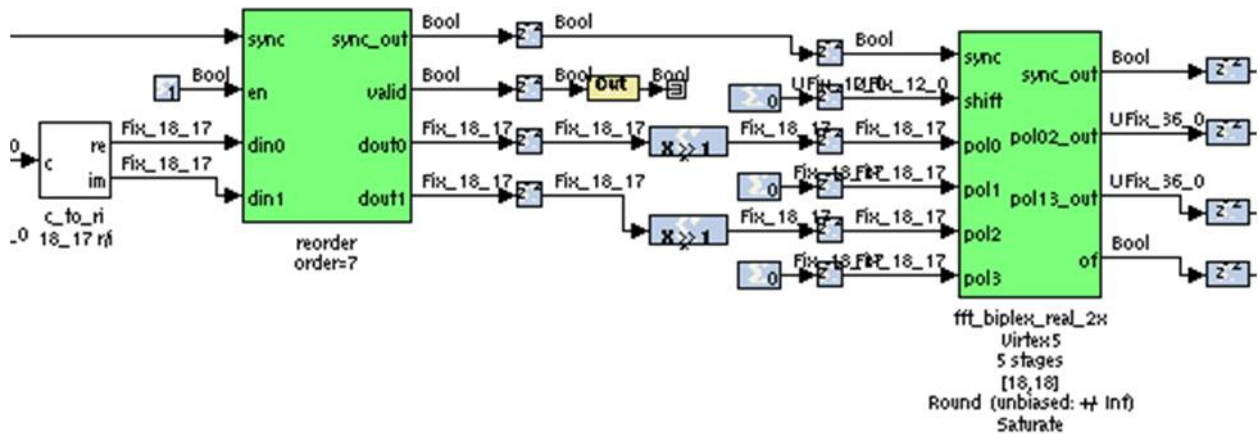


Figure 6.5: Design environment snapshot. CASPER toolflow provides Simulink blocks with signal processing blocks like polyphase filter banks (PFBs), FFTs, and correlation engines.

- * 8 inputs, 20 MHz BW (50 Msamp/sec), 16 frequency channels,
 - * Power per frequency channel accumulated in BRAM on the ROACH. A software register allows the user to select groups of 4 antenna inputs out of 64. Accumulate for up to 214 samples of an FFT bin output before overflowing the integer data type. This corresponds to $137\mu\text{sec}$ accumulations (the decimation by 2 going from real to complex data gives a decimated sample rate of 25 Msamp/sec so each of the 16 frequency channels repeats once every 640ns meaning 214 accumulations require $137\mu\text{sec}$). Read out accumulated spectra from the ROACH over 100 base-T Ethernet using a Python *katcp* client that interacts with *tcpborphserver* running on the ROACH. Spectra are plotted and animated using modified Python code provided by CASPER spectrometer tutorial.
- Near Future Development Strategies:
 - * Testing to see how the system scales when using all 64 inputs, increasing number of frequency channels output by the polyphase filter bank (PFB), increasing accumulation length and data bit widths. Determining at what point the design will run out of BRAMs and logic devices on the FPGA.
- Streaming raw spectral subband samples to disk:
 - Current System:
 - * Streams raw samples from ADC into a BRAM where they can be read out over 100 base-T Ethernet. Available BRAM space limits this to several thousand samples (8k samples = 40 microseconds).
 - * System is useful for understanding ADC behavior and verifying correct synchronization in the design; 10GbE transmission required for stream-to-disk data acquisitions.
 - Near Future Development Strategies:
 - * Stream complex samples directly from PFB to disk without any accumulation; software correlator to get beamformer calibrators that feed back into a hardware beamformer or store the data for future analysis

- * Limitation: Disk write speed. A single low cost enterprise file server PC can write 250 MB/sec with a 4-drive RAID 0 array. Streaming 64 inputs at 50 Msamp/sec with 12-bit raw samples produces 38.4Gbits/sec (4.8GB/sec). Reduced bandwidth is required for writing to disk at 250 MB/sec, at least initially.
 - * Several iterations will be required to determine the best way to format the 10GbE packets and headers to optimize disk write speed (large packets) and break up spectra for groups of antennas in a sensible way to allow easy packetization logic in hardware and easy reassembly of data from packets written to disk.
 - * Initially stream a single frequency channel per spectrum to disk, set up logic to insert meta data into 10GbE packets including antenna number and frequency channel numbering with the goal of obtaining beamformer calibrators
 - * Expand the system to stream a user-selected range of frequency channels to disk for all antenna inputs, either to get a more wideband time series signal, or to permit calibrating separate beamformers to different frequency subbands.
- Real-time Hardware Beamformer
 - Near Future Development Strategies:
 - * Add one more ROACH board and a 20 port 10GbE network switch to the system (ROACH and network switch already in hand).
 - * Begin forming a small number of beams to find out how much space this will require in hardware. The main difficulty is routing all the frequency channelized data from all antennas to the beamforming hardware. Multiplying by beamformer coefficients and summing will have a relatively small cost in hardware resources.
 - * Beamformer coefficients will feed into the ROACH using software registers so they can be updated with the calibrations found with the software correlator.
 - Long term goals for the x64adc system:
 - This is a 20 MHz Bandwidth system used as an intermediate step toward the final 300 MHz, 160-input AO40 system.

- The immediate long term goal is a 64-input full correlator/beamformer for a joint project with the University of Massachusetts to build a 64-element millimeter wave array
- A rack of 10 computers will be connected to the ROACH via 10GbE switch, streaming all frequency channels for all antenna inputs for software correlation, beamformer calibrators, and future analysis.
- Hardware correlator
 - * Correlator for all 64 antenna inputs, all frequency channels
 - * Correlator for subsets of antennas and/or subsets of frequency channels

6.3 Feasibility Assessment of a CASPER Platform AO40 Back End

6.3.1 Background

From January, 2010, to February 28, 2011, a phased array feed feasibility study was conducted by Brigham Young University and the Cornell National Atmosphere and Ionosphere Center (NAIC) under a subcontract from NAIC. The goal of the feasibility study was to collect modeling results, experimental data, and preliminary design concepts necessary to enable development of an L-band wide-field PAF for the Arecibo Telescope. The key performance parameter is a field of view equivalent to 40 beams spaced at the 3 dB overlap point, so the planned full-scale PAF on Arecibo is referred to as AO40.

A BYU prototype 19-element phased array feed was deployed on the telescope. The PAF was mechanically positioned to move the feed throughout the focal region of the Arecibo optics. Array calibration and source observation data was collected in order to characterize the focal region, determine the optimal focus location for the feed, and determine the required size and number of elements of a full-scale PAF for Arecibo. During the Focal Region Mapping Experiment all data was sampled and streamed to disk storage for post-processing and analysis, rather than being processed in real time. Sampled data was limited to an analog bandwidth of approximately 425 kHz for each of the 19 dual-polarized antennas. Data from the focal region characterization experi-

ment has been processed, analyzed, and will be reported at the IEEE International Symposium on Antennas and Propagation in July 2011.

In addition to understanding the focal region of the telescope optics and optimizing PAF size and configuration, the second major element of a full-scale scientific PAF is the required back end signal processing hardware. PAF calibration requires correlation of all array output signals during calibration, which is a significant computational challenge. Beamforming requires complex multiplies and summation of array output signals over narrow subbands across the full operating bandwidth, which is also computationally intensive. For these reasons, understanding the hardware requirements, development time, and cost of a full-scale PAF back end are essential to laying the groundwork for the AO40 PAF on the Arecibo Radio Telescope. Accordingly, this section considers a prototype digital back end system currently in development at BYU and maps out design options for a full-scale back end.

6.3.2 System Performance Specifications

In the proposed AO40 system real-time back end, the target bandwidth is 300 MHz for approximately 80 dual-polarized antennas (160 inputs). The actual number of elements for the final AO40 PAF design may be different, but the Focal Region Mapping Experiment results indicate that an array size of 60 to 80 elements will achieve the target 40 beam field of view. Signal processing modes that the hardware should support include narrowband correlation for array calibration, broadband correlation for continuum observations and mapping, and real-time beamforming for pulsar surveys.

6.3.3 Overview

The following sections address requirements for a real-time back end processor to support beamforming and calibration functions in the proposed AO40 system and identify several existing technology development efforts that could provide the needed hardware capability. Current progress in real-time beamformer development at BYU is presented. The fundamental approach is an incremental improvement in both bandwidth and input channel capacity to gain experience and develop software and hardware components needed to move toward the full AO40 capability.

As an intermediate step toward a full scale PAF back end, we are currently building a 64 input 20 MHz bandwidth system. The current and proposed hardware platforms are FPGA based, and at our current development stage the system is hosted on CASPER boards.

After describing the current prototype back end, an architecture for the full 300 MHz, 160 input system will be proposed, a brief overview will be presented of related development efforts at other institutions for PAF back end processors, and alternative architectures and development paths will be listed.

6.3.4 Feasibility Assessment

When considering the development path from our current 20 MHz CASPER PAF back end to the final AO40 300 MHz system, two important questions must be addressed: how much can be reasonably expected from a CASPER-based system, and how long will it take to develop? We can begin to answer these issues from our own experience, and by surveying other existing related CASPER systems.

An important observation about ROACH-based back ends is that while a significant number of useful gateway blocks are available in the CASPER community, modifying, extending, debugging, and integrating gateway blocks for large data acquisition and correlator systems is extremely time-consuming. Large-scale FPGA-based signal processing systems are in no way “plug-and-play”. The BYU narrowband 64 channel system has required more than one man-year for a team of two PhD students, one of whom had already spent six months working in Berkeley with the leading CASPER programming experts on the PAPER correlator, just to implement a working spectrometer (no X or B engines). Gateway development time for FPGA-based back ends is significant and should not be underestimated when scoping the resources required to develop a full-scale PAF back end.

