

# Observational and Theoretical Advances in Cosmological Foreground Emission

Thesis by  
Matthew A. Stevenson

In Partial Fulfillment of the Requirements  
for the Degree of  
Doctor of Philosophy



California Institute of Technology  
Pasadena, California

2014  
(Defended July 12, 2013)

© 2014

Matthew A. Stevenson

All Rights Reserved

Dedicated to Jennifer, Bronson, and Edmund.

# Acknowledgments

I carry a debt of gratitude to many people.

First I must thank my supervisor Tony Readhead for his support of this work and his patience for the variety of crazy ideas I came up with (and sometimes discarded) along the way. I must also thank Tim Pearson, who acted as my mentor in the day-to-day grind of actually getting this work done. His patience was inexhaustible, despite my many false starts and dead ends.

Being a part of C-BASS has afforded me the pleasure of working with many smart people. Foremost is Oliver King, my analog counterpart, with whom I have enjoyed many frantic troubleshooting sessions. Jaya John John set me on the path of FPGA programming, which has been great fun. Clive Dickinson's insights and support of my efforts have been a huge inspiration. Stephen Muchovej's tireless work on the telescope, control system, and pipeline were indispensable. This project has also given me the privilege of working with Rachana Bhatawdekar, Roy Booth, Charles Copley, Rod Davies, Richard Davis, Paul Grimes, Yaser Hafez, Christian Holler, Melis Irfan, Justin Jonas, Dayton Jones, Mike Jones, Paddy Leahy, Jamie Leech, Erik Leitch, Angela Taylor, and Joe Zuntz.

Outside of the C-BASS collaboration, I have had the opportunity to work with many other fine people. I have been fortunate to benefit from the enthusiasm and insights of Kieran Cleary, Rohit Gawande, Peter Goldreich, Keith Grange, Branson Hensley, Chris Hirata, Talvikki Hovatta, Andrew Lange, Roberta Paladini, Vaso Pavlidou, Yvette Perrott, Rodrigo Reeves, Clare Rumsey, Anna Scaife, Martin Shepherd, Chris Tibbs, and Jackie Villadsen. This work has also benefited from data provided by John Gaustad and Wolfgang Reich.

This work would have been impossible without the incredible staff at the Owens Valley Radio Observatory. Russ Keeney's tireless work has been essential to C-BASS. The experiment would not exist without him. Dave Hawkin's constant willingness to sit down and talk about digital hardware has been invaluable. I also wish to thank Tomi Hovatta, James Lamb, and Dave Woody for their help and expertise. Visiting OVRO has always been a pleasure, due in no small part to the hospitality provided by Mary Daniel, Lori McGraw, Cecil Patrick, and Terry Sepsey.

I owe many thanks to the Cahill administrative staff. Nothing would happen without them. I am grateful for the steady help of Althea Keith, on whom I imposed most often. I also wish to thank Gina Armas, Sandy Lester, Judith Mack, and Gita Patel.

I must also thank my classmates, Yacine Ali-Haimoud, Walter Max-Moerbeck, Laura Perez, and Varun Bhalerao, with whom I started this adventure. They have each, at various times, been study-mates, idea-testers, and immense inspirations.

There are two students who were senior to me and acted as valued mentors. I am lucky to have received such help, advice, and solidarity in parenting from Joey Richards and Milan Bogosavljevic.

Many of my fellow students have had the burden of sharing an office with me. I offer thanks and apologies to Brian Cameron, Adam Kraus, Tim Morton, Krzysztof Findeisen, Matt Schenker, and Yi Cao.

I would not be here without my family. I thank my parents, Diane and Terry, Alan and Lorraine, for everything. I thank my grandparents, Helen and Victor, Ann and Bronson, for even more. I thank my sister Teri for doing what she does. I thank my parents in-law, Tim and Val, for their unwavering support.

Finally, I offer my greatest thanks my wife and children, to whom this work is dedicated. I hope that they find it worth all the late nights and overseas trips.

# Abstract

Observational and theoretical work towards the separation of foreground emission from the cosmic microwave background is described. The bulk of this work is in the design, construction, and commissioning of the C-Band All-Sky Survey (C-BASS), an experiment to produce a template of the Milky Way Galaxy's polarized synchrotron emission. Theoretical work is the derivation of an analytical approximation to the emission spectrum of spinning dust grains.

The performance of the C-BASS experiment is demonstrated through a preliminary, deep survey of the North Celestial Pole region. A comparison to multiwavelength data is performed, and the thermal and systematic noise properties of the experiment are explored. The systematic noise has been minimized through careful data processing algorithms, implemented both in the experiment's Field Programmable Gate Array (FPGA) based digital backend and in the data analysis pipeline. Detailed descriptions of these algorithms are presented.

The analytical function of spinning dust emission is derived through the application of careful approximations, with each step tested against numerical calculations. This work is intended for use in the parameterized separation of cosmological foreground components and as a framework for interpreting and comparing the variety of anomalous microwave emission observations.

# Contents

<b>Acknowledgments</b>	<b>iv</b>
<b>Abstract</b>	<b>vi</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 The C-Band All-Sky Survey</b>	<b>4</b>
2.1 Overview of C-BASS . . . . .	4
2.2 Scientific Background . . . . .	5
2.2.1 Foregrounds . . . . .	5
2.2.2 Galactic science . . . . .	7
2.3 Design Goals . . . . .	8
2.3.1 Frequency & Bandwidth . . . . .	8
2.3.2 Resolution . . . . .	8
2.3.3 Sensitivity . . . . .	8
2.3.4 Sky Coverage . . . . .	9
2.3.5 Polarization . . . . .	9
2.3.6 Systematic Effects . . . . .	9
2.4 Experiment Design . . . . .	10
2.4.1 Antennas & Optics . . . . .	10
2.4.2 Receivers . . . . .	11
2.4.3 Scanning . . . . .	15
2.4.4 Pipeline . . . . .	16
2.5 Northern Telescope Performance . . . . .	17
<b>3 The C-BASS Digital Backend</b>	<b>19</b>
3.1 Context . . . . .	19
3.2 Requirements . . . . .	19
3.3 Hardware . . . . .	21

3.4	Firmware . . . . .	25
3.4.1	Timing . . . . .	27
3.4.2	1PPS Handling . . . . .	28
3.4.3	Parameters, Registers, and Diagnostics . . . . .	30
3.4.4	Control Signals . . . . .	34
3.4.5	Data Acquisition Modes . . . . .	35
3.4.6	DSP Chain . . . . .	38
3.4.6.1	Multiplicative Corrections . . . . .	40
3.4.6.2	Demodulation . . . . .	45
3.4.6.3	Filtering and Integration . . . . .	45
3.4.7	USB Microcontroller . . . . .	51
3.5	Operation in the Field . . . . .	54
<b>4</b>	<b>The C-BASS Pipeline</b>	<b>55</b>
4.1	Context . . . . .	55
4.2	The MATLAB Pipeline . . . . .	56
4.2.1	Initial Flags and Plots . . . . .	57
4.2.2	$\alpha$ Corrections . . . . .	59
4.2.3	Cold Cycle Correction . . . . .	63
4.2.4	$r$ Factor Correction . . . . .	64
4.2.5	RFI Flagging . . . . .	64
4.2.6	Astronomical Calibration . . . . .	68
4.2.7	Summary Plots . . . . .	73
4.2.8	FITS Output . . . . .	74
4.3	The Descart Map Maker . . . . .	75
<b>5</b>	<b>The C-BASS Cold Cycle Correction</b>	<b>82</b>
5.1	Context . . . . .	82
5.2	Assumptions . . . . .	83
5.3	Parameter Measurement . . . . .	84
5.4	Template Estimation . . . . .	86
5.5	Template Fitting . . . . .	88
5.6	Discussion . . . . .	88
<b>6</b>	<b>The C-BASS RFI Rejection Strategy</b>	<b>93</b>
6.1	Context . . . . .	93
6.2	Analog Filtering . . . . .	94



6.3	Spatial Flagging . . . . .	94
6.4	Statistical Flagging . . . . .	94
6.4.1	Assumptions . . . . .	96
6.4.2	Flagging Procedure . . . . .	97
6.4.3	Parameter Tuning . . . . .	102
6.5	Discussion . . . . .	105
<b>7</b>	<b>The North Celestial Pole with C-BASS</b>	<b>108</b>
7.1	Context . . . . .	108
7.2	C-BASS Data . . . . .	109
7.3	C-BASS Analysis . . . . .	113
7.4	Multiwavelength Comparison . . . . .	126
7.5	Discussion . . . . .	130
<b>8</b>	<b>The Spinning Dust Emission Spectrum</b>	<b>132</b>
8.1	Context . . . . .	132
8.2	Overview . . . . .	134
8.3	Dust Grains . . . . .	135
8.3.1	Size . . . . .	135
8.3.2	Shape . . . . .	137
8.3.3	Temperature . . . . .	137
8.3.4	Charge . . . . .	138
8.3.5	Dipole Moment . . . . .	138
8.4	Distribution Function . . . . .	139
8.5	Emissivity . . . . .	143
8.5.1	Grain Tumbling . . . . .	144
8.5.2	Grain Emissivity . . . . .	146
8.5.3	Total Emissivity . . . . .	147
8.6	Discussion . . . . .	152
<b>9</b>	<b>Conclusion</b>	<b>157</b>

# Chapter 1

## Introduction

Observation of the cosmic microwave background radiation is humankind’s most powerful probe of the cosmology of the Universe. The cosmic microwave background (CMB) was discovered less than 50 years ago (Penzias & Wilson 1965), and has since been shown to be the most perfect blackbody ever observed (Mather et al. 1990), to be isotropic to one part in  $10^5$  (Smoot et al. 1992; Lange et al. 2001), and to have polarization at one part in  $10^6$  (Kovac et al. 2002; Readhead et al. 2004). Studies of the anisotropy and polarization of the CMB continue to allow very precise measurements of the Universe’s fundamental cosmological parameters (Hinshaw et al. 2012; Planck Collaboration 2013a). It is remarkable that, in less than 50 years, observations of this radiation have advanced from discovery to detailed interrogation.

It is commonly held that the next steps in the study of the CMB will be in constraining or measuring primordial B-mode polarization and primordial non-Gaussian statistics in temperature anisotropy (Bock et al. 2005). These effects are key predictions of inflationary cosmology models, which postulate that the Universe underwent an early, explosive phase of expansion. Inflation can be thought of as the “Bang!” of the Big Bang, itself. Such a phase can account for the Universe’s remarkable flatness and uniformity, as well as the curious absence of magnetic monopole particles. Indeed, a tantalizing hint of inflationary models already exists in the form of the slope of the primordial density power spectrum,  $n_s$  (Planck Collaboration 2013c). It is the goal of testing these models that inspires current and future CMB experiments.

There is a great deal of Universe between us and the CMB surface of last scattering. Discerning primordial CMB from foreground effects is a challenge, and it is one that will only get harder as we seek to measure more subtle CMB features. Often, foreground effects are best separated from the CMB through multi-wavelength modeling, which relies upon the different spectral behaviours of the foreground emission and the CMB (Planck Collaboration 2013b). Synchrotron emission from the Milky Way Galaxy is a key component of this foreground emission, requiring low frequency measurements to constrain it. Also at low frequencies is an anomalous component of microwave emission, discovered by Leitch et al. (1997) and often attributed to spinning dust grains (Draine & Lazarian

1998a), which is comparable in magnitude to synchrotron emission at centimetre wavelengths (Lu et al. 2012).

The C-Band All-Sky Survey (C-BASS) is an observational experiment designed to address this need. The aim is to produce an all-sky map at 5 GHz in both total intensity and linear polarization. The 5 GHz frequency is low enough that synchrotron emission should dominate the map over most the sky, but high enough that Faraday rotation will not corrupt the polarization information.

Low frequency, all-sky maps of the Galaxy’s synchrotron emission already exist at 408 MHz and 1.4 GHz (Haslam et al. 1982; Reich 1982). These maps have found use in correcting foreground emission in total intensity, though they have deficiencies in terms of systematic errors. There are also polarization data at 1.4 GHz, although they suffer both from Faraday rotation and Faraday depolarization (Wolleben et al. 2006; Testori et al. 2008). Beyond these, the WMAP 23 GHz map is the lowest frequency, all-sky map (Bennett et al. 2012). However, by 23 GHz the synchrotron emission is confused with free-free emission and anomalous microwave emission, and the polarization sensitivity of this map is insufficient at high Galactic latitudes, limiting the use of this map in correcting polarized synchrotron emission in future CMB surveys. So there is need for a new survey to come into the breach.

The strength of the B-modes polarization is unknown, but constraining it to the order of 10 nK will delve deeply into the inflationary parameter space. Assuming cosmological observations at 100 GHz and a synchrotron power-law of around  $-3$ , then a sufficient synchrotron measurement requires 0.1 mK sensitivity at 5 GHz. Reaching this sensitivity over the whole sky is not easy, particularly when considering that cosmological studies require preservation of emission up to the largest spatial scales. Great effort is needed to reach low system temperatures and to avoid systematic effects.

The bulk of the work in this thesis has been aimed at measuring and minimizing systematic effects in the C-BASS experiment. This work has focused on the digital signal processing in the instrument and novel algorithms in the data pipeline. The digital signal processing was written for the instrument’s digital backend and runs on the backend’s field programmable gate array (FPGA). This work took place during a crucial time for digital electronics in radio astronomy, as technological advances in FPGAs and analog-to-digital converters (ADCs) have begun to allow wide-bandwidth, analog circuits to be replaced by relatively inexpensive digital alternatives (Parsons et al. 2009). The new algorithms in the pipeline have taken advantage of the continuing, explosive growth in computing power.

A preliminary, limited-area survey of the North Celestial Pole region, the discovery field for the anomalous microwave emission (Leitch et al. 1997), was also performed as part of this thesis. This test survey was designed to allow multi-wavelength characterization of the anomalous microwave emission in this field and to distinguish between flat spectrum and peaked spinning dust models.

The collection and reduction of these data provided the additional benefit of allowing for practical tests of the experiment before the full survey data is collected.