Table 6.1 is a survey of current correlator systems from various groups which use CASPER hardware. Building a beamformer does not require a real-time correlator but these numbers give an idea of what is currently being accomplished with CASPER systems. Array correlations are needed for beam calibration, but these can be performed with only a fraction of the beamformer bandwidth, and can be computed off-line in post processing if narrowband samples can be streamed to hard drive. If the number of formed beams for a PAF is roughly equal to the number of antenna

inputs, then the computational complexity for a beamformer is similar to that of a full correlator, which implies the systems listed in the table are roughly indicative of what could be accomplished for beamforming.

Table 6.1: Summary of existing array correlator projects using CASPER FPGA systems

	Inputs	BW (MHz)	Freq Chans	Dump time	ADC	Samp Rate Msamp/s	Config	Hardware	Dump over 10GbE?
BYU	64	20	64+	Var.	x64adc (12 bits, 64 inputs)	50	1 ADC per ROACH	2 ROACH (1F, 1X/B)	In progress
KAT7 (South Africa)	16	400	512	30ms	katADC (8 bits; 2 inputs)	800	1 ADC per ROACH	16 ROACH (8F, 8X)	Yes (?)
PAPER (UC Berkeley)	32	100	1024	8 sec	quadADC (8 bits; 4 inputs)	200	2 ADCs per iBOB	4 iBOB (F), 4 ROACH (X)	No
MEDICINA	32	20	1024	1 sec	x64adc (12 bits; 64 inputs)	40	1 ADC per ROACH	2 ROACH (1F, 1X)	No
WEST FORKS	32	25	512	3 sec	x64adc (12 bits; 64 inputs)	50	1 ADC, 1 ROACH	1 ROACH	No
GMRT (Oxford, India)	64	20	256	100ms	x64adc (12 bits; 64 inputs)	40	1 ADC per ROACH	2 ROACH (1F, 1X)	In progress

The PAPER project, based at UC Berkeley where CASPER originated, aims to continue increasing the number of inputs to their correlator while keeping the same 100 MHz bandwidth. They are working toward a 64-input correlator that may be completed this year. A 128-input correlator may be as much as two years away. For a group starting from scratch, the development time will be significantly longer. In addition to the hardware development, supporting software must be written to control the correlator, read from the ROACH, and write to disk.

One important conclusion to be drawn from this table is that although bandwidths similar to the AO40 design goals have been achieved (KAT7) and about 1/3 of the target input channels have been implemented on CASPER hardware (GMRT), no existing system even approaches the combined requirements for 300 MHz BW and 160 inputs planned for AO40. A CASPER-based

system of this magnitude, though theoretically feasible, would be entering new territory, and would pose appreciable development risks. Figure 6.6 illustrates a possible architecture for such a system. Each FPGA pictured represents one CASPER ROACH board.




Figure 6.6: Straw man block diagram of possible 300 MHz AO40 back end. Each FPGA can be implement with a CASPER ROACH Board. F engines perform digital receiver operations, including complex base banding and frequency channelization. X/B Engines perform real-time beamforming or correlation processing.

The biggest unknown in this design is the number of required X/B engines (40 shown in the figure). For smaller systems the number of F engine and X/B engine FPGAs is roughly equal, but X/B processing and crossbar data communications requirements grow as the product of the number of input channels and the number for formed beams. Thus this relationship may not hold for such a large system. Limiting the number of antenna inputs (i.e., to a small sub-array) used to compute a given beam may help keep the growth more manageable. As pictured, this system with 80 roach boards and a very expensive 160 port 10GbE network switch is large enough to push the limits of the CASPER design methodology. The cost of the hardware alone without programming is expected to be on the order of \$700,000 at today’s performance and prices. As mentioned above, gateway development time is also challenging, and based on our experience with CASPER systems would require several full-time staff over two or three calendar years to complete, even considering possible re-use of functional blocks developed by other CASPER users and our own narrow-band system development.

6.3.5 Alternative Hardware Approaches and Related Development Efforts

Several other groups are making progress in hardware for PAF back end processors. These should certainly be considered as options for the full performance AO40 implementation. The DRAO and CSIRO systems described below, like the CASPER approach above, are FPGA-based. However, they differ in that custom processor boards and backplanes for data distribution have been developed. Perhaps the biggest argument in favor of these systems is that the crossbar bus backplane approach, when compared to the 10GbE network switch proposed above, is likely a lower risk method for the high bandwidth, high crossover inter-processor communications needed in a large-N beamformer/correlator. Their designs are also far more physically compact than could be achieved with a CASPER implementation. Issues of intellectual property rights and cooperative development arrangements would of course need to be worked out. The following list outlines basic design features and the development status for three available hardware platforms.

- Software back end (CPU and/or GPU-based system)
 - Swinburne University (Adam Deller’s DiFX software correlator) [114, 115]
 - * A simple implementation that has been well tested.
 - * Would require a farm of networked PCs for near real-time operation.
 - * ADC hardware must be developed separately for AO40.
 - * Biggest challenge is getting the full 160 input, 300 MHz data stream from ADCs into the PC farm in real time, or to stream the data to disk.
- Custom-designed hardware back ends [116]
 - Dominion Radio Astronomy Observatory (DRAO), Brent Carlson
 - * Based on the ”Kermode” board– (pronounced ”Kurr-mo-dee”). This is an 8-FPGA ATCA card being developed in partnership with a commercial company.
 - * The board and its support package will be marketed as a turnkey system by this company, with preferential pricing for designated radio astronomy/science colleagues/collaborators.

- * The board support package would allow users to develop applications within the MATLAB/Simulink environment, download test vectors, and compare results with simulations.
 - * Board performance specifications are currently proprietary and are only released by DRAO under a nondisclosure agreement.
- Commonwealth Scientific and Industrial Research Organisation (CSIRO), for the ASKAP Focal Plane Array.
- * A complete Virtex 5 based (ES1) set of beamformer hardware has been developed. Recent upgrades have been made to ES2, using the Virtex 6.
 - * Currently building two complete beamformers using the Rev 2 of these boards. Expected completion is April 2011.
 - * Well on the way to completing the firmware for the Parkes Testbed. After this DRAO will move to an ASKAP capable system.
 - * Gateware for the Parkes Testbed implements 192 inputs, 304 MHz BW, 4 beams, using all inputs per beam.
 - * ASKAP implementation will be 192 inputs, 304 MHz BW, 36 dual pol beams, 63 inputs/beam. Beams can be traded for more inputs/beam. Channelization for beamforming is 1 MHz bin width, with final channelization of 18.5 kHz. The correlator load is independent of channel bandwidth, only affecting data output.
 - * A simple packet-based command and control system talks directly to the registers in the FPGAs.
 - * Performs array correlation during real-time beamforming using 20% of total FPGA capacity, on a few spectral channels and decimated time samples.

6.3.6 Conclusion

Nearly all major PAF back end efforts have converged on an FPGA-based solution. The software correlator on CPU effort described briefly in the previous section is the only important exception. Given the desired bandwidth and number of PAF elements, it is straightforward to

estimate the raw number of FPGA gates or chips required for a PAF back end, and to scale a few years into the future using the growth curve in the number of logic blocks per FPGA chip.

The difficult question is which architecture to use to inter-connect the FPGA chips. It is not clear at this time whether the CASPER packet switching approach is scalable to 80 channels and 300 MHz bandwidth, due to rapid growth of the required gateway development time with the size of the system. The low space efficiency of the ROACH architecture is another drawback. The CSIRO and DRAO efforts are designed for larger-scale total system throughput requirements, and so will be more space-efficient and possibly offer a shorter time to a working full-bandwidth back end. For these hardware efforts, unlike the CASPER platform, gateway programming expertise is at present confined only to the developing institutions. By the time that funding is requested or received for the AO40 PAF, more information will certainly be available that will be useful in selecting the ideal hardware solution.

As the CASPER, DRAO, and CSIRO hardware efforts progress, purchasing or otherwise acquiring a working PAF back end hardware platform and some of the required gateway may become feasible. In the interim, a reduced-bandwidth system similar to the BYU 64 channel back end could be used to provide initial science capability. For any of the approaches discussed above, hardware and gateway sharing with other back end development groups will save Arecibo time in bringing the AO40 back end on line. Even with a collaborative approach, additional custom programming will be required to accommodate to the detailed requirements of the AO40 PAF. Adequate staff resources for gateway programming must be factored into cost proposals for system development.

6.4 Future Development at BYU of a 64-Input Correlator/Beamformer Using CASPER Hardware

This section lays out a road map for the next stages of development for a 64-input correlator/beamformer using our existing CASPER hardware. This constitutes the relatively near-term plan for incremental capability improvements to be undertaken by the graduate students who take over the project. Figure 6.7 gives an overview of the full system architecture using three of the four ROACHs and one of the two x64adc's currently in hand, with one ROACH dedicated to each of the F/X/B engines. Phase 1 represents the initial completed work to configure the ROACH and




Figure 6.7: Overview of correlator/beamformer full system using the x64adc and three ROACH boards. ROACH 1 has the frequency channelization (F engine) gateway, ROACH 2 has a hardware correlator (X engine), and ROACH 3 has beamforming (B engine) gateway. Power per beam accumulated on the B engine ROACH can be formed into images on the display computer or raw samples can be streamed to disk over 10 GbE for further processing.

x64adc, culminating with demonstration of a real-time spectrometer using the x64adc. In Phase 2 of development, frequency channelized data are streamed from the ROACH over 10GbE to a receiving computer and written to disk. This phase also includes software development to interpret and reassemble packetized data, and development of a software correlator to find beamformer calibrators and produce beamforming coefficients. Phase 3 is the work to create a hardware beamformer, and Phase 4 is the work to create a hardware correlator. These two efforts can happen concurrently or in reverse order since neither relies on the other. The element weights needed by the beamformer are written into BRAM by the control computer with calibrations based on correlations from the software correlator of Phase 2 or the hardware correlator of Phase 4. In Phase 5 the software display is developed with power per beam read out of BRAM by the control computer, and optional raw data from each beam streamed to disk over 10GbE. Sections 6.4.1–6.4.5 present

a development road map for a basic architecture. Increasing the number of frequency channels or antenna inputs may require some of the alternatives presented in Section 6.4.6 to deal with resource limitations.