The separation of foreground components from the CMB is aided by parameterized descriptions of each. An analytical parameterization of the spectrum of spinning dust emission was derived as part of this thesis. This new description of the spinning dust model combines the advantages of a small number of adjustable parameters with the benefits of a physically motivated form. Further, its derivation provides a framework for an intuitive understanding of the spinning dust spectrum.

A detailed overview of the C-BASS experiment is provided in Chapter 2, which provides the context for the discussions of the instrument's digital backend, pipeline, and data analysis in following chapters. Chapter 3 describes the experiment's digital backend, a major portion of this thesis, with focus on its novel timing and filtering algorithms. A broad survey of the C-BASS data pipeline is given in Chapter 4. This leads into thorough descriptions of the cold cycle correction and radiofrequency interference (RFI) mitigation strategy in Chapters 5 and 6; the algorithms described in these chapters are wholly due to the work of the author. The NCP survey and analysis is described in Chapter 7 and is one part of the science focus of this work. The other major science result is a contribution to spinning dust theory in Chapter 8, which is a modified version of Stevenson (2013). A concluding discussion of the full body of work can be found in Chapter 9.

## Chapter 2

# The C-Band All-Sky Survey

### 2.1 Overview of C-BASS

The C-Band All-Sky Survey is an experiment to map the Galaxy's polarized synchrotron emission at 5 GHz. When complete, it will form the highest fidelity map of this emission yet produced. This will allow dramatic advances in the study of the Galaxy's cosmic ray electrons and magnetic field. As this emission forms an important foreground for studies of the polarized cosmic microwave background radiation, this map will also be of great use to future experiments to map the polarized cosmic microwave background.

To achieve the required sensitivity we must minimize a number of systematic effects, including atmospheric and instrumental correlated noise, thermal radiation from the Earth, radio frequency interference from unrelated transmitters, and low frequency interference from the instrument itself. These are not small challenges, but have inspired a number of novel approaches and techniques.

Solving these issues has led to an experiment with several novel design elements. Instrumental effects have been minimized through symmetries in the optics and receiver, atmospheric effects through scanning and mapping strategies, and interference through frequency- and time-domain filtering. Through these efforts, the above issues have been significantly mitigated.

This amount of work has required a large team. The C-BASS collaboration, spanning four continents and six institutions, contains more than twenty members. Combining our varied expertise in radiofrequency optics, radiotelescope control, analog and digital electronics, survey design, and radio data analysis, we have been able to attack these problems.

## 2.2 Scientific Background

### 2.2.1 Foregrounds

Study of the cosmic microwave background (CMB) allows us to study the physics of the early Universe with unprecedented detail (Hinshaw et al. 2012; Planck Collaboration 2013a). Astounding results on the composition and geometry of the Universe have already resulted, but many questions remain. In particular, the initial conditions in the  $\Lambda$ CDM framework pose a substantial challenge. One must ask why the Universe is so exquisitely flat, why it is so incredibly uniform, and why there are no observed magnetic monopoles (Bock et al. 2005). The commonly postulated answer is that the early Universe existed in a state of rapidly accelerating expansion, called inflation. In this model, a single Planck length was expanded beyond the size of the Hubble volume, stretching out any non-flat geometry, ensuring that the initial conditions were uniform over the Hubble volume, and decreasing the density of magnetic monopoles to unobservable levels.

The inflationary model already boasts a confirmed prediction in the primordial matter power spectrum  $n_s$ , predicted to be less than unity and measured as such (Planck Collaboration 2013c). This has limited power to constrain inflationary models, however, and so further observational handles are needed. The polarization of the CMB may provide such a handle in the form of the B mode polarization. This signal is predicted to arise from a primordial gravitational wave background excited during the inflationary period. In interacting with the surface of last scattering, these tensor waves ought to introduce a polarization signal with both curl-free and divergence-free components. The former, known as the E-mode, is buried under polarization that occurs due to density perturbations. The latter, called the B-mode, is not predicted to occur from any other primordial cause, and so it provides a potential means of measuring the gravitational wave background at that time and thus constrain inflationary models.

Nothing is ever easy. Observing the CMB requires looking past (and, in many cases, through) the entire observable Universe. Although very beautiful, the Universe thus acts as an undesirable foreground to the cosmological signal of interest. For temperature and E-mode measurements, these effects are of lesser importance. For the B-modes, however, the foreground effects represent a major impediment. In particular, polarized synchrotron and dust emission from the Milky Way Galaxy are expected to dominate the B-mode signal by an order of magnitude in the most optimistic cases (see Figure 2.1). These foregrounds need to be characterized and removed if there is to be any hope of reaching the B-mode signal.

Fortunately, there is a path forward. The spectral and spatial shapes of the foreground emission are substantially different from those of the primordial B-mode signal. While the CMB has a purely thermal spectrum and has a spatial structure consistent with Gaussian noise, the Galactic signals have falling (synchrotron) or rising (dust) spectra and are distributed according to local interstellar

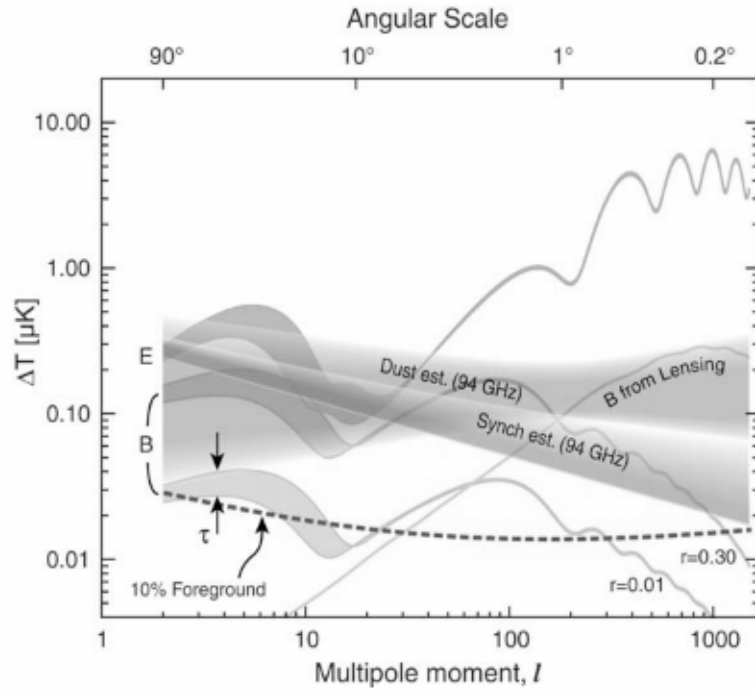


Figure 2.1: Predicted polarization power spectra of the CMB. E-modes, already detected, are shown as the upper curve. Two plausible predictions for the primordial B-modes are shown below. Non-primordial B-modes from gravitational lensing are shown peaking at  $l \sim 1000$ . Polarized, Galactic foregrounds are illustrated by the two broad curves. The dashed line illustrates the probable foreground level if reduced by a factor of 10, which would require C-BASS to accomplish. Figure reproduced from Bock et al. (2005).

processes. Separation of the CMB from foreground emission can therefore be achieved via cross-analysis of maps spanning a range of observing frequencies.

Cosmology experiments, such as Planck, reflect this in their design, though there are limits in how wide a range of frequencies can be covered by a single telescope. Extending this frequency range will improve the foreground subtraction and thus the sensitivities of cosmological efforts. C-BASS will do just this. By observing at a far lower frequency than is possible with Planck, we aim to provide a much stronger handle on the polarized synchrotron emission and so facilitate a much deeper exploration of the primordial B-mode polarization.

### 2.2.2 Galactic science

The interstellar medium of the Galaxy consists of many constituents: neutral, molecular, and ionized gas, relativistic cosmic rays, silicate and carbonaceous dust, the interstellar radiation field, and the Galaxy's magnetic field. These can account for the observed characteristics of the Galaxy, yet many details have escaped capture.

The cosmic ray population interacts with the Galactic magnetic field to produce the synchrotron emission, so study of the synchrotron emission allows us to infer the properties of these components. Basic questions remain unsettled, such as the structure of the magnetic field and its relationship to Galactic structure, and the distribution of the cosmic rays and how well they propagate through the Galaxy. A great deal of data have already been brought to bear on these issues (Jaffe et al. 2013; Ackermann et al. 2012), yet significant degeneracies continue to exist between model parameters. These matters stand to be dramatically improved by new, high sensitivity, polarization and temperature maps.

Leitch et al. (1997) discovered that the Galaxy contains an anomalous emission component at centimetre wavelengths, inconsistent with the well-known synchrotron, free-free, and thermal dust emission components. They suggested that this new emission could be explained by flat-spectrum synchrotron emission or free-free emission from excessively hot gas. Draine & Lazarian (1998b) countered that the emission is more naturally interpreted as electric dipole radiation from very small, very rapidly spinning dust grains. These different possibilities give very different predictions for low frequency radiation: spinning dust will peak at 30 GHz and fall off very rapidly by 5 GHz, while the free-free and synchrotron possibilities predict no such drop. Distinguishing between the latter two is also simple, as synchrotron emission will be polarized and free-free emission will not be.

It is clear, then, that C-BASS promises progress on several exciting avenues of Galactic research. Advances on the Galaxy's cosmic ray population, magnetic field, and smallest dust grains are all made possible.



## 2.3 Design Goals

### 2.3.1 Frequency & Bandwidth

The main aim of mapping the Galaxy’s polarized synchrotron emission determines the optimal frequency and bandwidth of observation. Optically thin synchrotron emission is brighter at lower frequencies, yet differential Faraday rotation across the received bandwidth will cause depolarization if too low of a frequency is observed. A wider bandwidth exacerbates the latter effect, yet provides more sensitivity.

Faraday rotation itself is a problem, even without depolarization. If a low-frequency, linearly polarized map is subject to significant Faraday rotation, then this rotation must be removed before the map can be used as a template for higher frequencies. Thus it is necessary to observe at a high enough frequency that neither Faraday depolarization nor Faraday rotation are significant.

An excellent compromise is found by using a 1 GHz bandwidth at 5 GHz. The chosen frequency, which happens to be the geometric mean between the existing 1.4 GHz and 23 GHz maps (Reich 1982; Reich & Reich 1986; Reich et al. 2001; Bennett et al. 2012), takes advantage of the brighter synchrotron emission at this frequency, while maintaining less than  $10^\circ$  of Faraday rotation across the whole sky.

### 2.3.2 Resolution

For any given map, higher spatial resolution is usually better. The catch is that higher resolution often comes at a price. Most obvious is the required aperture size: going to higher resolution requires having a larger telescope. This in turn imposes mechanical constraints on the scanning strategy: it is harder to scan a larger telescope at a high rate.

The first peak of the primordial B-mode power spectrum occurs at degree scales, and existing, all-sky, radio maps have had resolutions measured in tens of arcminutes. Thus, to achieve the primary aim of this experiment and to allow optimal comparison with existing maps, sub-degree resolution is required. A fully-illuminated aperture of 5 m provides a resolution of  $\sim 40'$ , which we consider adequate. It’s worth noting that at an angular diameter of  $30'$ , the Moon is barely resolved at this resolution.

### 2.3.3 Sensitivity

The primordial B-mode signal is predicted to be as bright as a tenth of a  $\mu\text{K}$  in the best case. Pushing well below this thus requires a foreground sensitivity better than a hundredth of a  $\mu\text{K}$ . As the synchrotron spectral index tends to be around  $-3$  (Bennett et al. 2012), moving to lower frequency gives an advantage. Assuming cosmological analysis is performed at 100 GHz, a synchrotron map

at 5 GHz need only reach a tenth of a mK. This then sets the sensitivity requirements for C-BASS.

### 2.3.4 Sky Coverage

Cosmological measurements at degree scales face a limitation in the number of independent, observable resolution elements. This limits the statistics and therefore the maximum significance of any given measurement. Cosmological experiments therefore typically attempt to measure as much of the sky as possible in order to minimize this issue. To be of greatest use, it is thus necessary that C-BASS cover as much of the sky as possible.

An all-sky map provides additional benefits. Study of the Galaxy is best done with all-sky data, as the entire Galaxy is not visible from either of the Northern or Southern hemispheres alone. Further, many wide area maps already exist, and so all-sky coverage allows for maximal exploitation of these resources.

The Earth is optically thick, so mapping the whole sky from the surface requires both Northern and Southern instrumentation, and the means to integrate data from the two resulting data sets together. This has prompted a careful optical design, so that the primary beams, sidelobes, and cross-polarization behaviours of the two antennas are well-matched.

### 2.3.5 Polarization

As described above, the primordial B-mode signal is a linear polarization feature. It is therefore the linearly polarized foreground emission that is of most importance. Total intensity signal is also of value in that it provides additional information of Galactic interest and an important diagnostic for the performance of the entire experiment. Circular polarization does not carry the same weight, however. Strong, circularly polarized emission is not predicted for the targeted science, and is not emphasized in the experiment design.

### 2.3.6 Systematic Effects

Reaching the sensitivity goals requires minimizing sources of systematic noise. A wide variety of these exist. The atmosphere introduces  $1/f$  noise through slow variations in opacity. The instrument's amplifiers also introduce  $1/f$  noise due to electrons interacting with semiconductor traps (Van der Ziel 1988), a known behaviour for high electron mobility transistor (HEMT) based amplifiers (Weinreb & Schleeh 2013). The beam shape can introduce artefacts in the map, through asymmetry, sidelobes, and cross-polarization features. Radiofrequency anthropogenic interference can be especially confounding, as it can masquerade as astronomical signal. Other sources of correlated noise are also of concern, such as the omnipresent 60 Hz electrical lines. Avoiding, correcting,




Figure 2.2: Photographs of the Northern (left) and Southern (right) telescopes. The Northern telescope has the receiver, subreflector assembly, and absorbing baffles installed.

or otherwise getting past these effects is not easy, but is necessary if we are to reach our sensitivity goals.