6.4.1 Phase 1: Real-Time Spectrometer

As an initial demonstration of hardware functionality a hardware spectrometer has been created using the x64adc and a single ROACH, showing that several milestones have been reached. Red Hat Linux, MATLAB, Simulink, and XILINX[®] System Generator successfully configured on the CASPER development computer, and basic designs can be compiled into *.bof* files and loaded onto the ROACH. VNC servers and viewers were installed, allowing multiple users to create designs simultaneously and remotely without interfering with each other. Network filesystem (NFS) was properly set up on the control computer, enabling all of the ROACHs to see the same filesystem, and each receives the proper IP address and domain name with the mini-DHCP server set up on the control computer, giving SSH access to each ROACH. The x64adc calibration script that synchronizes the 8 ADC chips runs successfully, giving a sync pulse once every four clocks verifying that the chips are still synchronized. Because the x64adc runs at 1/4th the rate of the ROACH, the CASPER yellow block that provides the interface between the physical ADC and the MathWorks[®] Simulink[®] design environment cycles through 4 different ADC inputs on successive clocks. This allows the ROACH to be run at full rate without using as many FPGA resources, but requires some reordering to reassemble the four interleaved ADC inputs on each yellow block line. A sync pulse that shows up once every four clocks is used to denote the first of the four ADC inputs is present. Integer samples written to BRAM can be converted into numbers that match the physical input voltage on the ADC.

Modifying the CASPER spectrometer tutorial to use blocks developed by the team at Oxford, a spectrometer was created at BYU to use the x64adc. Figure 6.8 shows the basic components of the spectrometer. The F engine consists of an FIR filter and Polyphase Filter Bank (PFB) connected to give better frequency channelization than could be obtained with a simple FFT, and a BRAM accumulator for the squared magnitude of each frequency channel. The spectrum is read out of BRAM using the control computer over the 100 Mbps Ethernet port on the ROACH to periodically read from the file on the ROACH filesystem mapped to the BRAM hardware location.

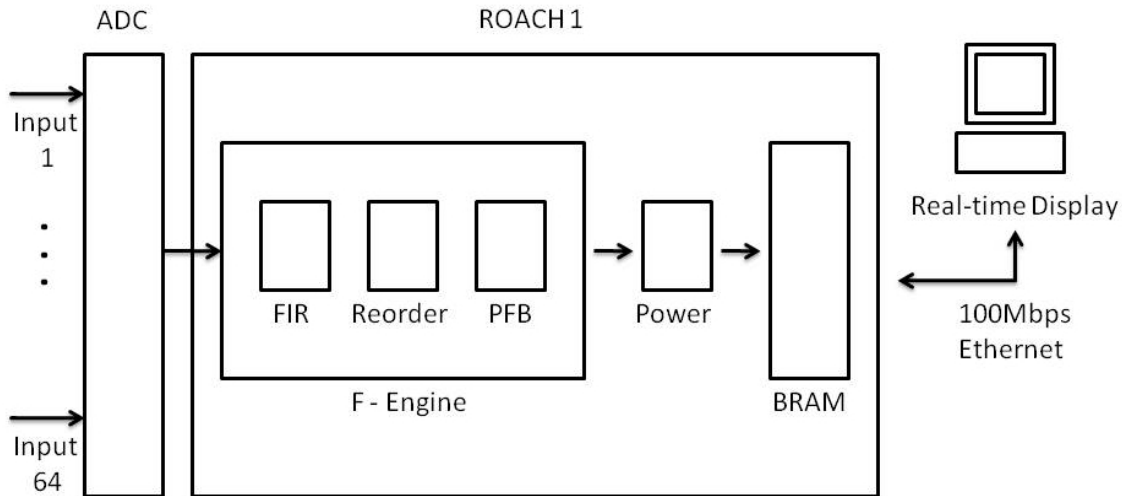


Figure 6.8: Initial hardware demonstration: real-time spectrometer using the x64adc. The F engine performs an FIR filter and polyphase filter-bank implementation to give better frequency channelization than an FFT could provide. Frequency channels stream out serially from the F engine to BRAM where power-per-channel is accumulated and then read out periodically over 100 Mbps Ethernet by the control computer.

To simplify design and debug, the design was created with a single set of spectrometer logic and a mux to select which ADC input goes into the F engine, and the number of frequency channels was kept to only 16 (i.e., 16 positive frequencies of a 32-point FFT-based PFB).

6.4.2 Phase 2: Streaming Raw Samples to Disk

This current phase of CASPER development at BYU aims to stream raw complex frequency channels from the F engine directly to disk over 10GbE and record them on a receiving computer running Gulp [117]. Reduced bandwidth will be streamed to the receiving computer in order to stay under the 250 MB/sec disk write speed, as illustrated in Figure 6.9. Streaming over the 4 CX4 ports on the ROACH will eventually allow for connecting over a 20-port network switch to a rack of receiving computers, allowing for the full bandwidth to be streamed to disk and later aggregated, as in Figure 6.10. The steps toward creating a limited bandwidth stream-to-disk design are:

1. Stream a single frequency channel from a subset of ADC inputs to a single computer.
2. Write software to interpret data streamed to disk and reassemble into a useful file format for analysis and correlation.

3. Add the ability to select multiple frequencies from a subset of ADC inputs.
4. Expand to stream multiple frequencies from all ADC inputs to a single computer.

Creating a full-bandwidth stream-to-disk design includes:

1. Based on a lookup table select different IP addresses for sending 10GbE packets.
2. Create a design with 4 10GbE transmit blocks, one for each CX4 port, and each with a lookup table for IP addresses.
3. Stream a single frequency from all ADC inputs to all receiving computers.
4. Write aggregation software to reassemble data from all computers into a single data stream.
5. Increase the number of frequencies until the entire bandwidth is streamed to disk.




Figure 6.9: Complex-valued frequency channelized data is sent without accumulation to the 10GbE packetizer for streaming to disk. Streamed data can be used for further spectral analysis or software correlation to get beamformer calibrators. The receiving computer runs Gulp, a multi-threaded program for capturing data, but disk write speed limits the bandwidth a single computer can handle so a subset of frequency bins are selected for streaming to disk.




Figure 6.10: With a rack of receiving computers to receive the frequency channelized data, the full bandwidth can be streamed to disk. Packets are sent to the appropriate IP addresses through the 4 CX4 connectors on the ROACH to a 10GbE network switch that connects the receiving computers.

At this point the main motivation in streaming voltages to disk is to create a replacement for the current 40-input data acquisition system used in PAF experiments. This requires only a narrow 425 kHz subband from the 20 MHz analog bandwidth available, so the reduced bandwidth stream to disk system has immediate utility. Eventually the ability to stream the full bandwidth to disk may be useful for system diagnostics and pulsar work. The first priority is to develop the software on the receiving computer to reassemble 10GbE packets arriving from the ROACH into a meaningful data stream that can be written to disk similar to the current data acquisition system and correlated to find beamformer calibrators and weights.

One major emphasis of current work to stream data to disk is the effort to insert meta-data into the 10GbE packets to identify which ADC inputs and frequency channels are contained in each packet, along with a master counter number to identify packet loss and allow packet reordering where necessary. Packet loss is probable because packets are streamed using UDP which is popular in the CASPER community for its higher speed over TCP/IP since there is no handshaking or

requesting retransmission of lost packets. The CASPER 10GbE block takes care of adding the necessary network headers for each packet, so the meta-data is added to the packet payload using a state machine to control whether ADC data or meta-data are being clocked into the 10GbE buffer at the appropriate times.

6.4.3 Phase 3: Beamformer Development

Figure 6.11 gives an overview of the next phase of development. Frequency channelized data are streamed from ROACH 1 to ROACH 3 for beamforming, and to an external receiving computer for software correlation to compute beamformer calibrators and coefficients. Beamformers are passed to the control computer and written into BRAM on ROACH 3 where they are used in beamforming. Streaming data between ROACHs is done over the CX4 ports using the XAUI (“zowie”) CASPER block for 10GbE data transfer that can be used in connecting two XILINX® devices and is leaner than the regular 10GbE block that must be used when connecting the ROACH to an arbitrary receiving computer. It is expected that ROACH 1 will utilize nearly all of its FPGA resources to implement the F engine, XAUI, and 10GbE blocks, and that the number of frequency channels may need to be small in order for the design to compile. Preliminary tests suggest ROACH 1 will be BRAM-limited and 256 frequency channels is the current maximum. To help reduce BRAM utilization, Figure 6.11 illustrates where XAUI and 10GbE are to be used. More testing will be required to identify what obstacles will be encountered in getting the data from ROACH 1 to ROACH 3, and solutions may require streaming data from the ADC to multiple ROACHs with identical F engines.

In many ways the beamformer is the easiest part of the system to understand and implement—a single complex multiplication per ADC input stream and sum over all streams. The difficulty arises in forming multiple beams and running the interconnections between all the ADC data streams and the beamformer coefficients. For a narrowband array one set of beamformer coefficients is required for each desired beam pointing, as in Figure 6.12. A wideband array requires either time-delay beamforming or with the frequency channelization step done first in this design, a different set of beamformer coefficients can be applied to each frequency channel. The complicated phase structure at the focus of a large reflector requires a new beamformer every 400kHz or so, rather than the 25 MHz that might be expected based on the small physical extent of the array.




Figure 6.11: ROACH 1 streams data over XAUI to ROACH 3 where beams are formed. Beamformer coefficients are written to BRAM on ROACH 3 by the control computer. Beamforming calibrations are found either by streaming data from ROACH 1 to disk and using a software correlator, as shown here, or streaming to the hardware correlator described in Section 6.4.4. Beamformer coefficients are then sent to the control computer for use in the B engine.

In fact, a new beamformer may be required every 40 kHz to correct for ADC sample clock misalignments, as recent investigations have shown on experimental data with the BYU PAF. Because of the large number of multiplications and sums for a large number of beams, the beamformer will likely be as computationally intensive as the correlator. The correlator described in Section 6.4.4 uses all the resources of ROACH 2, so it is likely that the beamformer will fill up ROACH 3 in a similar way.

A beamforming engine block (similar to the existing X engine block) is currently under development in the CASPER community to implement beamforming with scalable numbers of antenna inputs and output beams. The South African KAT effort has placed a high priority on beamforming. Glenn Jones at CalTech has demonstrated preliminary results with a “Pocket” Correlator/Beamformer/Imager (PoCoBI), where “pocket” refers to designing for a small number of inputs connected to a single ROACH with F,X,B, and Imaging engines all on the same board. BYU developers should pay attention to the CASPER mailing list to be aware of international beamforming developments.

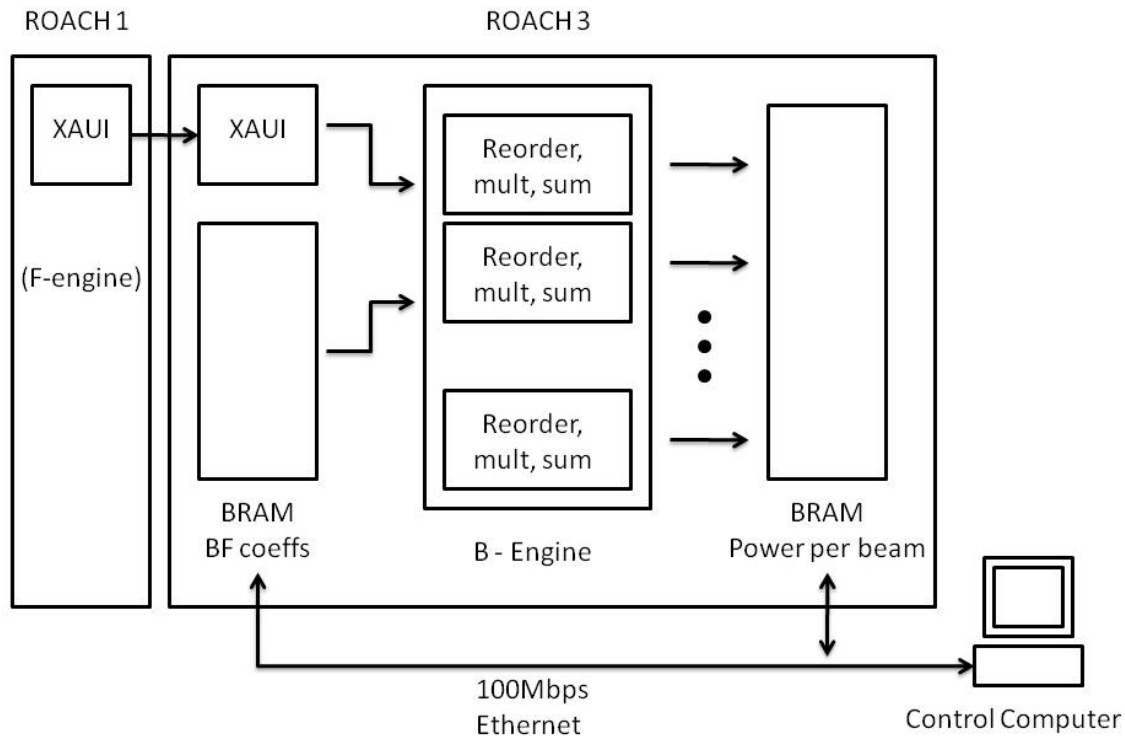


Figure 6.12: Beamformer detail. Beamformer coefficients are written to BRAM by the control computer for use by the B engine. Forming each beam requires the current frequency bin for all antennas, along with the beamforming coefficients for each simultaneously formed beam. The reordering and data connections required to form multiple beams makes the B engine of similar complexity as the X engine, and will likely use most of the available FPGA resources of ROACH 3. Accumulated power per beam will be stored in BRAM and periodically read by the control computer.

Beamformer coefficients will be written by the control computer into BRAM and updated as needed. During beamformer development, formed beam power can be accumulated in another BRAM for diagnostic purposes. Eventually unaccumulated formed beams may be streamed to disk over 10GbE for further processing, as in Figure 6.7.

6.4.4 Phase 4: Hardware Correlator Development

In the previous section Figure 6.12 shows beamformer coefficients are written to BRAM by the control computer. These coefficients can come from the software correlator developed in Phase 2 using raw samples streamed to disk, or from a hardware correlator if it exists. With sufficient manpower, Phases 3 and 4 can happen in parallel. Figure 6.13 shows where the X engine fits in the system.




Figure 6.13: ROACH 1 passes channelized data over XAUI to ROACH 2 for real-time correlation used in finding beamformer calibrators, and to ROACH 3 for beamforming. Correlations are read periodically from ROACH 2 by the control computer, and beamformer calibrators are then written to ROACH 3. Correlations are used by the PAF in finding beamformer calibrations that only need to be updated periodically, and the software correlator in Section 6.4.3 will suffice for preliminary work, so the hardware correlator is a lower priority for development.

Correlation is a very simple process of creating a matrix of all pairs of antennas multiplied together and accumulated. One correlation matrix is created for each frequency bin. The CASPER X engine block implements this process in a very efficient way with a minimum accumulation of 128 iterations. A good block diagram can be found at [118] with general information about the CASPER packetized correlator development at [119].

A hardware correlator based on the CASPER X engine block has been designed for the GMRT using the x64adc with 20 MHz BW and 256 frequency channels. The design uses a single ROACH for the F engine and a single ROACH for the X engine. The X engine ROACH is completely full and requires clever re-use of some multipliers to get the design to compile and fit into one ROACH, but can be spread across multiple ROACHs. A major benefit of using the CASPER X engine block is that it makes it easy to create correlators that can be spread across multiple ROACHs, with the F engine streaming groups of frequency channels from all antennas to different X engine ROACHs. Even with an existing hardware correlator design, several months will be

required to become familiar with the design and modify it as needed for the BYU development effort.

The primary purpose of the X engine is to create correlation matrices needed in finding beamformer calibrators, which happens on a long timescale with accumulations read out of BRAM over 100 Mbps Ethernet. The other purpose of the X engine is to create correlations that can be written to disk for reprocessing later. These are used for calibration processing or, for example, in post-correlation beamforming for beamshape modification or interference mitigation. For accumulations on the order of tenths of a second up to several seconds, the Python software package *corr* [120] developed by the CASPER community can be used. The *corr* package starts a process on each X engine ROACH in a correlator system that periodically reads the vector accumulator BRAMs, creates UDP packets, and streams them to an aggregation computer running multi-threaded receiving software which writes the correlations into a common astronomical data format, MIRIAD. The packets are formatted to implement the SPEAD protocol, a simple header/payload scheme for easily identifying the data contained in the packet so the receiving computer knows how to collect and write the packets to disk. For applications where correlation dumps are needed on the order of microseconds or milliseconds, streaming over 10GbE may be required and significant effort will then be needed to create the receiving software.

The full system layout in Figure 6.7 shows the configuration for the case where the correlations are recorded to disk for later processing and the real-time beamforming/imaging operations are going on as well. If correlations are only needed for beamformer calibration, ROACH 2 could be used for a short time to find correlations and then reprogrammed to be used for a different function in the beamforming system. This may be necessary if finer frequency channelization is required and the F engine needs to be spread across more than one ROACH.

6.4.5 Phase 5: Imaging Software Development and Streaming Raw Beams to Disk

Phase 5 is to create the software required for creating and displaying images. Processes such as cross-beam calibration and normalization, noise field flattening, and interpolation for mosaicing of adjacent image snapshots could be performed here. Since beamforming is done separately on each frequency channel, spectral analysis can be done on each beam separately or each image can be viewed at different frequencies.

A fourth ROACH may be required at this stage of development to act as an extension of the B engine ROACH if additional BRAM is required for forming beams or if compiling the 10GbE blocks is limited by the BRAM requirements of the B engine. Streaming formed beam time series samples to disk would be necessary for doing pulsar work or having a record of the un-accumulated data that could be reprocessed later. The work done in Phase 2 to stream data to a rack of computers over the 10GbE network switch, and the associated aggregation software can be reused here. The correlator/beamformer described in the preceding sections could also serve as a front end to the CASPER pulsar processor being developed at NRAO.