## 2.4 Experiment Design

### 2.4.1 Antennas & Optics

Covering the whole sky from the surface of the Earth requires a Northern map and a Southern map and thus two telescopes. Telescopes are expensive, and so we counted ourselves fortunate to receive both telescopes through donations. The Northern telescope, which is 6.1 m in diameter, was built by the Jet Propulsion Laboratory as a low-cost prototype for the Deep Space Network and donated for C-BASS use. The Southern telescope, which is 7.6 m in diameter, was donated by Telkom, a South African telecommunications company. The Northern telescope has been relocated to the Owens Valley Radio Observatory ( $-118^{\circ}17'00.2'' + 37^{\circ}13'59.9''$ ) and the Southern to the Karoo in South Africa ( $21^{\circ}19'16.7'' - 30^{\circ}42'22.7''$ ). These telescopes are shown in Figure 2.2.

With the telescopes' designs as different as their origins, nontrivial effort was needed to ensure matching beam and cross-polarization performance between the Northern and Southern data sets.




Figure 2.3: Simulated co- and cross-polar beams of the Northern (black) and Southern (grey) antennas. Figure reproduced from Holler et al. (2013).

The agreement was achieved through custom feedhorns and subreflectors. Through allowing these optical elements to be different between the two telescopes, effectively correcting for the differences in telescope shapes, it was possible to bring the resulting beam shapes and cross-polarization response into close agreement, as shown in Figure 2.3.

This freedom in optical design did result in qualitative differences between North and South. While the former telescope was optimized by a Gregorian reflector and absorbing baffles about the primary and secondary reflectors, the latter required a Cassegrain reflector and no baffles. These designs are illustrated in Figure 2.4. The most striking element of both designs, however, is that the support structure was designed for circular symmetry. Ordinarily, metal struts are used to support antenna subreflectors, which can affect beam shape and cross-polarization through diffraction and reflection. The C-BASS design instead supports the subreflectors atop radio-transparent foam cones, avoiding such sources of systematic error. Predictions and measurements of the Northern main beam and sidelobe response are shown in Figure 2.5.

## 2.4.2 Receivers

The Northern and Southern receivers are both designed to measure the Stokes parameters  $I$ ,  $Q$ , and  $U$ . They differ, however, in approach. The Northern receiver, which was designed earlier, produces the Stokes parameters via analog manipulation of the radiofrequency signal and digitizes post-detection. The Southern receiver, taking advantage of recent advances in ADC and FPGA technology, uses a heterodyne strategy to make the signal available to wide-bandwidth digitization. The Stokes parameters are calculated through digital signal processing, which also allows for spectral




Figure 2.4: Schematic illustrations of the C-BASS optics. Northern design on the left, Southern on the right. Geometric and diffractive rays are shown. Figure reproduced from Holler et al. (2013).




Figure 2.5: Measured beam response of the Northern antenna. Main beam is shown on the left. Solid line is the prediction, scatter plot shows measurement via scans across the radio source Cas A. Wide-angle, measured sidelobes are on the right, as measured using a far-field transmitter. Solid black line is without absorbing baffles, solid grey is with subreflector baffle, and dotted black is with both baffles installed. Versions of these figures appear in Holler et al. (2013).

resolution not available in the Northern design.

The Northern receiver is a novel combination of a continuous comparison radiometer and correlation polarimeter. The former aspect allows a temperature-controlled, reference load to be subtracted from the total intensity sky signal, dramatically reducing correlated, receiver noise in Stokes  $I$  measurements. The latter aspect minimizes correlated, receiver noise in the measurements of Stokes  $Q$  and  $U$ .

A schematic drawing of the Northern receiver is shown in Figure 2.6. The opening of the horn is at ambient temperature, while the base of it is cryogenically stabilized at 40 K. Linear polarization is received, but immediately converted to right- and left-circular polarization using the ortho-mode transducer. A warm noise diode can insert noise at this stage for calibration purposes. Analog hybrids are then used to add and subtract the reference signals (calibrated to match the sky temperature) prior to amplification. This ensures that the reference signals will be subject to the same additive and multiplicative amplifier noise as the sky signal.

Entering the temperature-controlled warm stage, the Stokes  $I$  signals (which are actually right- and left-circular polarization) are modulated. This modulation introduces symmetry in the system, as the line carrying sky plus load is periodically switched with the line carrying sky minus load. The separate sky and load signals are then recovered after passing through another hybrid. Detector diodes convert these RF signals to near-DC prior to digitization. Through such a symmetric design, the difference between sky and load can be expected to have an enormous reduction in correlated receiver noise. It should be clear that four channels are in need of digitization: the right- and left-circularly polarized sky signals, and the corresponding reference load for each.

The Stokes  $Q$  and  $U$  take different paths through the warm stage. Analog hybrids are used once again to separate the sky and load signals, but in this case the load signals are discarded. The sky signals (once again, right- and left-circular polarization) are then modulated and combined with each other through analog  $90^\circ$  and  $180^\circ$  hybrids to get  $Q$  and  $U$ , respectively. As in the Stokes  $I$  case, the signals of interest take the form of paired, modulated RF lines which are detected via diodes before digitization. In this case, eight channels must be digitized: symmetry in the calculations of  $Q$  and  $U$  requires two measurements of each, with each measurement represented by paired RF signals.

Assuming frequency-independent phase response and gain, this design allows us to meet our desired bandwidth of 1 GHz, which is challenging in the absence of high-speed digitization. However, as it provides no frequency resolution, it does leave us susceptible to narrow-band, radiofrequency interference from man-made transmitters. After deploying the receiver in Owens Valley, we found that this last point was significant and required additional bandpass and notch filtering following the receiver amplifiers to ameliorate it. As a consequence, we have had to settle for an effective receiver bandwidth of 487 MHz.

The 3 dB video bandwidths of the detector diodes were measured to be 150 kHz. The video

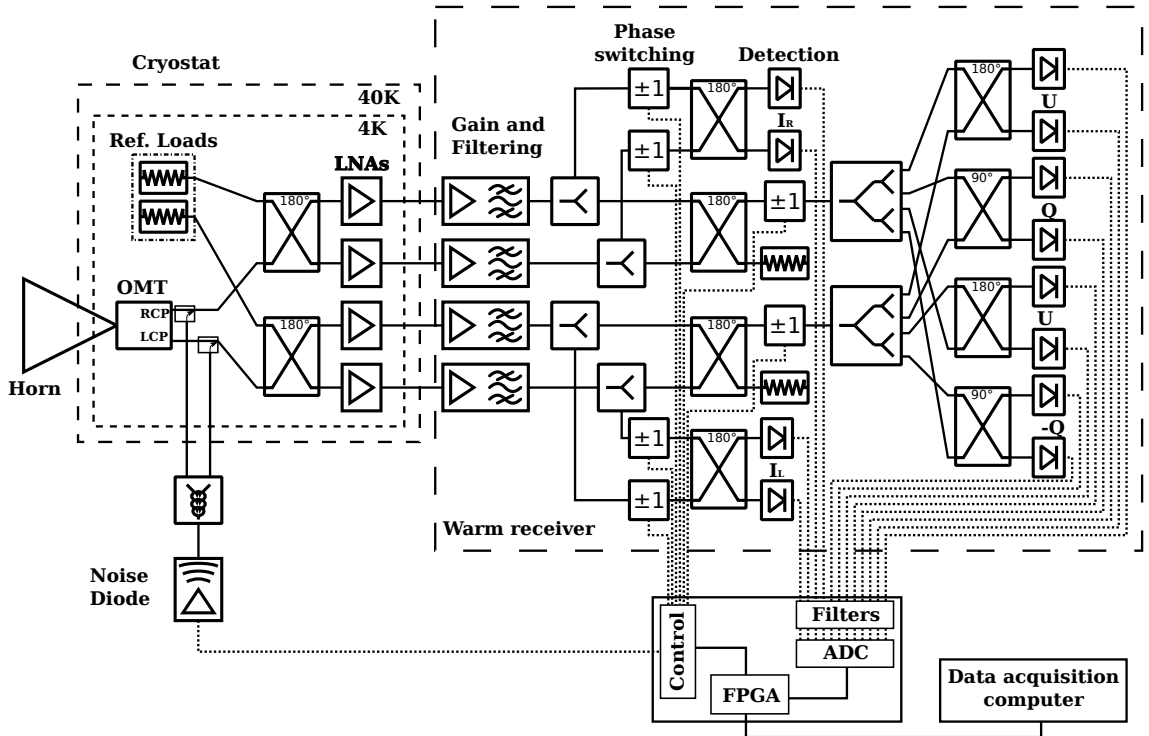


Figure 2.6: Schematic illustration of the Northern C-BASS receiver. Radiation enters the system through the feedhorn on the left and is immediately separated into right-circular and left-circular polarizations (RCP and LCP) by the ortho-mode transducer (OMT). A directional coupler allows the noise diode to inject noise into each line. The reference loads are coupled to the two polarizations through  $180^\circ$  hybrids prior to amplification by the low noise amplifiers (LNAs). Further amplification and filtering occurs in the warm receiver, as well as  $180^\circ$  phase switching in order to reduce systematic errors. The  $90^\circ$  and  $180^\circ$  hybrids combine the signals in various ways to linear polarization and both circular polarizations available. The radiofrequency power is converted to near-DC signals by the detector diodes before anti-alias filtering and digitization in the backend.

response at higher frequencies is not negligible, however, so the actual Nyquist rate is significantly greater than 300 kHz. The frequency response does fall off fast enough, though, that a rate faster than a few MHz is unnecessary. For reasons internal to the digital backend, it was convenient to set the digitization rate at 2 MHz. To prepare the detected signals for digitization, a combination anti-alias filter and amplification circuit were used. For each of the twelve channels, this consisted of a 1 MHz, low-pass filter and amplification to maximize use of the ADC dynamic range. Following this, the signals were finally ready for digitization.

The digital backend itself consists of twelve parallel ADCs, an FPGA, and a USB interface. The twelve channels were digitized independently, minimizing cross-contamination of signals, and piped directly into the FPGA. Upon arrival, the signals were demodulated, pairwise subtracted, filtered, and downsampled to the final data rate of 100 integrations per second. The data were then packaged and sent to the control computer for storage via USB. The programming of this FPGA constituted

a major portion of this thesis and is discussed in great detail in Chapter 3.

The Southern receiver is designed with a similar eye towards correlated noise reduction, but through digital rather than analog means. This translates to an identical cold stage design (including reference loads), but a post-amplification heterodyne stage to prepare for full bandwidth digitization of the IF. CASPER-designed ROACH boards (Parsons et al. 2009) are used for the signal processing, which calculates the Stokes  $I$ ,  $Q$ , and  $U$  directly from the IF channels. In performing these calculations digitally, it is natural to work in the frequency domain and so gain frequency resolution for free. This, in turn, obviates the need for RFI notch filters and thus widens the effective bandwidth of the receiver.

The Southern receiver did not form a substantial part of this PhD, and further discussion is beyond the scope of this thesis.

### 2.4.3 Scanning

The receiver has a single horn and thus a single pixel on the sky. Making an all-sky map requires moving this pixel about by scanning the telescope. Essential to this is the understanding that scanning a single-pixel telescope converts spatial information on the sky into time-series information in the observed data product. There are many ways to cover the whole sky, but minimizing systematic effects requires careful planning.

The two dominant effects are correlated noise from variable atmospheric opacity and from ground pickup through beam sidelobes. As atmospheric column depends strongly on telescope elevation, azimuthal scans at constant elevation are optimal. Scanning in azimuth will give significant variability in ground pickup, but repeated scans at a single elevation will have a repeatable ground pickup signal. Azimuth scans will still be susceptible to variations in atmospheric opacity, but these are due to spatial variations in water column and vary only slowly with time. Scanning rapidly, then, ensures that the correlated noise from the atmosphere appears at different frequencies than the sky signal in the observed, time series data. Rapid scanning gives the additional benefit of returning to individual sky positions faster than the atmosphere opacity is changing, and thus giving an additional handle on atmospheric noise separation.

Coverage of a single hemisphere is guaranteed by scanning at the elevation of the celestial pole. The resulting sky coverage is not uniform, however. The pole will receive a disproportionate amount of integration time, while declinations midway between the equator and pole will receive far less than their fair share. Furthermore, most sky positions will be crossed by scans at similar position angles on the celestial sphere, which is not optimal for atmospheric noise suppression. A compromise can be reached by scanning at a small number of different elevations. This spreads out the observing time more evenly and provide a wider variety of scan angles for given sky positions. Too many elevations lead to larger systematic effects, as described above, and we settled on scanning at just



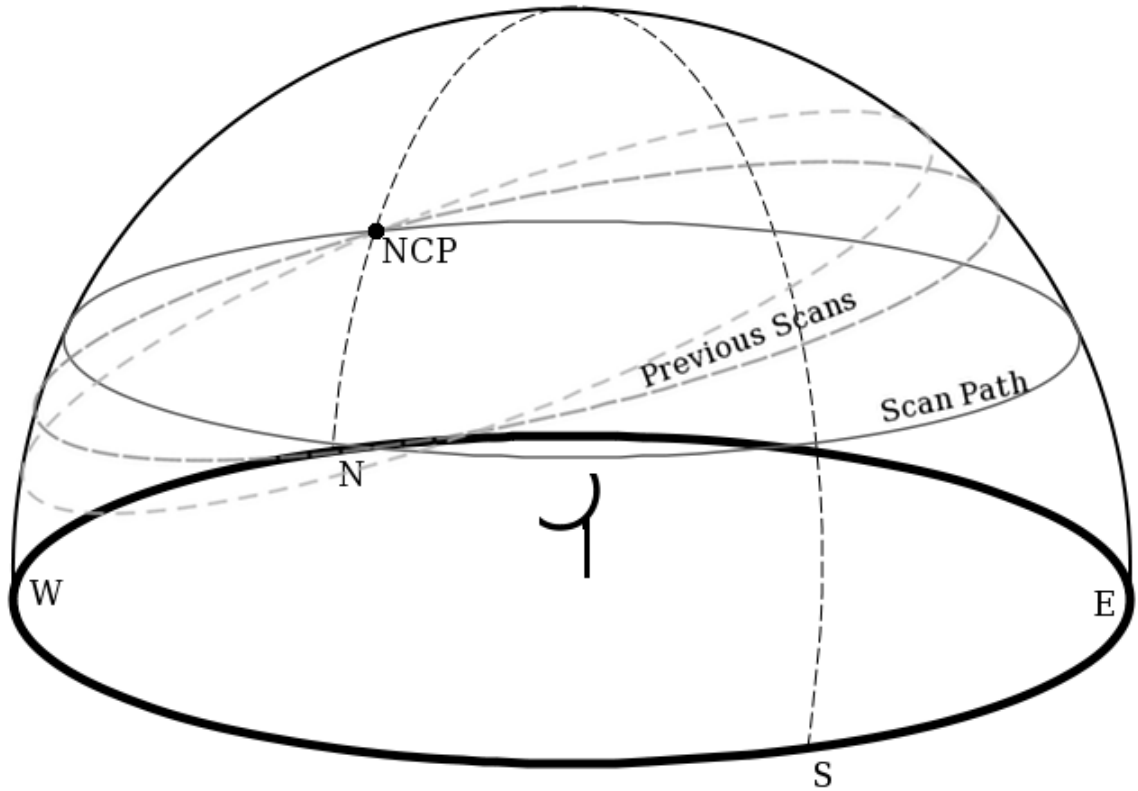


Figure 2.7: Illustration of the C-BASS scanning strategy. The telescope’s azimuth scans cover a circle on the sky. The Earth’s rotation causes successive circles to eventually cover the entire hemisphere.

two: the elevation of the celestial pole, and an elevation  $20^\circ$  above. This strategy is illustrated in Figure 2.7.