6.4.6 Tradeoffs and Alternative Architectures

This chapter has presented one potential architecture for back end development but along the way hinted at some alternatives in design choices, such as:

- Trading off the number of antenna inputs and frequency channels. The current F engine implementation (FIR filter, reorder, PFB) using modified CASPER blocks relies on BRAM to store frequency bins at certain points, limiting the number of frequency channels that can be computed. Applications requiring fine channelization may tolerate a reduced number of antenna inputs to reduce the BRAM requirement.
- Reconfiguring the system for different goals within an observing session. If correlations are only needed for beamformer calibrators, the B engine ROACH could be programmed at one time with an X engine design to get correlations and calibrations, then reprogrammed with a B engine design for the F/B engine configuration for real-time imaging. Reprogramming a ROACH is a simple matter of running a different *.bof* file. This would free up ROACH 2 for use elsewhere in the system, potentially to splitting the F engine across two ROACHs thus freeing up BRAM and allowing finer channelization.
- Additional hardware. Depending on budget constraints, it may be more cost-efficient to buy additional ROACH/x64adc hardware so the focus can be on getting a prototype system running earlier without needing such austere focus on efficient use of existing hardware. This could include the following options:

- Dedicating one ROACH to digitizing the raw data and streaming to two or four dedicated F engine ROACHs could enable much finer frequency channelization.
- Two x64adc boards using external clocking to synchronize their sampling, each connected to its own ROACH, followed by two X/B engine ROACHs. This could free up the resources to do imaging operations on the X/B ROACHs, so the four ROACHs in hand would be sufficient for the full correlator/beamformer.
- Implementing signal processing blocks in new ways. Resources can be freed up by implementing some memory storage operations using logic slices instead of BRAM where CASPER blocks may have been developed with an eye toward flexibility instead of efficiency. The time involved in getting the required detailed understanding of each CASPER block could be spent more effectively elsewhere, so the wisdom of this investment would depend on the individual developing the system.

Other design alternatives exist beyond those suggested in this section. The architecture road map of Sections 6.4.1 through 6.4.5 is one possible suggestion to use the hardware currently in hand, but as requirements for numbers of inputs and frequency channel bandwidth become clearer the final design will likely require alternative design choices.

6.5 Conclusions

This chapter has covered three main areas: progress at BYU with CASPER hardware including the x64adc and a single ROACH, recommendations to Arecibo Observatory on large scale digital back end development, and a road map for correlator/beamformer development here at BYU. An important result is that digital back end development is within reach for the BYU PAF design team. This chapter proposed a plan for increasing the digital back end sample rate by a factor of 20 and the number of antenna inputs by 50%. Designing custom FPGA boards would have been beyond the capabilities of the BYU team, but CASPER hardware is facilitating the creation of preliminary digital back ends in a relatively short time frame. Laying out a road map for building a 64-input, 20 MHz bandwidth digital back end with the hardware in hand is a major step for BYU PAF work.

The encouraging steps taken so far also reveal that development is slow and requires full-time effort and manpower. Even in the simplified plug-and-play world of CASPER with international collaborators developing packaged astronomy signal processing blocks, there is a substantial learning curve involved and each system modification requires step-by-step verification of logic and timing. After a six-month visit to Berkeley, it still took four months to assemble the BYU hardware and build a demonstration spectrometer. The following six months have been spent developing confidence with the spectrometer control and display software and working on 10GbE streaming to disk. With such a steep learning curve, full-time attention must be paid to ongoing digital back end development in order to deliver a science-ready instrument.

The analysis done for the AO40 feasibility study shows that a full 300 MHz bandwidth, 80 dual-pol input system will require major effort and expense. Continuing to use the CASPER 10GbE packetized correlator approach requires that a sufficiently large 10GbE switch can be obtained within budget. It is not yet clear that the CASPER approach will easily scale up to the size required for a large scale correlator, so the approach taken by CSIRO of developing a dedicated hardware backplane should continue to be an option. If a dedicated hardware backplane approach is taken, the CASPER tools may still be valuable in creating prototypes to assess hardware requirements of a large system and gain experience with digital design.

Future work on digital back ends will continue the preliminary effort described in this chapter. In addition to building the hardware back end described in Section 6.4, major effort will be required to develop software to make the instrument useful to the end user. That software development will require more time and people than can be offered by the BYU PAF team which will continue to focus on building prototype demonstration systems.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

The experimental results presented in this dissertation are some of the first astronomical observations in the world using a phased array feed mounted at the focus of the Green Bank 20 Meter Telescope at NRAO in Green Bank, WV, demonstrating the highest sensitivity and lowest noise temperature of their time. This was a major group project involving more than 30 people at BYU and NRAO in which I played a major leadership role including primary responsibility for the design of the experimental system. Several results from the large group effort have been included in this dissertation to emphasize the importance of the system as a meaningful milestone in astronomical PAF development.

The BYU PAF is an array of 19 L-band linear dipoles backed by a ground plane with system bandwidth of 425 kHz feeding into a 20-input ADC sampling at 1.25 Msamp/sec streaming data to disk for post-processing. Embedded dipole radiation pattern measurements demonstrate the perturbed pattern due to mutual coupling as compared with isolated dipole modeled results. The experimental platform built from commercial off-the-shelf components has receiver electronics stable enough in gain and phase to permit re-use of calibration vectors obtained on different days. Total power measurements are shown using continuum sources (Cygnus A) and OH sources (W49N and W3OH), as well as a spectral detection of the OH line at 1665 MHz.

One of the most important results coming from this work is the development of a Y-factor technique using room temperature microwave absorber and cold empty sky to characterize the isotropic noise response of an astronomical array. This technique accounts for the added system noise due to mutual coupling and loss in the antennas, and using the measured isotropic noise allows beamformer outputs to be tied to physically meaningful results.

Increasing sky survey speed by forming multiple beams on the sky is one of the primary motivations for astronomical PAF development, and this dissertation presents several radio camera images made from grids of MaxSNR beams using calibration vectors obtained experimentally. Another promising aspect of PAFs is the ability to spatially cancel RFI using adaptive beamforming. RFI mitigation is demonstrated, showing the frequency spectrum and total power images made before and after cancelation.

Improved RFI mitigation is achieved in the challenging cases of low INR and moving interference by combining SP beamforming with a polynomial model to track a rank 1 subspace. Deeper beampattern nulls are obtained by incorporating all of the available data in the model instead of considering each block STI independently. Choosing the appropriate STI length a priori to achieve optimal null depth is very difficult in practice. An important contribution of this dissertation is the investigation of factors that limit covariance estimation in SP on a per-STI basis including sample estimation error, subspace smearing, and correlated noise bias. The polynomial model-based approach presented here performs well over a wide range of initial STI lengths, overcoming the need for knowledge of the optimal STI length. The PSP algorithm is highly robust to initial STI choice and EPSP has a threshold at very short STIs so that for reasonable STI lengths its performance is very similar to PSP with reduced computational burden. Performance achieved with the polynomial model in the particular motion scenario considered in the simulated results is better by 9dB for EPSP than the best achievable null depth for SP without the polynomial model, and over the range of STIs considered here, it is better by several tens of dB. Overall, PSP can be expected to give nulls by a factor of between 6 and 30dB depending on the particular signal scenario and the blind choice of STI length for block-estimated SP. Real data experiments provide experimental verification of the principles described in this dissertation.

Several implementation details are considered for PSP. Simply normalizing to zero the phase on the first array element ignores the true phase variation (unobservable through a covariance estimate) due to physical interferer motion and reflector focusing effects. The smoothness constraint used here does a better job of preserving the continuous variations expected for the signal model. Expressions are given for the gradient, Hessian, and Jacobian of the algorithm's objective function for use in numerical optimization. The Cramer-Rao Bound gives insight into the ability to estimate the polynomial coefficients.

Preliminary work to increase system bandwidth is presented here with an FPGA-based digital back end using a 64-input analog-to-digital converter running at 50 Msamp/sec and the ROACH processing board. This platform puts a significant bandwidth increase within reach, jumping from 425 kHz to 20 MHz. The road map toward a 64-input 20 MHz correlator/beamformer lays out incremental steps from initial frequency channelization experiments, to a software correlator and hardware beamformer, and then adding a software imaging block. Progress made by other groups with CASPER hardware in correlator/beamformer development is included here, along with alternative approaches to digital back end development for large scale PAFs.

7.2 Future Work

Future experimental PAF work will focus on increased sensitivity and bandwidth, and improved calibration procedures. These results used a single-polarization array with narrow bandwidth. Extending the array to dual-polarization feeds and wider bandwidth presents additional beamforming challenges, particularly because the orientation of the feed changes with respect to an astronomical source as it moves across the sky. Beampatterns must be characterized farther out into the sidelobes and system stability must be improved to enable weaker source detection. Work has been done to incorporate array calibrations into deterministic beamformer design [95], and continued work focuses on calibrator stability over time as well as new methods for calibration that do not require repeated time-consuming scans over bright sources to update the calibrator grids. Imaging artifacts must be removed along the mosaic tile boundaries in mosaiced images seen in Figures 4.12 and 4.13(a). Variations in sensitivity with beam steering angle must be removed to produce a flat response to an isotropic noise field over the entire field of view.

Sensitivity has been the focus of recent work to extend these results including minimizing the mutual coupling penalty through active noise matching [42] and designing antenna elements which are optimally matched to the noise parameters over a range of beam steering angles [45]. NRAO is working with BYU to develop a cryo-cooled array for improved sensitivity. Increased bandwidth affects the digital back end design and computational burden, as well as the antenna design. Optimal noise matching for antennas across a much broader band will continue to be a research topic.