Simulation of this scanning strategy over a year of observing time shows that it gives 20s of observing time per  $27.5'$  sky pixel, which corresponds to an RMS noise level of 0.42 mK in Stokes  $I$  and 0.30 mK in Stokes  $Q$  and  $U$  (assuming  $T_{\text{sys}} = 30\text{ K}$ ). Simulated sky coverage is shown in Figure 2.8.

Making an all-sky map requires significant overlap between the Northern and Southern surveys. This scan strategy guarantees  $44^\circ$  overlap in declination, which is essential for cross-calibration efforts between the two data sets.

#### 2.4.4 Pipeline

Despite our best efforts, the data as they comes off the telescope are not science-ready. Considerable processing is still required, and so we have built a data processing pipeline for this. The pipeline is tasked with several important steps, including correction of receiver balancing and cross-polarization, removal of spurious, periodic signals, flagging of sporadic RFI, calibration, and mapmaking.




Figure 2.8: Simulated C-BASS sky coverage (in Galactic coordinates) assuming 12 hours observing per day, 15 days per month, and 9 months per hemisphere. Half of observing time spent  $22^\circ$  above pole elevation. Integration time is per  $27.5'$  pixel.

Work on the pipeline formed a large part of this PhD and is discussed in detail in Chapters 4, 5, and 6.

## 2.5 Northern Telescope Performance

It is one thing to describe the components of such an endeavour as this, and another entirely to actually make it work. Commissioning of the Northern telescope began in Fall of 2009 and has stretched on through the variety of instrumentation challenges that tend to plague such experiments.

Testing was completed to demonstrate the performance of the instrument using data from early 2012. At this time, the first total intensity channel, denoted Stokes  $I_1$ , was working well. However, the Stokes  $I_2$ ,  $Q$ , and  $U$  channels were not performing well due to a poorly performing LNA and phase mismatches in the polarimeter. The key performance metrics for Stokes  $I_1$  are compiled in Table 2.1.

The measured system temperature of 66 K is poor for a modern receiver at 5 GHz. The reason for this disappointing performance is a combination of the greater-than-expected gain temperatures and lower-than-expected gains of the LNAs. In the Fall of 2012, new amplifiers with lower noise temperatures and higher gains were installed. Preliminary tests have shown that the system temperature is

Table 2.1: C-BASS North Performance

Measurement	Value
Scan Rate	4 °/s
Pointing Accuracy RMS	4.6'
System Temperature	67.1 K
Bandwidth	487 MHz
Beam FWHM	44'

indeed improved. Table 2.2 shows the system temperature budgets for the system before and after the amplifier upgrade. The table lists, in order, the various noise components of the system and carries through a calculation of the running gain and running system temperature. The difference in system temperature due to the differences in LNA gain and gain temperature is striking.

Table 2.2: C-BASS Noise Budget

Component(s)	Gain (dB)	Physical Temperature (K)	Gain Temperature (K)	Running Gain (dB)	Running Temperature (K)
Winter 2012					
CMB	...	2.7	...	0.00	2.7
Atmosphere	-0.06	290	...	-0.06	6.7
Optics	-0.11	290	...	-0.17	14.3
Cold Electronics	-4.33	3.6	...	-4.50	20.7
LNA	27.00	...	12	22.50	54.5
40 K Electronics	-6.60	38	...	15.90	55.3
Warm Electronics	-0.90	290	...	15.00	57.0
Warm Amplifier	30.00	...	200	45.00	63.3
Fall 2012					
CMB	...	2.7	...	0.00	2.7
Atmosphere	-0.06	290	...	-0.06	6.7
Optics	-0.11	290	...	-0.17	14.3
Cold Electronics	-4.33	3.6	...	-4.50	20.7
LNA	36.00	...	3	31.50	29.2
40 K Electronics	-6.60	38	...	24.90	29.3
Warm Electronics	-0.90	290	...	24.00	29.5
Warm Amplifier	30.00	...	200	54.00	30.3

A demonstration of the full system was performed in the Winter of 2012, through an extremely deep integration at the North Celestial Pole. The reduction of these data demonstrated the proper operation of the instrument, while allowing for the first Galactic science from C-BASS. The analysis forms one part of the science content of this PhD and is discussed in detail in Chapter 7.

## Chapter 3

# The C-BASS Digital Backend

### 3.1 Context

Although a largely analog affair, the Northern C-BASS receiver nevertheless requires a digital backend for demodulation, filtering, and integration of the time series data. In Figure 2.6 of Section 2.4.2, the backend can be seen as the rectangle-enclosed components near the bottom of the schematic. The post-detection data enters the backend, is anti-alias filtered, is digitized in the Analog-to-Digital Converters (ADCs), and enters the Field Programmable Gate Array (FPGA) for digital processing. At the same time, the FPGA provides control signals for the noise diode and the receiver's phase switches.

An FPGA-based system is optimal for this purpose. An initial investigation into the requirements for a CPU-based backend showed that it would be a challenge to fit the necessary hardware onto the antenna, due to space and electrical power constraints. One of the strength of FPGAs, however, is that they can easily perform real-time processing of time series data through massively parallel computation and without need for operating system overheads. Thus, what would have required many 2 GHz CPU cores was more easily, compactly, and inexpensively achieved using a single, 50 MHz FPGA.

Fortunately, it was not necessary to design the backend hardware from scratch. Rather, it was possible to repurpose hardware designed at the University of Oxford for the LiCAS experiment (Reichold et al. 2006), which was well-suited to our needs. Minimal hardware modifications were needed, but significant modification was required for the firmware. This is described in detail in Section 3.4.

### 3.2 Requirements

The digital backend accepts twelve channels of data. Each channel has been anti-alias filtered below 1 MHz, which restricts the bandwidths to allow a Nyquist frequency of 2 MHz. The twelve channels

are modulated at 1 kHz, with the astronomical signal being found via pairwise subtraction of the channels as shown in Table 3.1 (wherein the modulation is noted by the  $\pm$ ). The bandwidth of the astronomical signal is comparatively small at less than 40 Hz (the actual bandwidth being set by the scan rate of the telescope), but is contained at the modulation frequency (and its harmonics).

Table 3.1: Calculation of the Polarization Channels

Polarization	Calculation
$I_1$	$\pm(\text{Channel 1} - \text{Channel 2})$
$Q_1$	$\pm(\text{Channel 3} - \text{Channel 4})$
$U_1$	$\pm(\text{Channel 5} - \text{Channel 6})$
$Q_2$	$\pm(\text{Channel 7} - \text{Channel 8})$
$-U_2$	$\pm(\text{Channel 9} - \text{Channel 10})$
$I_2$	$\pm(\text{Channel 11} - \text{Channel 12})$

The desired outputs from the backend are the six polarization channels (two each of Stokes  $I$ ,  $Q$ , and  $U$ ) and the diagnostic information needed to evaluate the fidelity of these data. There is such a thing as too much data: the backend is connected to the control computer via USB, so the control computer needs to initiate transfers of data. As the control computer is busy with many other tasks, it is only able to request 25 USB transfers from the backend per second. A single USB bulk transfer contains 512 B (where B is the abbreviation for Byte), so we are limited to an output bandwidth of 10 kB/s.

An important consideration is the actual information content of the data time series. As the bandwidth of the astronomical signal is less than 40 Hz, the Nyquist theorem tells us that it is unnecessary to have an output rate faster than 80 Hz (after demodulation). Extending this to 100 Hz is useful, as it accommodates the smooth attenuation above 40 Hz from the filtering and downsampling.

Timing is an important issue for the backend, which must act as a bridge between the digitization and integration timescales. Of these, the first is tied to the backend clock and the second the GPS clock. Bringing data across this divide in real-time in a consistent and reliable manner is a nontrivial matter, especially when, as is the case for C-BASS, the GPS clock signal is extremely noisy. Nevertheless, the digital backend must be designed to deal with these problems.

Careful filtering of the data is necessary during downsampling, as any non-white noise components (such as 60 Hz interference from power lines) will otherwise survive and appear in our data via aliasing. As our scanning strategy maps spatial structure to received time series data, such non-white noise can introduce artificial features into our final maps. At the same time, unfiltered data can be useful for diagnostic purposes, both because non-white noise can serve as a bellwether of instrument health and because filtering on the limited, backend hardware destroys other useful

information (see Section 3.4.6). As such, it is necessary for the backend to produce both filtered and unfiltered time series.

The analog receiver design minimizes noise under two assumptions. First, the cold load is assumed to be the same temperature as the sky signal (post ortho-mode transducer). Second, the paired channels are assumed to have perfect isolation from each other. A failure of the former compromises the subtraction of correlated noise, while a failure of the latter reduces the sky signal relative to the receiver noise. While the signal-to-noise ratio cannot recover if the latter fails, the failure of the former can be partially compensated for post-detection. By using the noise diode to measure and correct paired-channel leakage, then scaling the cold load channel to match the sky channel amplitude (called the  $\alpha$  and r-factor corrections, respectively; see Sections 4.2.2 and 4.2.4 for detailed discussions), the correlated noise can be reduced at the cost of increased white noise. For unfiltered data, these corrections can be applied offline in the data pipeline, but the same is not true for filtered data. Thus, facilities for these corrections are needed in the backend itself.

The real-time, digital control of the noise diode and phase switch modulation within the receiver requires a real-time, digital controller. As the backend controls the sample timing and demodulation, it is natural for it to be in control of these receiver functions. Indeed, controlling the modulation from the backend allows for easy demodulation within the backend.

The functioning of the receiver can be monitored through diagnostic signals, a number of which are available to the backend. These include the status of the GPS time signal and the state of the internal memory. Packaging these signals with the data time series facilitates monitoring the backend, thus forming an essential function in and of itself.

Sometimes more detailed troubleshooting is necessary. As is the case for any complex, electronic device, this can even require probing of internal, raw voltages. The digital backend sees such information in the form of the raw ADC outputs. Preserving and making these measurements available at all times is impossible within the bandwidth constraints, but the hardware does allow for short bursts of the ADC outputs to be saved and transferred to the computer upon request.

These various constraints form something of a wish list of backend features. Meeting these needs has been a challenge, mainly due to limits in the backend hardware itself. This challenge has been met, however, as is described below.

### 3.3 Hardware

The backend hardware consists of four cards housed in a single assembly, shown in Figure 3.1. Component A is an attenuator board, which allows for incoming signals to be reduced in amplitude to avoid ADC saturation (if necessary). The signal amplitudes coming from the C-BASS analog receiver are not so strong as to require attenuation, and so this board is bypassed for C-BASS




Figure 3.1: The C-BASS digital backend assembly. Labeled components are (A) the attenuator board, (B) the digital processing (ADC/FPGA) card, (C) the analog filter board, and (D) the digital I/O board.

operation. Component C is the analog filter board, providing the parallel, 1 MHz, anti-alias filters and a voltage gain of 6.5 to fully use the dynamic range of the ADCs. The analog signals are conducted to the digital processing card (B) via the backplane, whereupon they are digitized by the ADCs and sent to the FPGA for processing. The filtered and integrated signals are retrieved directly from the digital card via Universal Serial Bus (USB). Component D provides direct access to several FPGA input and output (I/O) pins, allowing for interfacing to control and timing signals.

The digital processing card is the workhorse of the backend and so the focus of the remaining discussion. Pictured in Figure 3.2, this card has five components of interest: (A) the 16 parallel ADCs, (B) the main FPGA, (C) the secondary FPGA, (D) 128 MB of Synchronous Dynamic Random-Access Memory (SDRAM), and (E) the USB microcontroller. The card must be clocked at 50 MHz, which can either be provided externally or taken from an on-card oscillator. C-BASS operates in the latter mode. The C-BASS receiver has just 12 channels, so four of the ADCs are unused. The filtering and integration occurs on the main FPGA. The secondary FPGA acts solely as an interface to the SDRAM, itself only used for the occasional collection of raw ADC outputs for diagnostic purposes. A list of the important chips and part numbers is found in Table 3.2.

The ADCs provide 14 b (where b is the abbreviation for bit), differential measurements of signals between  $\pm 1.25$  V with a maximum sampling rate of 2.78 MSps (where MSps abbreviates Million




Figure 3.2: The C-BASS digital processing (ADC/FPGA) card. Labeled components are (A) the 16 ADCs, (B) the main FPGA, (C) the secondary FPGA, (D) the SDRAM chips, and (E) the USB microcontroller.