Further research on PSP is required to extend these results to higher dimensional subspaces. This dissertation focused on the goal of getting very deep beam pattern nulls for the rank 1 interference subspace due to a single point source line-of-sight far-field interferer. Rank increases will be caused by canceling multiple interferers, spatially distributed sources, multipath and diffraction for an experimental PAF system. The experimental results presented here used an array in an anechoic chamber to demonstrate the principles discussed for a rank 1 interferer. On-reflector experiments with a single transmitter showed higher-rank interference than expected. The rank increase could have been due to some combination of multipath, antenna near-field effects and diffraction, or on-site RFI. Further investigation may reveal the cause of the rank increase.

The EPSP approach does not immediately transfer to using the first two or more eigenvectors. The dominant eigenvector is a good first estimate to the whole subspace that will vary smoothly for slowly changing subspaces, but the remaining eigenvectors can be thought of as estimates of the error between the first eigenvector and the true subspace. The series of second eigenvectors acting like error vectors change in a much more erratic way, and simply fitting a polynomial to them does not perform well. Possibly a hybrid approach using an algorithm that tracks a larger subspace, like PASTd, could be combined with an approach that uses the entire data record instead of a sequential estimation approach.

Experimental results were shown here for observations on the order of seconds, but astronomical detections may require observations on the order of hours, so real-time processing will need to be considered before this algorithm can be implemented on an operational astronomical instrument. Real-time concerns include rapid polynomial parameter estimation and beamforming as well as automating the parameter estimation problem so that the proper polynomial order can be chosen for the particular dataset. Further work can be done to analyze the computational burden of PSP and EPSP.

Astronomers who may use the algorithm would benefit from clearer guidance on choosing the appropriate initial STI length and polynomial order for a given INR and motion rate, as well as the types of scenarios where this algorithm is most applicable. The noise whitening step requires good noise-only calibrations. Information about these types of parameter choices can be brought together into a formal document.

In this dissertation only SP beamforming was considered, but polynomial fitting can be extended to other interference mitigation algorithms such as LCMV since any algorithm requiring better covariance estimates for a changing subspace will also benefit.

The future work required for digital back end development is to complete development of a 20 MHz correlator/beamformer, and continue looking toward the proper approach to building a full 300 MHz bandwidth back end for an 80-input dual-polarization system. It is not yet clear that an FPGA-based hardware correlator approach is more desirable than a GPU-based software correlator approach. The ease of development for a software correlator is attractive and if transfer rates are sufficient to get data into the various GPUs included in a full correlator, the performance may be comparable to the FPGA approach. Staying with the FPGA approach, there are still alternatives to the 10GbE data transfer approach used by the CASPER community, notably the method used by CSIRO with large multi-FPGA boards inter-connected by a dedicated backplane. The packetized CASPER mentality relies on buying a sufficiently large 10GbE network switch and alternative approaches will have their own infrastructure costs. A certain amount of further analysis can be done with the available information about FPGA resources and data transfer rates, but the final choice must be based on experience gained through additional work to build preliminary designs. Keeping track of international developments and challenges faced with each of the various approaches will be important for ongoing work at BYU.

Immediate future work includes expanding the streaming-to-disk system to include recording wider bandwidths to more receiving computers. Significant work is required to write the code for a software correlator initially for a single receiving computer and eventually to aggregate and correlate data from multiple receiving computers. This software correlator will ultimately automate the additional step of finding beamformer weight vectors from the correlations. Initially beamformer coefficients will be written to BRAM manually but eventually the process could be streamlined. Eventually a convenient user-interface will need to be developed for the final instrument to be more useful for observations in the field.

Bibliography

- [1] CASPER, “Casper – collaboration for astronomy signal processing and electronics research,” <https://casper.berkeley.edu/>, May 2011. 2, 118
- [2] J. D. Kraus, *Radio Astronomy, Second Ed.* Powell, Ohio: Cygnus-Quasar Books, 1986. 3, 8, 19, 20, 23
- [3] S. Snezana and NAIC-NRAO Summer School on Single Dish Radio Astronomy, *Single-Dish Radio Astronomy : Techniques and Applications : Proceedings of the NAIC-NRAO Summer School.* San Francisco: Astronomical Society of the Pacific, 2002. 3, 8
- [4] A. Thompson, J. Moran, and G. Swenson Jr., *Interferometry and Synthesis in Radio Astronomy, Second Edition.* New York: Wiley-Interscience, 2001. 3, 8
- [5] R. Perley, F. Schwab, and A. Bridle, Eds., *Synthesis Imaging in Radio Astronomy, A collection of Lectures from the Third NRAO Synthesis Imaging Summer School.* Astronomical Society of the Pacific, 1989, vol. 6. 3
- [6] G. Taylor, C. Carilli, and R. Perley, Eds., *Synthesis Imaging in Radio Astronomy II : A Collection of Lectures from the Sixth NRAO/NMIMT Synthesis Imaging Summer School.* Astronomical Society of the Pacific, 1999. 3
- [7] NRAO, “National Radio Astronomy Observatory Archives: NRAO Timeline,” <http://www.nrao.edu/archives/Timeline/timeline.shtml>, Feb 2009. 3
- [8] Wikipedia, “Arecibo Observatory,” http://en.wikipedia.org/wiki/Arecibo_Observatory, June 2011. 3
- [9] International SKA Steering Committee (ISSC), “SKA, square kilometre array home page,” <http://www.skatelescope.org/>, 2008. 4, 9, 84
- [10] L. Staveley-Smith *et al.*, “The Parkes 21-cm multibeam receiver,” *Publications of the Astronomical Society of Australia, PASA*, vol. 13, p. 243, 1996, see also <http://www.atnf.csiro.au/research/multibeam/>. 4, 8
- [11] 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. Schröder, F. Stootman, R. L. Webster, W. E. Wilson, and T. Ye, “The HI Parkes All Sky Survey: southern observations, calibration

and robust imaging,” *Monthly Notices of the Royal Astronomical Society*, vol. 322, pp. 486–498, Apr. 2001. 4

- [12] 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,” *The Astrophysical Journal*, vol. 637, no. 1, p. 446, 2006. [Online]. Available: <http://stacks.iop.org/0004-637X/637/i=1/a=446> 4
- [13] R. Minchin, “ALFA: Arecibo L-Band Feed Array,” <http://www.naic.edu/alfa/>, Sep 2009. 4, 8
- [14] 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/i=3/a=1154> 5, 8, 9, 10, 39, 52, 58, 74, 75, 77, 78, 79, 80, 105
- [15] ATNF, “ASKAP Technologies: Antennas,” <http://www.atnf.csiro.au/projects/askap/antennas.html>, Jun. 2011. 7
- [16] G. Langston, “Green Bank Telescope K-band (18-26.5 GHz) Focal Plane Array,” <https://safe.nrao.edu/wiki/bin/view/Kbandfpa/WebHome>, April 2011. 8
- [17] S. R. Dicker, P. M. Korngut, B. S. Mason, P. A. R. Ade, J. Aguirre, T. J. Ames, D. J. Benford, T. C. Chen, J. A. Chervenak, W. D. Cotton, M. J. Devlin, E. Figueroa-Feliciano, K. D. Irwin, S. Maher, M. Mello, S. H. Moseley, D. J. Tally, C. Tucker, and S. D. White, “MUSTANG: 90 GHz science with the Green Bank Telescope,” in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 7020, Aug. 2008, pp. 702 005–702 005–9. 8
- [18] J. R. Fisher, R. F. Bradley, R. P. Escoffier, and K. S. Saini, “Studies and Experiments on Phased Array Feeds,” in *Workshop on High Sensitivity Radio Astronomy*, Jodrell Bank, England, January 1996. 8, 33
- [19] ———, “Phased Array Feed Design and Prototype,” in *URSI General Assembly and Scientific Symposium of International Union of Radio Science*, Lille, France, Aug 1996. 8, 33
- [20] J. Fisher and R. Bradley, “Full sampling array feeds for radio telescopes,” *Proceedings of the SPIE, Radio Telescopes*, vol. 4015, pp. 308–318, 2000. 8, 33, 83
- [21] J. R. Nagel, K. F. Warnick, B. D. Jeffs, J. R. Fisher, and R. Bradley, “Experimental verification of radio frequency interference mitigation with a focal plane array feed,” *Radio Science*, vol. 42, p. RS6013, Dec. 2007. 8, 9, 12, 83