Table 3.2: Notable Components of the C-BASS Digital Backend

Component	Manufacturer	Part Number	Comment
Serial 14-Bit ADC	Linear Technology	LTC1403A-1	2 MSps sampling
Spartan-3 FPGA	Xilinx	XC3S1000-4	Main and Secondary
Synchronous DRAM	Micron Technology	MT48LC32M16A2	Two installed
USB Microcontroller	Cypress Semiconductor	CY7C68013A	“Default” Configuration

Samples per second) when clocked at 50 MHz. They rely upon external control signals to initiate a single conversion, allowing fine control of the exact sampling rate. Conversion and serial transfer of a single reading requires 18 clock cycles (setting the maximum sampling rate). For C-BASS, samples are requested every 25 clock cycles, giving a sampling rate of 2 MSps.

The main FPGA interfaces to and controls the ADCs, can communicate with the secondary FPGA when desired, and sends the processed data products to the USB microcontroller. Being of the Xilinx Spartan series, it has relatively few resources for an FPGA: just 24  $18 \times 18$  multipliers, 24 18 kb blocks of RAM (which share inputs and outputs with their adjacent multipliers), and 17 280 equivalent logic cells. Despite this, the remarkable parallelization possible on the FPGA allows these modest resources to be translated into a powerful digital signal processing (DSP) chain. A description of necessary firmware is provided in Section 3.4.

The SDRAM consists of two identical, 64 MB chips, and requires buffering during reading and writing, provided by the secondary FPGA. This FPGA is identical to the main FPGA in terms of resources, but with firmware tuned to this more modest role. The 128 MB of memory has no use in regular DSP, but can be tasked with buffering raw ADC outputs when needed.

Transfer of processed data to the control computer is achieved via USB. 25 512 B packets are sent per second via USB 2.0 Bulk transfers. These packets are assembled on the main FPGA and sent to the USB microcontroller in 16 b segments. Upon the USB microcontroller, they are reassembled and made available upon request to the control computer. The number of transfers per second is limited not by the USB microcontroller or the USB standard, but by the control computer itself, which is multitasking several parallel control processes and cannot be relied upon to request many more transfers during any given second.

A number of I/O signals interface with the main FPGA. First amongst these is the 1 pulse-per-second (1PPS) signal provided by the antenna’s Global Positioning System (GPS) receiver, used for synchronizing the integrated data product with the antenna’s pointing and diagnostic data. The FPGA also drives the six  $180^\circ$  phase switches in the analog receiver, as well as the on/off switch for the receiver’s calibration noise diode. Each of these I/O signals are interfaced through the digital I/O board (component D in Figure 3.1).

### 3.4 Firmware

The firmware of the C-BASS digital backend’s main FPGA, representing a large portion of this PhD work, is now described in detail. The FPGA was programmed using VHDL<sup>1</sup> in the Xilinx Integrated Software Environment (ISE Version 10.1). The Xilinx ISE offers easy access to a variety of customizable DSP components, in-depth debugging tools, FPGA hardware constraints, and the FPGA compiler. VHDL is a concurrent programming language, meaning that, by default, every line of code is evaluated simultaneously rather than sequentially. This is a natural choice for highly-parallel hardware, such as an FPGA, but is nonetheless quite different from the more popular sequential paradigm of C, Fortran, Python, and others.

The FPGA firmware has gone through a number of revisions during its design, testing, and commissioning. The current version number is known to the firmware, and is included with outgoing data packets. The current version number is 7.

Of order  $10^4$  lines of code comprise the digital backend firmware. This code is split amongst a number of “components”, each of which is contained in an individual text file. The higher-level components can call the lower level ones, and so is the entire system built up. Several of the filter and memory components were not coded by hand, but were instead provided using graphical “wizards” provided by Xilinx. These wizards allowed for customization of licensed Xilinx IP and provided useful, high level estimates of FPGA resource usage. A simplified schematic drawing of the FPGA firmware, including its most important components, is shown in Figure 3.3.

The chief components are briefly described:

**dcm\_cbass** Provides the clocking signals for the entire FPGA, as well as for the ADCs and the USB microcontroller. Discussed further in Section 3.4.1.

**pps\_control** Contains algorithms for detecting the GPS 1PPS rising edge in the presence of noise. Provides GPS-synchronized timing pulses to the data processing components. Discussed in Section 3.4.2.

**register\_control** Manages the variety of firmware registers used by the FPGA, listed in Section 3.4.3.

**noise\_control** Controls the noise diode on/off signal. See Section 3.4.4.

**acquisition\_control\_filt** Provides the timing pulses for ADC triggering, data trimming, and phase switching.

**rom\_walsh** Accepts timing pulses from `acquisition_control_filt` to generate the phase switching and demodulation signals, which are described in Section 3.4.4.

---

<sup>1</sup>VHDL is a two-level acronym: VHSIC (Very High Speed Integrated Circuit) Hardware Description Language.

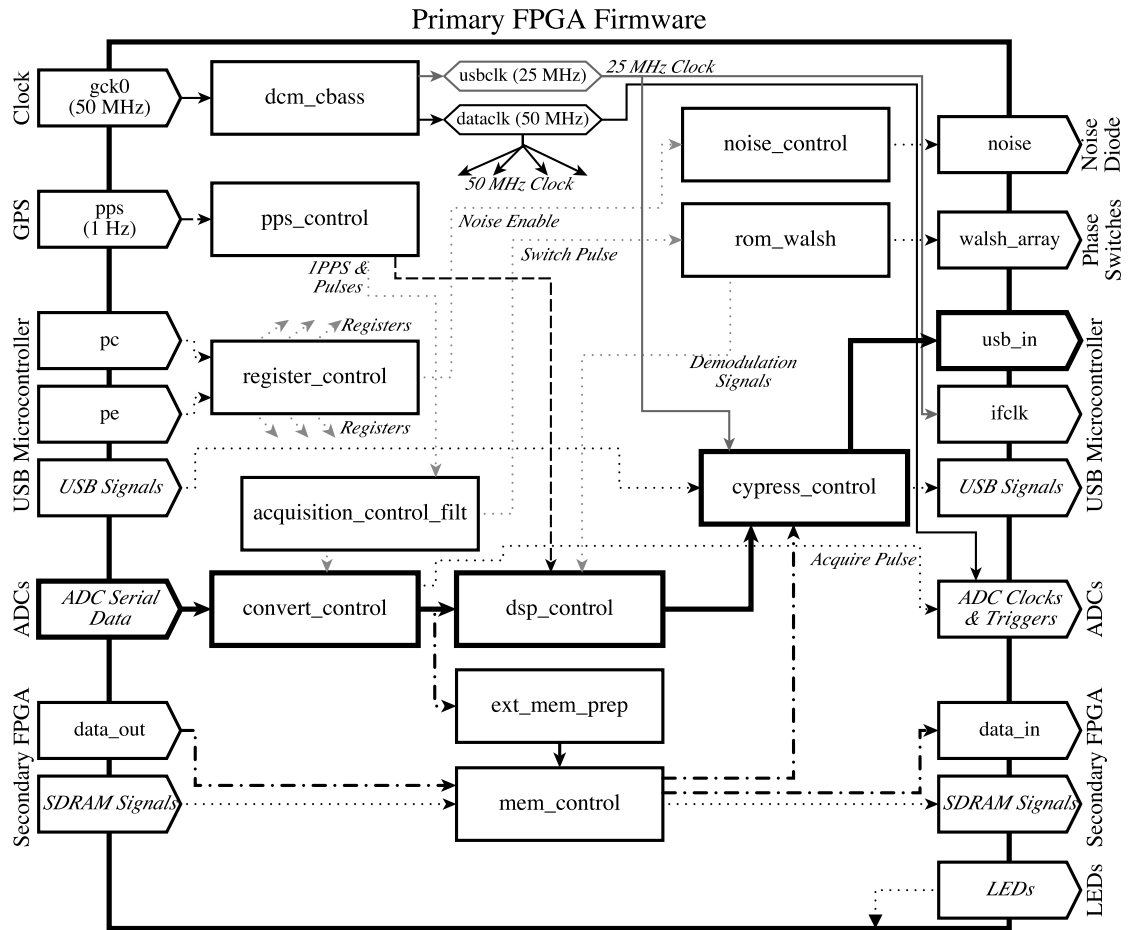


Figure 3.3: Schematic illustration of the primary FPGA's firmware. The components are represented as rectangles. Pentagons represent the FPGA's input and output signals. Internal signals, and groups thereof, are represented by arrows. The bold lines and components mark the path of continuous data. The dot-dashed line is for burst data.

**convert\_control** Sends acquisition pulses to the ADCs and then receives the serial data therefrom.

Provides the parallel, time series data to the various DSP options (thoroughly discussed in Section 3.4.6).

**dsp\_control** Filters and integrates the time series data. Described in Section 3.4.6.

**cypress\_control** Sends the processed data to the USB microcontroller.

**ext\_mem\_prep** When a burst of raw, ADC data are requested, prepares the data for temporary storage in the backend's SDRAM. Discussed further in Section 3.4.5.

**mem\_control** Sends data to the SDRAM when a burst of data is requested. After the data have been collected, forwards the data from the SDRAM to **cypress\_control** to be sent out via

USB.

One structure often used in this firmware is First in, First out (FIFO) memory. In a FIFO, data are written sequentially, and then read out in the same order. This has proven very useful for constructing data filters and queues.

The nature of FPGAs makes them inefficient at floating-point arithmetic, but adept at fixed-point arithmetic. Floating-point numbers consist of a *mantissa* and an *exponent*, and floating-point calculations are allowed to vary both. For a fixed-point number, the *exponent* is implicit and is not stored, and neither is it allowed to vary during calculations. Care is required to avoid under- or over-flow of significant bits. Beyond this, the value of the *exponent* is unimportant, as the data are not calibrated to a physical scale until processed through the offline pipeline.

Fixed-point arithmetic can represent negative numbers in a variety of ways. The convention used by the ADCs and on the FPGA is two's complement. For an  $n$ -bit number, the allowable values range between  $-2^{n-1}$  and  $2^{n-1} - 1$ , inclusive. Positive numbers are represented as expected, with the most significant bit set to zero. Negative numbers are represented by subtracting the absolute value from  $2^n$  and trimming the most significant bit. In this scheme, the most significant bit of a negative number is 1. Two's complement is convenient because the sum of positive and negative numbers is handled naturally.

### 3.4.1 Timing

The main FPGA is responsible for coordinating several different timescales: the backend clock, the USB microcontroller clock, the ADC sampling rate, the phase switch rate, the integration timescale, the USB transfer rate, and the 1PPS from the GPS. Although there are several constraints on these timescales, only the backend clock, the USB clock, and 1PPS timescales were absolutely fixed. The timescales were thus chosen to satisfy a number of constraints:

- the ADC sampling rate cannot exceed 2.78 MSps,
- the USB transfer rate cannot exceed 25/s,
- there must be an integral number of phase switch periods per integration,
- the phase switch frequency must be much less than the analog video bandwidth, and
- the higher frequencies need to be multiples of the lower.

The chosen frequencies and periods for the various timescales are summarized in Table 3.3.

Although the FPGA can phase-lock its internal 25 MHz 50 MHz clocks to the on-board 50 MHz clock (using its Digital Clock Manager, or DCM), there is no way to phase lock these clocks to the 1PPS. Any variation in the backend clock thus goes uncorrected, which can be significant. Indeed,

Table 3.3: Timescales of the C-BASS Digital Backend

Timescale	Family	Frequency	Period
Backend Clock	Clock	50 MHz	20 ns
USB Clock	Clock	25 MHz	40 ns
ADC Sampling	Clock	2 MHz	500 ns
Modulation	Clock	1 kHz	1 ms
Integration	1PPS	100 Hz	10 ms
USB Transfer	1PPS	25 Hz	40 ms
GPS 1PPS	1PPS	1 Hz	1 s

careful testing showed that the backend clock, while precise to 1 ppm, is inaccurate at 5 ppm. It is useful to think of this as two separate timing families, one driven by the backend clock and one driven by the real-time GPS signal. In this picture, it is necessary to carefully delineate which signals belong to which family and where the interface sits. This demarcation is included in Table 3.3; the interface found a natural home in the integration stage, and is described in Section 3.4.5.

Within the clock family, there are two clock domains: the 50 MHz main clock and the 25 MHz USB microcontroller clock. The latter is a simple division of the former, but it is a nontrivial distinction nonetheless. Processed data transferred from the main FPGA to the USB microcontroller must be provided on the slower clock scale, which requires that the FPGA itself be divided into two clock domains. The two are bridged through one of the FPGA’s RAM blocks, with the necessary configuration provided via a Xilinx-provided, drop-in component. This is discussed further in Section 3.4.6.

### 3.4.2 1PPS Handling

The 1PPS signal from the GPS is crucial because it allows the received sky signal to be matched to the antenna pointing and diagnostic signals. As the mapmaking relies tying together the time series sky and pointing data, the final map can potentially be limited by the 1PPS handling. It must, therefore, be handled with extreme care. In particular, the rising edge of the 1PPS square wave, marking the precise beginning of each second of real time, must be determined precisely and accurately.

Unfortunately, the C-BASS antenna suffers from strong 60 Hz pickup on the 1PPS line. 60 Hz, being due to the power transmission lines, is notoriously difficult to eliminate. Every effort was expended to eliminate this signal from the 1PPS line, yet it remained strong enough to confound precise measurements of the 1PPS rising edge.

A solution, however, was found and implemented. This solution required the joint assumptions that the 1PPS noise consisted of low amplitude 60 Hz contamination and white noise, that this noise was able to cause the 1PPS line to appear high when it was indeed low but not vice versa, and

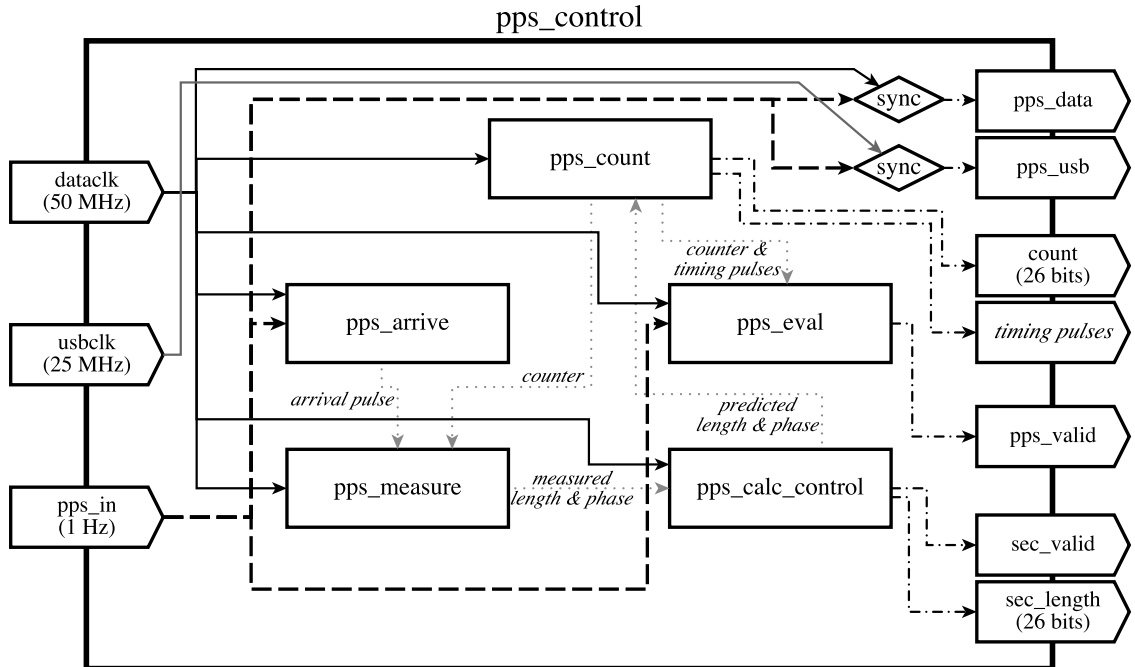


Figure 3.4: Schematic illustration of the 1PPS handling component. Internal components are marked by rectangles. Pentagons represent input and output signals. Internal signals are represented by arrows.

that the digital backend's 50 MHz clock was fairly accurate and even more precise. A multi-step algorithm was developed which then robustly detected the 1PPS rising edges and took a running mean of the number of clock cycles per second through a damping filter. Any discrepancies between internal timekeeping and the GPS signal were iteratively resolved, providing a stable time reference, accurate to  $2\ \mu\text{s}$ . As this is much, much less than the receiver integration time, the potential for systematic impact on astronomical data is exceedingly small.

This algorithm is contained within the FPGA firmware in a component called `pps_control`. A schematic illustration of `pps_control` and its subcomponents can be found in Figure 3.4. `pps_control` has five subcomponents. These are listed and briefly described:

**pps\_arrive** Awaits the arrival of the next 1PPS rising edge. The rising edge is only accepted if the 1PPS line has been low for at least 500 clock cycles prior to the rising edge and the 1PPS line stays high for at least 500 clock cycles following the rising edge. If these conditions are met, then an *arrival pulse* is sent to `pps_measure`.

**pps\_measure** Upon receipt of the *arrival pulse*, calculates the length of the previous second (in clock cycles) and any phase difference between the 1PPS and internal counter, as reported by `pps_measure`, is taken into account. The measured second *length* and *phase* are sent to

`pps_calc_control` for use in predicting the *length* of future seconds.

**pps\_calc\_control** Calculates the running mean of the last 32 second lengths to predict the length of the next second. The predicted second *length* and corrective *phase* shift (which is the measured *phase* divided by 32 to avoid sudden discontinuities) is sent to `pps_count` for use during the next second.

**pps\_count** Maintains an internal counter representing the backend firmware’s best guess of the current time between 1PPS arrivals. The counter counts up to the *predicted second length*, modified by the predicted second *phase*, as calculated by `pps_calc_control`. This counter is used by `pps_measure` to measure the current second’s length, as well as to drive several *timing pulses* for use in the firmware’s DSP chain.

**pps\_eval** Ensures the health of the *counter & timing pulses* coming from `pps_count` in comparison to the 1PPS signal itself. If all appears to be well, the `pps_valid` flag is raised, required for the proper operation of the rest of the backend firmware.

These five components work together in a negative feedback loop which is robust to noise on the 1PPS line and quickly damps any errors resulting from rare errors when detecting the 1PPS rising edge. Indeed, the system is designed to work using predicted second lengths for eight full seconds in the event that no 1PPS rising edge is detected.

The `pps_valid` flag is high when all is working well, and is used by the rest of the firmware to ensure proper timing. The *predicted second lengths* are bundled with the astronomical data product and are saved in the C-BASS data archive. An illustrative sample of these are shown in Figure 3.5. It is clear from these time series that the timing varies over the day, likely due to temperature variations of the backend hardware. Errors in the detection of the 1PPS rising edge occur on second timescales, but result in timing errors of less than  $2\ \mu\text{s}$  and are quickly damped.

### 3.4.3 Parameters, Registers, and Diagnostics

A wide variety of backend options and parameters are available to the user. These comprise choices of observing mode (continuous or burst), data modulation, and nonlinearity and balance correction. The parameters are stored in the backend in a set of registers, managed by the `register_control` component. Writing to the registers is achieved via the parallel inputs `pc` (the register address) and `pe` (the updated register) from the USB microcontroller.

The backend registers are listed in Table 3.4. Most of these registers are described elsewhere in this chapter (with relevant section shown in the table), with the exception of the `reset` bit of the `acquire_reg` register. This signal resets all of the filters and memory on the FPGA and is intended for use when switching between continuous and burst mode.

Table 3.4: Primary FPGA Firmware Registers

Register Address	Register Name	Number of Bits	Description	Section
0	<code>nonlin_1A_reg</code>	18	Nonlinearity coefficients for Channel 1	(3.4.6)
3	<code>nonlin_1B_reg</code>	18		
6	<code>nonlin_2A_reg</code>	18	Nonlinearity coefficients for Channel 2	(3.4.6)
9	<code>nonlin_2B_reg</code>	18		
12	<code>nonlin_3A_reg</code>	18	Nonlinearity coefficients for Channel 3	(3.4.6)
15	<code>nonlin_3B_reg</code>	18		
18	<code>nonlin_4A_reg</code>	18	Nonlinearity coefficients for Channel 4	(3.4.6)
21	<code>nonlin_4B_reg</code>	18		
24	<code>nonlin_5A_reg</code>	18	Nonlinearity coefficients for Channel 5	(3.4.6)
27	<code>nonlin_5B_reg</code>	18		
30	<code>nonlin_6A_reg</code>	18	Nonlinearity coefficients for Channel 6	(3.4.6)
33	<code>nonlin_6B_reg</code>	18		
36	<code>nonlin_7A_reg</code>	18	Nonlinearity coefficients for Channel 7	(3.4.6)
39	<code>nonlin_7B_reg</code>	18		
42	<code>nonlin_8A_reg</code>	18	Nonlinearity coefficients for Channel 8	(3.4.6)
45	<code>nonlin_8B_reg</code>	18		
48	<code>nonlin_9A_reg</code>	18	Nonlinearity coefficients for Channel 9	(3.4.6)
51	<code>nonlin_9B_reg</code>	18		
54	<code>nonlin_10A_reg</code>	18	Nonlinearity coefficients for Channel 10	(3.4.6)
57	<code>nonlin_10B_reg</code>	18		
60	<code>nonlin_11A_reg</code>	18	Nonlinearity coefficients for Channel 11	(3.4.6)
63	<code>nonlin_11B_reg</code>	18		
66	<code>nonlin_12A_reg</code>	18	Nonlinearity coefficients for Channel 12	(3.4.6)
69	<code>nonlin_12B_reg</code>	18		
72	<code>alpha_I1_A_reg</code>	18	Balance coefficients for Polarization $I_1$	(3.4.6)
75	<code>alpha_I1_B_reg</code>	18		
78	<code>alpha_21_A_reg</code>	18	Balance coefficients for Polarization $I_2$	(3.4.6)
81	<code>alpha_21_B_reg</code>	18		
112	<code>burst_reg</code>	21	Number of ADC samples per burst	(3.4.5)
116	<code>switch_reg</code>	8	Modulation frequency	(3.4.6)
120	<code>trim_reg</code>	8	Trim length following phase switch	(3.4.6)
124	<code>acquire_reg</code>	8	0 <code>adc_read</code>	(3.4.4)
			1 <code>reset</code>	(3.4.3)
			6 <code>cont_mode</code>	(3.4.5)
			7 <code>simulate_en</code>	(3.4.5)
125	<code>trigger_reg</code>	8	0 <code>trigger</code>	(3.4.5)
			1 <code>nonlin_en</code>	(3.4.6)
			2 <code>alpha_r_en</code>	(3.4.6)
			4 <code>noise_on</code>	(3.4.4)
			5 <code>switch_en</code>	(3.4.4)
			6 <code>switch_alt</code>	(3.4.4)
			7 <code>switch_full</code>	(3.4.4)






Figure 3.5: Timing errors in detecting the 1PPS rising edge. Lower panel are data from January 13-19, 2012, smoothed over 64 s. Upper panel is zoomed in, without smoothing.

The contents of these registers are included with the data products produced by the backend, allowing for real-time monitoring of the backend settings and historical settings in the C-BASS archive. In continuous mode, the parameter bits of output Packet #1 (see Table 3.8) correspond to several bits of the `trigger_reg` register, as spelled out in Table 3.5. The continuous data also include the modulation rate register `switch_reg` in Packet #3, the demodulation trimming length `trim_reg` in Packet #4, and the nonlinearity and balance correction coefficients in Packet #8 (further described in Table 3.9). A number of diagnostic signals exist outside of the registers, and are also made available to the user. Packet #2 includes several, also listed in Table 3.5, such as the “empty” flag for the output FIFO, the health of the FPGA’s DCM, and the current state of the 1PPS input. Further included are the output FIFO status (also in Packet #2), the backend version number (in Packet #6) and the health of the unfiltered delay memory (Packet #7).

For burst mode, the parameters (see Table 3.10) are the contents of the `trigger_reg` register, as spelled out in Table 3.6. Some diagnostics are also provided. These are listed again in Table 3.6 and include the FPGA’s DCM health, the 1PPS state, the data-synchronized modulation state, and the data trimming flag.

Finally, the backend is outfitted with four LEDs which are visible from the front of the backend assembly. These LEDs were useful during the programming and debugging of the backend firmware,

Table 3.5: Continuous Data Parameters &amp; Diagnostics

Location	Description	Bit	Signal	Comment
Packet #1 Bits 5 $\rightarrow$ 0	Parameters	5	<code>switch_alt</code>	Alternate switching register
		4	<code>switch_en</code>	Switch enable register
		3	<code>noise_on</code>	Noise diode register
		2	<code>switch_full</code>	Full switching register
		1	<code>alpha_r_en</code>	Balance correction register
		0	<code>nonlin_en</code>	Nonlinearity correction register
Packet #2 Bits 3 $\rightarrow$ 0	Diagnostics	3	<code>empty_sync</code>	Output FIFO empty flag
		2	0	Reserved
		1	<code>dcm_locked</code>	FPGA DCM health
		0	<code>pps</code>	Current 1PPS state

Table 3.6: Burst Data Parameters &amp; Diagnostics

Location	Description	Bit	Signal	Comment
Packets #1 & #2 Bits 17 $\rightarrow$ 12	Parameters	17	<code>switch_full</code>	Full switching register
		16	<code>switch_alt</code>	Alternate switching register
		15	<code>switch_en</code>	Switch enable register
		14	<code>noise_on</code>	Noise diode register
		13	<code>alpha_r_en</code>	Balance correction register
		12	<code>nonlin_en</code>	Nonlinearity correction register
Packets #1 & #2 Bits 11 $\rightarrow$ 8	Diagnostics	11	1	Reserved
		10	1	Reserved
		9	<code>dcm_locked</code>	FPGA DCM health
		8	<code>pps</code>	Current 1PPS state

but have in the current version been turned off. Activity in these LEDs was seen to have a measurable effect on the data channels by introducing small offsets into the electrical ground level, which was unacceptable.

### 3.4.4 Control Signals

The digital backend provides a natural platform from which to control the digital signals used by the analog receiver. These are the ADC triggers, the noise diode control and the six phase switch modulation signals. These are now described.

Controlling the ADCs consists of providing them with an FPGA-synchronized clock, sending an acquisition pulse when data are desired, and waiting for the serialized data to arrive. These jobs are handled by the `acquisition_control_filt` component, which is enabled by the `adc_read` register.

The noise diode injects a set level of thermal noise into the analog sky channels prior to combination with the temperature-controlled loads. This is an essential component of the C-BASS calibration strategy, being used to measure the C-BASS gain and balance on minute timescales. The backend does little to regulate the noise diode signal: upon raising of the `noise_on` register, the signal is forwarded to the noise diode immediately. The chief advantage in using the backend for this is that the state of the noise diode is packaged with the backend data product. The actual state of the noise diode is thus forever recorded in the C-BASS data archive.

180° modulation of the analog data channels assures symmetry in the analog components. There are two phase switches each for the  $I_1$ ,  $I_2$ , and  $Q + U$  analog chains, for a total of six. Each pair works in concert, providing four possible phase switch states: 00, 01, 10, and 11. These are partially degenerate: in the absence of complex gain errors, 00 and 11 give the same output, as do 01 and 10. Each switch is modulated with a square wave, with each phase switch pair being out of phase by a quarter period. In this way, the output modulation varies between the four states at the individual phase switch frequency and between the degenerate states at twice the rate. The default individual phase switch frequency is 500 Hz, implying a modulation frequency of 1 kHz. The  $Q$  and  $U$  modulation phase shifted by 180° from the  $I_1$  and  $I_2$  modulation phase.

The modulation signals originate in the backend, motivated by the convenience of generating the modulation signal in the same hardware that performs the demodulation. In particular, if the modulated signals were generated externally, then there would be questions of modulator communication with the backend and jitter between the backend and modulator clocks. Indeed, the only issue that does remain is the latency between outgoing, modulation signal and the detected, modulated data. Experimentation showed this delay to be a constant  $0.98 \mu\text{s}$ , which has been hardcoded into the backend.