- [22] J. Landon, B. Jeffs, K. Warnick, J. Fisher, and R. Norrod, “Improved low INR interference cancellation using phased array feeds,” *Proc. URSI General Assembly, Chicago*, Aug 2008. [Online]. Available: <http://ursi-test.intec.ugent.be/files/URSIGA08/start.pdf> 8, 9, 12
- [23] J. Landon, D. Jones, B. Jeffs, K. Warnick, R. Fisher, and R. Norrod, “Interference cancellation and sensitivity optimization using an L-band focal plane array on the Green Bank 20m Telescope,” in *URSI National Radio Science Meeting*, Boulder, CO, Jan 2008. 8, 9, 12, 59, 60, 69
- [24] K. F. Warnick, B. D. Jeffs, J. Landon, J. Waldron, D. Jones, R. Fisher, and R. Norrod, “BYU/NRAO 19-Element Phased Array Feed Modeling and Experimental Results,” *Proc. URSI General Assembly, Chicago*, Aug 2008. [Online]. Available: <http://ursi-test.intec.ugent.be/files/URSIGA08/start.pdf> 8, 9, 10, 51
- [25] B. Jeffs, K. 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>. 8, 9, 10, 51
- [26] B. Jeffs, K. Warnick, J. Landon, M. Elmer, J. Fisher, and R. Norrod, “Recent progress with the BYU/NRAO phased array feed,” in *Calibration and Imaging Workshop, CALIM2009*, Socorro, New Mexico, Mar. 2009, <https://safe.nrao.edu/wiki/pub/Software/CalIm09Program/brian.jeffs.pdf>. 8, 9, 10, 51
- [27] R. Maaskant, M. V. Ivashina, R. Mittra, and N. T. Huang, “Parallel FDTD modeling of a focal plane array with Vivaldi elements on the highly parallel LOFAR BlueGene/L super-computer,” in *Proc. IEEE Antennas and Propagation Symposium*, Albuquerque, NM, June 20-25 2006, pp. 3861 – 3864. 8, 83
- [28] 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*, ser. American Institute of Physics Conference Series, R. Minchin & E. Momjian, Ed., vol. 1035, Aug. 2008, pp. 265–271. 8, 9, 10, 33, 51
- [29] M. Ivashina, O. Iupikov, R. Maaskant, W. 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. 8
- [30] B. Veidt and P. Dewdney, “A phased-array feed demonstrator for radio telescopes,” in *Proc. URSI General Assembly, 2005*. 8, 84
- [31] 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. 8, 33
- [32] B. Van Veen and K. 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. 9, 27, 83, 87
- [33] H. Van Trees, *Detection, Estimation and Modulation, Part IV: Optimum Array Processing*. John Wiley and Sons, 2002. 9, 19, 27, 28, 83, 87
- [34] 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. 9, 11, 12, 80, 83, 84, 87
- [35] 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. 9, 51
- [36] 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>. 9, 10, 50, 51, 56
- [37] K. Warnick and B. Jeffs, “Efficiencies and system temperature for a beamforming array,” *IEEE Antennas and Wireless Propagation Letters*, vol. 7, pp. 565–568, 2008. 10, 62
- [38] K. Warnick, M. Ivashina, R. Maaskant, and B. Woestenburger, “Unified definitions of efficiencies and system noise temperature for receiving antenna arrays,” *IEEE Transactions on Antennas and Propagation*, vol. 58, no. 6, pp. 2121–2125, June 2010. 10, 61, 62, 74
- [39] K. F. Warnick and M. A. Jensen, “Optimal Noise Matching for Mutually Coupled Arrays,” *IEEE Transactions on Antennas and Propagation*, vol. 55, pp. 1726–1731, June 2007. 10
- [40] M. V. Ivashina, R. Maaskant, and B. Woestenburger, “Equivalent System Representation to Model the Beam Sensitivity of Receiving Antenna Arrays,” *IEEE Antennas and Wireless Propagation Letters*, vol. 7, pp. 733–737, 2008. 10
- [41] E.E.M. Woestenburger, “Noise matching in dense phased arrays,” ASTRON, Dwingeloo, The Netherlands, Tech. Rep. RP-083, Aug. 2005. 10
- [42] K. Warnick, B. Woestenburger, L. Belostotski, and P. Russer, “Minimizing the noise penalty due to mutual coupling for a receiving array,” *IEEE Transactions on Antennas and Propagation*, vol. 57, no. 6, pp. 1634–1644, June 2009. 10, 82, 151
- [43] K. Warnick and M. 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. 10, 83

- [44] K. Warnick, B. Jeffs, J. Landon, J. Waldron, D. Jones, J. Fisher, and R. Norrod, “Phased array antenna design and characterization for next-generation radio telescopes,” in *Antenna Technology, 2009. iWAT 2009. IEEE International Workshop on*, Mar. 2009, pp. 1–4. 10
- [45] 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. 10, 71, 82, 83, 151
- [46] Y. Bhattacharjee, “Radio astronomers take arms against a sea of signals,” *Science*, vol. 330, no. 6003, pp. 444–445, 2010. [Online]. Available: <http://www.sciencemag.org/content/330/6003/444.short> 11
- [47] J. R. Fisher, personal communication, Feb. 2011. 11
- [48] C. Beaudet, personal communication, Feb. 2011. 11
- [49] S. W. Ellingson, J. D. Bunton, and J. F. Bell, “Removal of the GLONASS C/A Signal from OH Spectral Line Observations Using a Parametric Modeling Technique,” *The Astrophysical Journal Supplement Series*, vol. 135, pp. 87–93, July 2001. 11, 84
- [50] A. J. Poulsen, B. D. Jeffs, K. F. Warnick, and J. R. Fisher, “Programmable Real-Time Cancellation of GLONASS Interference with the Green Bank Telescope,” *The Astronomical Journal*, vol. 130, pp. 2916–2927, Dec. 2005. 11, 46
- [51] W. L. Combrinck, M. E. West, and J. J. Gaylard, “Coexisting with GLONASS: Observing the 1612 MHz hydroxyl line,” *Publications of the Astronomical Society of the Pacific*, vol. 106, pp. 807–812, July 1994. 11, 84
- [52] W. Dong, B. D. Jeffs, and J. R. Fisher, “Radar interference blanking in radio astronomy using a Kalman tracker,” *Radio Science*, vol. 40, p. RS5S04, June 2005. 11
- [53] B. D. Jeffs, W. Lazarte, and J. R. Fisher, “Bayesian detection of radar interference in radio astronomy,” *Radio Science*, vol. 41, p. RS3006, June 2006. 11
- [54] J. Fisher, “Analysis of Radar Data from February 6, 2001,” National Radio Astronomy Observatory, Green Bank, WV, Tech. Rep., 2001. 11, 84
- [55] S. W. Ellingson and G. A. Hampson, “Mitigation of Radar Interference in L-Band Radio Astronomy,” *The Astrophysical Journal Supplement Series*, vol. 147, pp. 167–176, July 2003. 11
- [56] Q. Zhang, Y. Zheng, S. G. Wilson, J. R. Fisher, and R. Bradley, “Combating Pulsed Radar Interference in Radio Astronomy,” *The Astronomical Journal*, vol. 126, pp. 1588–1594, Sept. 2003. 11, 84
- [57] —, “Excision of Distance Measuring Equipment Interference from Radio Astronomy Signals,” *The Astronomical Journal*, vol. 129, pp. 2933–2939, June 2005. 11

- [58] J. R. Fisher, Q. Zhang, Y. Zheng, S. G. Wilson, and R. F. Bradley, “Mitigation of pulsed interference to redshifted HI and OH observations between 960 and 1215 Megahertz,” *The Astronomical Journal*, vol. 129, pp. 2940–2949, June 2005. 11, 84
- [59] A. Leshem, A.-J. van der Veen, and E. Deprettere, “Detection and blanking of GSM interference in radio-astronomical observations,” in *Signal Processing Advances in Wireless Communications, 1999. SPAWC '99. 1999 2nd IEEE Workshop on*, 1999, pp. 374–377. 11
- [60] A. Boonstra, A. Leshem, A. van der Veen, A. Kokkeler, and G. Schoonderbeek, “The effect of blanking of TDMA interference on radio-astronomical observations: experimental results,” in *Acoustics, Speech, and Signal Processing, 2000. ICASSP '00. Proceedings. 2000 IEEE International Conference on*, vol. 6, 2000, pp. 3546–3549. 11, 12, 84
- [61] Committee on Radio Astronomy Frequencies, “Literature on interference reduction methods,” <http://www.craf.eu/antemi.htm>, June 2005. 11
- [62] S. W. Ellingson, “Report on techniques for mitigation of radio frequency interference in radio astronomy,” in *IAU Working Group on RFI Mitigation*, IAU GA, Prague, Aug 2006. [Online]. Available: http://159.226.88.6/IAU_COM40/WG/rfi_iau_v1.pdf 11
- [63] A. Leshem, A.-J. van der Veen, and A.-J. Boonstra, “Multichannel interference mitigation techniques in radio astronomy,” *Astrophysical Journal Supplements*, vol. 131, no. 1, pp. 355–374, 2000. 11, 12, 80, 83, 84, 87
- [64] 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. 11, 12, 80, 83, 84, 87, 93, 108, 109
- [65] M. A. Lillrose, “High-speed data acquisition and FPGA detected pulse blanking system for interference mitigation in radio astronomy,” Master’s thesis, Brigham Young University, Provo, UT, 2007. 11
- [66] B. Jeffs, L. Li, and K. Wanick, “Auxiliary antenna assisted interference mitigation for radio astronomy arrays,” *IEEE Transactions on Signal Processing*, vol. 53, no. 2, pp. 439–451, Feb. 2005. 11, 12
- [67] A. Leshem and A.-J. van der Veen, “Radio-astronomical imaging in the presence of strong radio interference,” *Information Theory, IEEE Transactions on*, vol. 46, no. 5, pp. 1730–1747, Aug. 2000. 12
- [68] B. Jeffs and K. Warnick, “Bias corrected PSD estimation for an adaptive array with moving interference,” *IEEE Transactions on Signal Processing*, vol. 56, no. 7, pp. 3108–3121, July 2008. 12, 84
- [69] —, “Spectral bias in adaptive beamforming with narrowband interference,” *IEEE Transactions on Signal Processing*, vol. 57, no. 4, pp. 1373–1382, April 2009. 12, 83, 84, 88, 91, 92, 112