The signals are generated in the backend component `rom_walsh`, whose name refers to the original strategy of using Walsh functions for modulation, rather than square waves. Although the default

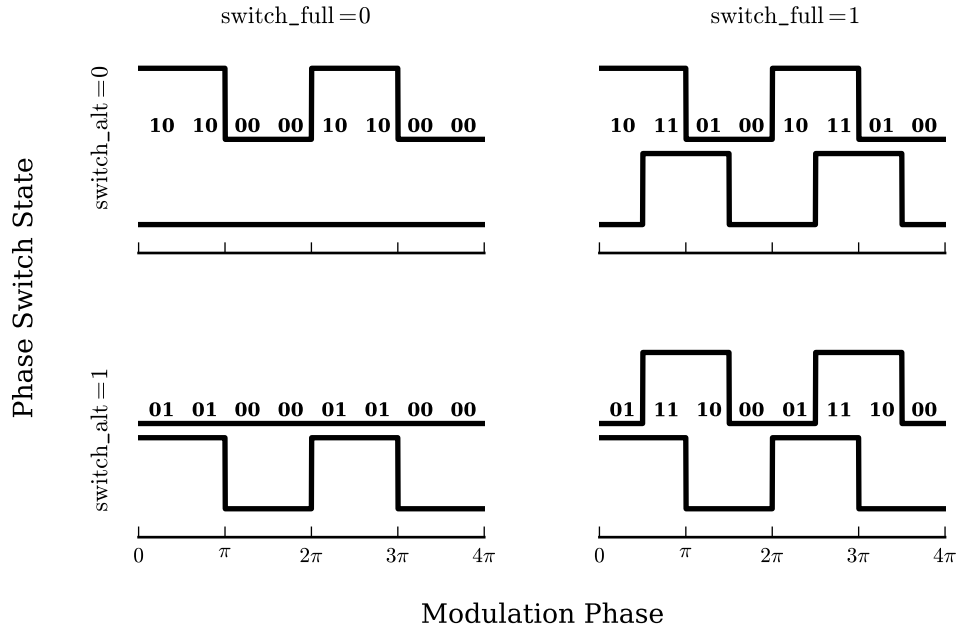


Figure 3.6: Illustration of the various phase switching options.

frequency is 1 kHz, several other, higher frequencies are allowed (based upon assumptions in the DSP chain). Valid choices can be selected by setting the `switch_reg` register, as described in Section 3.4.3, and are enforced by the component `timing_control_hc`. The valid choices are discussed further in Section 3.4.6 and are listed in Table 3.13. The backend allows for the phase switching to be turned off (via the `switch_en` register), for only one of each pair to be switched (via `switch_full`), and for the phase switching to be reversed (via `switch_alt`). The modulation waveforms, and their behaviour with respect to the various options, are illustrated in Figure 3.6.

### 3.4.5 Data Acquisition Modes

Survey observations call for continuous data acquisition during scanning with maximum suppression of systematic noise. There are times, however, when diagnostic tests require minimally processed data. It was therefore necessary to design the C-BASS backend with the flexibility to meet these needs. Two observation modes are available: continuous mode and burst mode. In the former, data are acquired, processed, and made available without interruption. In the latter, raw data are collected and stored for a short time, being transferred to the control computer only after collection has ceased. The choice of mode is set by the `cont_mode` register.

In regular operations, data arrives on the FPGA from the ADCs. During commissioning and testing of the backend firmware, well-defined signals were of great use. These signals, referred to as

“simulated data”, were generated within the backend firmware in the `simulate_control` component and enable via the `simulate_en` register. They were contrived to test various aspects of the DSP chain, and were not meant to resemble actual astronomical signal. They are described in Table 3.7. In this table, the “A” and “B” suffixes on the channel numbers represent modulation state.

Table 3.7: Simulated Data &amp; Output Predictions

Channel	Simulated Input	Unfiltered Output	Filtered Output
1A	6144	3 840 000	
1B	7168	4 480 000	81 920 000
2A	5120	3 200 000	
2B	4096	2 560 000	
3A	3072	1 920 000	
3B	3584	2 240 000	40 960 000
4A	2560	1 600 000	
4B	2048	1 280 000	
5A	8191	5 119 375	
5B	0	0	-655 280 000
6A	0	0	
6B	8191	5 119 375	
7A	8191	5 119 375	
7B	0	0	-655 280 000
8A	0	0	
8B	8191	5 119 375	
9A	768	480 000	
9B	896	560 000	10 240 000
10A	640	400 000	
10B	512	320 000	
11A	68	42 500	
11B	70	43 750	160 000
12A	66	41 250	
12B	64	40 000	

The various data sources and processing result in an array of observing choices. These are shown schematically in Figure 3.7. The continuous and burst modes are described in further detail below.

The backend has operated in continuous mode for the vast majority of its lifetime. Data are digitized, demodulated, filtered, downsampled, and prepared for USB transfer. The resulting, 100 Hz time series are bundled into 25 USB bulk transfers per second, easily requested and processed by the control computer. The filtered data are free of all systematic, periodic signals with the exception of 1.2 Hz (plus harmonics) from the cryostat’s cooling cycle. 60 Hz contamination is undetectable (see Section 3.4.6). The filtering comes at a cost, however: the paired channels must be differenced, destroying information on the analog receiver balance. To preserve this information, the demodulated data are simultaneously integrated without filtering, providing 24 additional time series (two phase

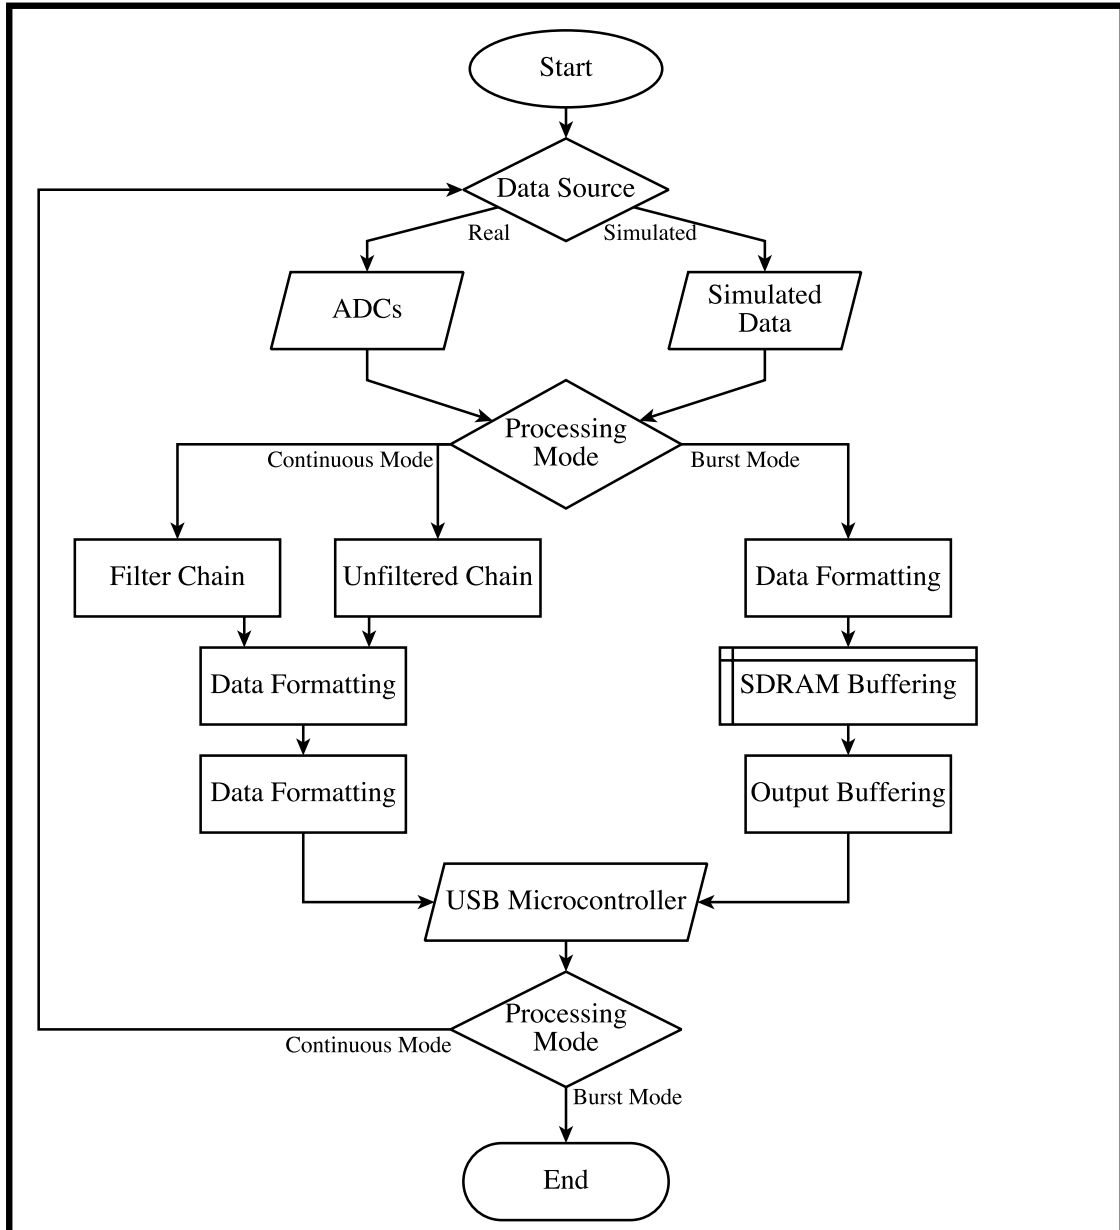


Figure 3.7: Schematic flowchart of the data acquisition modes. Shown are the choices between real and simulated data, and between continuous and burst mode.

states for 12 digitized channels). These data are also produced continuously, and transferred to the control computer and archived alongside the filtered series.

USB bulk transfers have a maximum size of 512 B, which requires careful packing in order to fit the data from four 10 ms periods. Each 10 ms period is allotted 128 B, which are themselves divided into eight 128 b packets. The first five of these packets are used to store the unfiltered data, the next two the filtered data, and the last the nonlinearity and balance corrections applied prior to filtering (as described in Section 3.4.6). In truth, 128 b are only enough for one quarter of these coefficients, so they are actually divided amongst four periods encompassed by the single bulk transfer. Amongst these are squeezed various control signals and diagnostics, which allow for real-time monitoring of the backend system. Table 3.8 shows the detailed  $8 \times 128$  b format, while Table 3.9 describes the nonlinearity/balance correction packaging. As 25 such transfers occur per second, the total data rate during continuous operation is 12.5 kB/s.

When raw ADC measurements are needed, the burst mode is used. The twelve 2 MHz time series are immediately packaged into two 128 b packets (described in Table 3.10) and shunted continuously through the secondary FPGA and into the 128 MB SDRAM. As in the continuous mode, a variety of control and diagnostic signals is included. The SDRAM has capacity to hold over 2 s of data in this mode, however the backend has been programmed to limit this to just 0.5 s due to the practical limitations in control computer bandwidth. The burst length is determined by the `burst_reg` register, and the collection of a burst is initiated by briefly raising the `trigger` register. This mode has found use in debugging and commissioning of the backend, but has not found common use in general observing. As such, it has remained unpolished.

### 3.4.6 DSP Chain

Robust processing of the time series data is the main purpose of the digital backend during continuous observing. The processing is split into two parallel chains, one for filtered and one for unfiltered results. The filtered chain is designed to give maximal rejection of systematic, periodic noise while preserving the astronomical signal. The unfiltered data allows systematic noise to pass, but preserves receiver balance information.

As described in Section 3.3, 14 b data enters the FPGA via serial transfers from each ADC. The first step in processing the data is to convert the serial data to parallel, which is done in the component `convert_control`. The data then enter the `dsp_control` component through the `data_mux`, which can be switched to allow simulated data instead (see Section 3.4.5). The unfiltered chain splits the data into 24 channels, representing the two phase states of each of the 12 analog inputs, while the filtered chain pairwise-subtracts and demodulates the data down to 6 channels. The unfiltered chain then integrates the data directly, while the filtered chain filters and downsamples the data four times before arriving at the final data rate. Finally, the outputs of the two chains are

Table 3.8: Continuous Data Format

Bit Range	Description	Bit Range	Description
Packet #1		Packet #2	
127 → 102	Timestamp	127 → 104	Channel 5A
101 → 78	Channel 2B	103 → 80	Channel 4B
77 → 54	Channel 2A	79 → 56	Channel 4A
53 → 30	Channel 1B	55 → 32	Channel 3B
29 → 6	Channel 1A	31 → 8	Channel 3A
5 → 0	Parameters	7 → 4	Output FIFO status
		3 → 0	Diagnostics
Packet #3		Packet #4	
127 → 104	Channel 7B	127 → 104	Channel 10A
103 → 80	Channel 7A	103 → 80	Channel 9B
79 → 56	Channel 6B	79 → 56	Channel 9A
55 → 32	Channel 6A	55 → 32	Channel 8B
31 → 8	Channel 5B	31 → 8	Channel 8A
7 → 0	Modulation rate	7 → 0	Trim Length
Packet #5		Packet #6	
127 → 104	Channel 12B	127 → 102	Timestamp
103 → 80	Channel 12A	101 → 72	Filtered U1
79 → 56	Channel 11B	71 → 42	Filtered Q1
55 → 32	Channel 11A	41 → 12	Filtered I1
31 → 8	Channel 10B	11 → 0	Version Number
7 → 0	Reserved		
Packet #7		Packet #8	
127 → 102	Length of Second	127 → 110	Coefficient 7
101 → 72	Filtered I2	109 → 92	Coefficient 6
72 → 42	Filtered U2	91 → 74	Coefficient 5
42 → 12	Filtered -Q2	73 → 56	Coefficient 4
11 → 9	Reserved	55 → 38	Coefficient 3
8 → 0	Unfiltered FIFO status	37 → 20	Coefficient 2
		19 → 2	Coefficient 1
		1 → 0	Period ID

Table 3.9: Nonlinearity/Balance Coefficients (Continuous Packet #8)

Bit Range	Description	Period #1	Period #2	Period #3	Period #4
127 → 110	Coefficient 7	nonlin_4A_reg	nonlin_7B_reg	nonlin_11A_reg	alpha_I2_B_reg
109 → 92	Coefficient 6	nonlin_3B_reg	nonlin_7A_reg	nonlin_10B_reg	alpha_I2_A_reg
91 → 74	Coefficient 5	nonlin_3A_reg	nonlin_6B_reg	nonlin_10A_reg	alpha_I1_B_reg
73 → 56	Coefficient 4	nonlin_2B_reg	nonlin_6A_reg	nonlin_9B_reg	alpha_I1_A_reg
55 → 38	Coefficient 3	nonlin_2A_reg	nonlin_5B_reg	nonlin_9A_reg	nonlin_12B_reg
37 → 20	Coefficient 2	nonlin_1B_reg	nonlin_5A_reg	nonlin_8B_reg	nonlin_12A_reg
19 → 2	Coefficient 1	nonlin_1A_reg	nonlin_4B_reg	nonlin_8A_reg	nonlin_11B_reg
1 → 0	Period ID	00	01	10	11



Table 3.10: Burst Data Format

Bit Range	Description	Bit Range	Description
Packet #1		Packet #2	
127 → 102	Timestamp	127 → 102	Timestamp
101 → 88	Channel 1	101 → 88	Channel 7
87 → 74	Channel 2	87 → 74	Channel 8
73 → 60	Channel 3	73 → 60	Channel 9
59 → 46	Channel 4	59 → 46	Channel 10
45 → 32	Channel 5	45 → 32	Channel 11
31 → 18	Channel 6	31 → 18	Channel 12
17 → 12	Parameters	17 → 12	Parameters
11 → 8	Diagnostics	11 → 8	Diagnostics
7 → 2	Phase switches	7 → 2	Phase switches
1	Trim flag	1	Trim flag
0	0	0	1

packaged together and prepared for USB transfer. A schematic drawing of these chains is shown in Figure 3.8.