- [70] T. Palka and D. Tufts, "Reverberation characterization and suppressions by means of principal components," *Proc. IEEE Conf. Oceans*, vol. 3, pp. 1501–1506, Sep 1998. 12, 80, 83
- [71] W. S. Youn and C. K. Un, "Robust adaptive beamforming based on the eigenstructure method," *Signal Processing, IEEE Transactions on*, vol. 42, no. 6, pp. 1543–1547, June 1994. 12, 80, 83
- [72] B. Friedlander, "A signal subspace method for adaptive interference cancellation," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 36, no. 12, pp. 1835–1845, Dec. 1988. 12, 80, 83
- [73] J. Raza, A.-J. Boonstra, and A.-J. van der Veen, "Spatial filtering of RF interference in radio astronomy," *Signal Processing Letters, IEEE*, vol. 9, no. 2, pp. 64–67, Feb 2002. 12
- [74] B. Yang, "Projection approximation subspace tracking," *IEEE Trans. Signal Processing*, vol. 43, pp. 95–107, Jan 1995. 12, 108, 109
- [75] K. Rohlfs and T. L. Wilson, *Tools of radio astronomy*. New York: Springer, 2000. 19, 22, 23
- [76] S. Haykin, *Adaptive Filter Theory, 3rd. Edition*. Prentice Hall, 1996. 19, 27, 28
- [77] R. Monzingo and T. Miller, *Introduction to Adaptive Arrays*. New-York: Wiley, 1980. 19, 27, 28
- [78] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave Remote Sensing, Active and Passive: Radar Remote Sensing and Surface Scattering and Emission Theory*. Artech House, 1986, vol. I & II. 20, 21, 24, 88
- [79] O. L. Frost, III, "An algorithm for linearly constrained adaptive processing," *Proc. IEEE*, vol. 60, pp. 926–935, Aug 1972. 27, 28
- [80] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *Proceedings of the IEEE*, vol. 57, no. 8, pp. 1408–1418, Aug. 1969. 27, 28
- [81] C. K. Hansen, "Beamforming techniques and interference mitigation using a multiple feed array for radio astronomy," Master's thesis, Brigham Young University, Provo, UT, 2004. 29, 69
- [82] J. R. Nagel, "A prototype platform for array feed development," Master's thesis, Brigham Young University, Provo, UT, 2007. 29, 39, 41
- [83] F. Ghigo, "Green Bank 20-meter Telescope," <http://www.gb.nrao.edu/fgdocs/20m/GB20m.html>, Oct 2000. 37
- [84] J. S. Waldron, "Nineteen-element experimental phased array feed development and analysis on effects of focal plane offset and beam steering on sensitivity," Master's thesis, Brigham Young University, Provo, UT, 2008. 39, 41, 57

- [85] A. J. Stemmons, “Radio frequency interference mitigation on the Very Small Array,” Master’s thesis, Brigham Young University, Provo, UT, 2009. 46
- [86] K. F. Warnick, D. Carter, T. Webb, J. Landon, V. Asthana, M. Elmer, R. Norrod, and J. R. Fisher, “Towards a high sensitivity, cryogenic phased array feed antenna for the Green Bank Telescope,” in *URSI General Assembly and Scientific Symposium of International Union of Radio Science*, Istanbul, Turkey, Aug 2011. 48
- [87] “IEEE standard definitions of terms for antennas,” IEEE Std 145-1993, p. i. 68
- [88] C. A. Balanis, *Antenna Theory*, 3rd ed. New York: John Wiley & Sons, 2005. 68
- [89] D. Jones, “Characterization of a phased array feed model,” Master’s thesis, Brigham Young University, Provo, UT, 2008. 69
- [90] —, “Relative Sensitivities of the SKA and Other Current or Near-Future Instruments,” SKA, SKA Memo 66, Oct. 2005. [Online]. Available: http://www.skatelescope.org/uploaded/9241_66_Jones.pdf 75, 76
- [91] R. Olsson, P.-S. Kildal, and S. Weinreb, “A novel low-profile log-periodic ultra wideband feed for the dual-reflector antenna of US-SKA,” in *Antennas and Propagation Society International Symposium, 2004. IEEE*, vol. 3, June 2004, pp. 3035–3038. 77
- [92] Y. Karandikar and P. Kildal, “Optimization of 200-800MHz eleven feed for GMRT,” in *Antennas and Propagation, 2007. EuCAP 2007. The Second European Conference on*, Nov. 2007, pp. 1–6. 77
- [93] J. Yang, “On conditions for constant radiation characteristics for log-periodic array antennas,” *Antennas and Propagation, IEEE Transactions on*, vol. 58, no. 5, pp. 1521–1526, May 2010. 77
- [94] A. Taylor *et al.*, “The Canadian Galactic Plane Survey,” *The Astronomical Journal*, vol. 125, no. 6, pp. 3145–3164, June 2003. 78, 79
- [95] M. Elmer and B. Jeffs, “Beamformer design for radio astronomical phased array feeds,” in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, March 2010, pp. 2790–2793. 82, 151
- [96] C. Granet, A. Grancea, A. R. Forsyth, and D. H. Hayman, “Focal plane array development for ASKAP (Australian SKA Pathfinder),” in *The Second European Conference on Antennas and Propagation, 2007. EuCAP 2007.*, CSIRO ICT Centre, Epping, NSW, Nov. 2007, pp. 1–5. 83
- [97] S. van der Tol, B. D. Jeffs, and A.-J. van der Veen, “Self-calibration for the LOFAR radio astronomical array,” *IEEE Transactions on Signal Processing*, [see also *Acoustics, Speech, and Signal Processing, IEEE Transactions on*], vol. 55, no. 9, pp. 4497–4510, Sep 2007. 84, 100, 101, 102

- [98] J. Rissanen, “Modeling by shortest data description,” *Automatica*, vol. 14, pp. 465–471, 1978. 87
- [99] T. F. Coleman and Y. Li, “An interior, trust region approach for nonlinear minimization subject to bounds,” *SIAM Journal on Optimization*, vol. 6, pp. 418–445, 1996. 96
- [100] ———, “On the convergence of reflective newton methods for large-scale nonlinear minimization subject to bounds,” *Mathematical Programming*, vol. 67, no. 2, pp. 189–224, 1994. 96
- [101] J. J. Moré and D. C. Sorensen, “Computing a trust region step,” *SIAM Journal on Scientific and Statistical Computing*, vol. 3, pp. 553–572, 1983. 96
- [102] M. A. Branch, T. F. Coleman, and Y. Li, “A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems,” *SIAM Journal on Scientific Computing*, vol. 21, no. 1, pp. 1–23, 1999. 96
- [103] R. Byrd, R. Schnabel, and G. Shultz, “Approximate solution of the trust region problem by minimization over two-dimensional subspaces,” *Mathematical Programming*, vol. 40, pp. 247–263, 1988. 96
- [104] S. J. Wijnholds and A.-J. van der Veen, “Effects of parametric constraints on the CRLB in gain and phase estimation problems,” *IEEE Signal Processing Letters*, vol. 13, no. 10, pp. 620–623, Oct 2006. 97
- [105] P. Stoica, B. Ottersten, M. Viberg, and R. L. Moses, “Maximum likelihood array processing for stochastic coherent sources,” *IEEE Transactions on Signal Processing*, vol. 44, no. 1, pp. 96–105, Jan 1996. 101
- [106] A.-J. Boonstra and A.-J. van der Veen, “Gain calibration methods for radio telescope arrays,” *IEEE Transactions on Signal Processing*, vol. 51, no. 1, pp. 25–38, Jan 2003. 102
- [107] X. Wang and H. V. Poor, “Blind multiuser detection: a subspace approach,” *Information Theory, IEEE Transactions on*, vol. 44, no. 2, pp. 677–690, Mar. 1998. 108
- [108] X. G. Doukopoulos and G. V. Moustakides, “Fast and stable subspace tracking,” *IEEE Transactions on Signal Processing*, vol. 56, no. 4, pp. 1452–1465, Mar 2008. 108
- [109] T. Chonavel, B. Champagne, and C. Riou, “Fast adaptive eigenvalue decomposition: A maximum likelihood approach,” *Signal Processing*, vol. 83, pp. 307–324, 2003. 108
- [110] P. Strobach, “Low-rank adaptive filters,” *IEEE Transactions on Signal Processing*, vol. 44, pp. 2932–2947, Dec 1996. 108
- [111] R. Badeau, B. David, and G. Richard, “Fast approximated power iteration subspace tracking,” *IEEE Transactions on Signal Processing*, vol. 53, pp. 2931–2941, Aug 2005. 108
- [112] CASPER, “64adcx64-12,” <https://casper.berkeley.edu/wiki/64ADCx64-12/>, Aug. 2010. 119
- [113] R. Raffanti, “Parallel Cross Strip ADC Board ADC Block,” <https://casper.berkeley.edu/wiki/images/b/b8/ADC64x12.block.pdf>, Sept. 2007. 119, 121

- [114] A. T. Deller, S. J. Tingay, M. Bailes, and C. West, “DiFX: A Software Correlator for Very Long Baseline Interferometry Using Multiprocessor Computing Environments,” *Publications of the Astronomical Society of the Pacific*, vol. 119, pp. 318–336, Mar. 2007. 132
- [115] A. T. Deller, “Precision VLBI astrometry: Instrumentation, algorithms and pulsar parallax determination,” *ArXiv e-prints*, Feb. 2009. 132
- [116] K. F. Warnick, personal communication, Mar. 2011. 132
- [117] C. Satten, “Lossless gigabit remote packet capture with linux,” <http://staff.washington.edu/corey/gulp/>, Mar. 2008. 137
- [118] J. Manley, “Casper packetized correlator,” https://casper.berkeley.edu/wiki/images/9/9e/Packetized_correlator_signal_flow.pdf, Oct. 2007. 143
- [119] CASPER, “Casper packetized correlator,” https://casper.berkeley.edu/wiki/CASPER_Correlator, Aug 2010. 143
- [120] —, “Corr,” <https://casper.berkeley.edu/wiki/Corr>, Mar. 2011. 144