The data precisions throughout the DSP chain were limited by the FPGA’s resources. More bits for a given data point allow higher precision, but at the cost of greater resources. The FPGA’s multipliers and random access memory blocks (BRAM) operate with 18 b; using more than 18 bits thus requires double the number of multipliers or RAM blocks. At the same time, truncating the number of bits can introduce roundoff noise Rabiner & Gold (1975). Care must be taken to ensure that the roundoff noise is less than the intrinsic noise of the data, lest the final signal-to-noise ratio be compromised.

With just 24 multipliers, effective filter design can only be achieved by multiplexing filter multiplications through a small number of multipliers. This technique allows a single multiplier, over a number of clock cycles, to act over many of a filter’s multiplications. For example, the input data rate is 2 MSps, but the FPGA clock rate is 50 MHz. Thus, a single multiplier can complete 25 multiplications for every sampling period. At the same time, the parallel data channels must each also be multiplexed. Reducing the number of channels can therefore conserve filter resources. This motivates reducing the number of filtered channels to six.

### 3.4.6.1 Multiplicative Corrections

Reducing the number of channels to six in the filter chain destroys information and precludes corresponding data corrections. In particular, any nonlinearity effects owing to differential detector diode response to the sky and reference load channels must be corrected beforehand, as well as any multiplicative corrections aimed at receiver balance. The filter chain allows for these corrections in

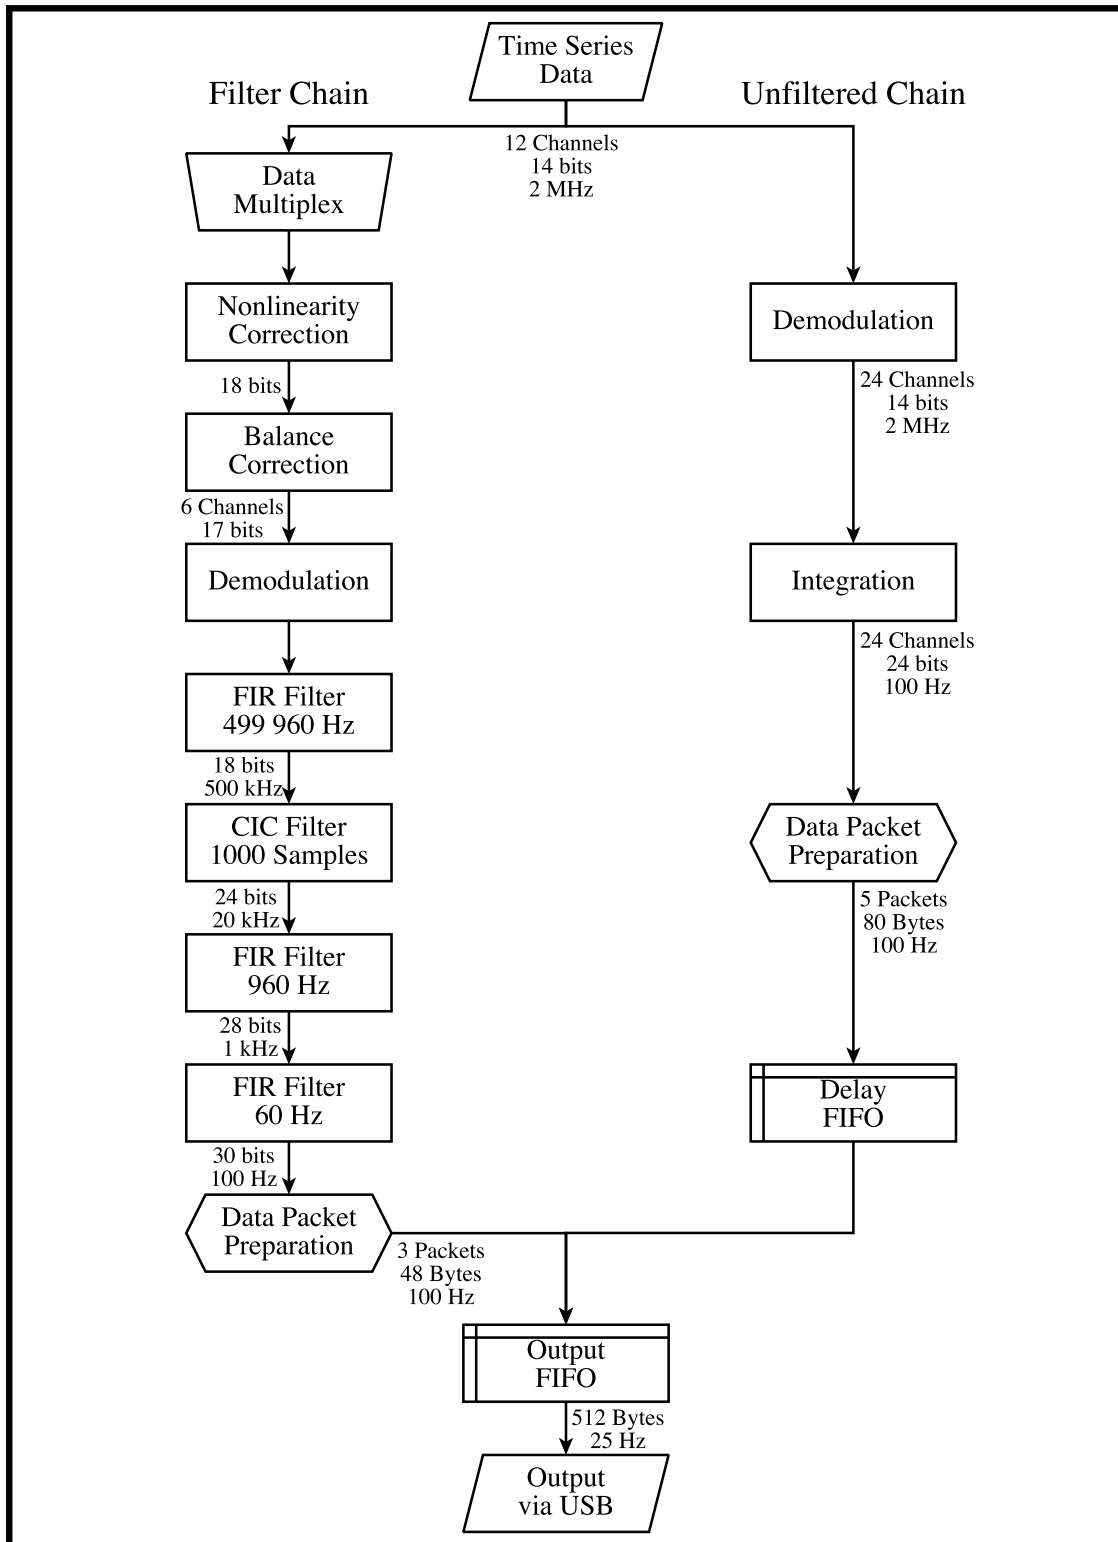


Figure 3.8: Schematic flowchart of the continuous mode DSP. Data bitwidths and rates at various stages are shown.

the `filter_condition` component, which applies pre-set nonlinearity and receiver balance corrections to the data. These corrections are enabled via the `nonlin_en` and `alpha_r_en` registers, and the corrections are contained in the `nonlin_XZ_reg` and `alpha_Y_Z_reg` registers (where  $X \in [1, 12]$ ,  $Y \in \{I1, I2\}$ , and  $Z \in \{A, B\}$ ).

Nonlinearity is modeled as a quadratic function:

$$P_{NL} = n_1 P_{ADC}^2 + n_2 P_{ADC} = (n_1 P_{ADC} + n_2) P_{ADC} \quad (3.1)$$

$P_{ADC}$  is the measured power and  $P_{NL}$  the power after nonlinearity correction. The  $n_1$  and  $n_2$  coefficients translate to the `nonlin_XA_reg` and `nonlin_XB_reg` registers. These coefficients were measured from the individual detector diodes in the lab. The measured nonlinearity coefficients are listed in Table 3.11, and an example curve is shown in Figure 3.9. On-sky tests demonstrated that, in regular observations, nonlinearity was not a detectable issue. This correction is therefore disabled by default.

Table 3.11: Measured Nonlinearity Coefficients

Channel	$n_1$	$n_2$
1	0.000336554779346	0.996896819239721
2	0.000316794934346	0.965089216446858
3	0.000326975352185	0.971295577967417
4	0.000318725389250	0.970519782777347
5	0.000361523210216	1.079131109387122
6	0.000314409183476	0.958882854926300
7	0.000342127601991	1.015515903801396
8	0.000316594604121	0.970519782777347
9	0.000331601159217	0.982156710628394
10	0.000316102884476	0.939487975174554
11	0.000317687314443	0.944142746314973
12	0.000296652640734	0.975950349107836

In the case where a nonlinearity correction is applied, care must be taken to ensure that the dynamic ranges of the data channels are preserved. Because the arithmetic is fixed-point rather than floating-point, careless choice of coefficients can lead to under- or over-flow of significant bits. The 14 b  $P_{ADC}$  inputs are made 18 b via zero-padding of the least significant bits (equivalent to multiplying by  $2^4$ ). Being two's complement, the maximum values are then  $2^{17} - 2^4$ . The maximum acceptable value of  $P_{NL}$  is  $2^{17} - 1$ .  $n_1$  should be maximal to guarantee maximum dynamic range, but no greater than  $2^{17} - 1$ . These requirements are satisfied by re-normalizing the nonlinearity coefficients via:

$$N_1 = x2^{17}n_1 \quad (3.2)$$

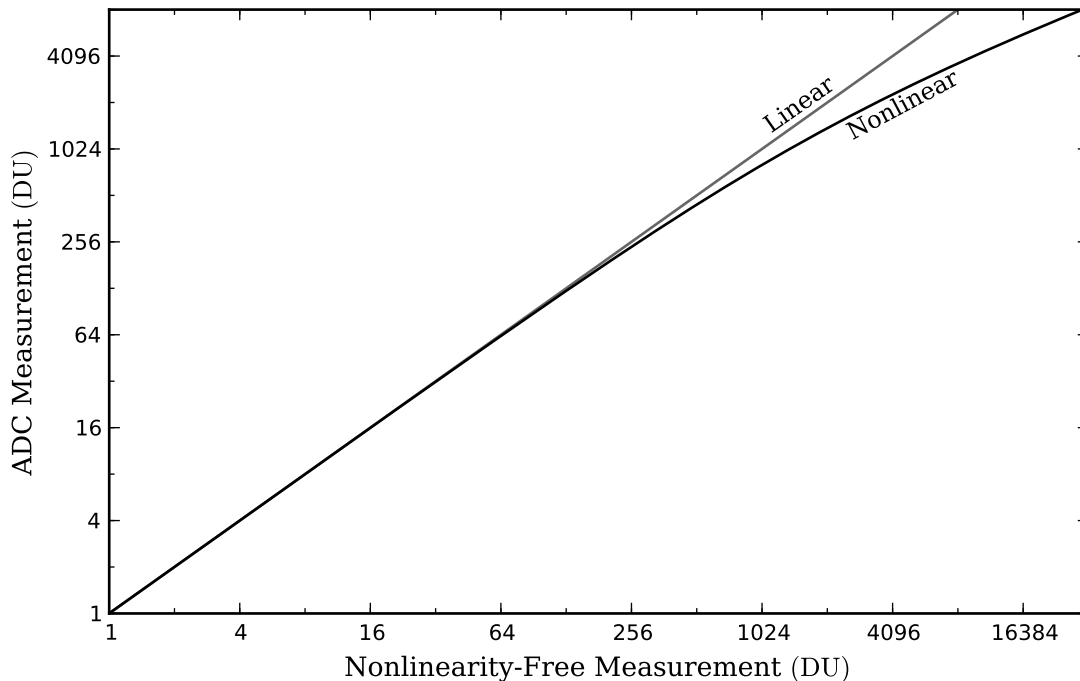


Figure 3.9: Sample nonlinearity curve for the detector diode of Channel 1. A Linear curve is shown for comparison.

and

$$N_2 = x2^4 n_2 \quad (3.3)$$

where the multiplicative constant

$$x = \min \left[ \frac{1}{2^4} \frac{2^{17} - 1}{n_1 (2^{13} - 1) + n_2}, \frac{1}{2^{17}} \frac{2^{17} - 1}{n_1} \right] \quad (3.4)$$

ensures the conditions above. A single re-normalization factor  $x$  should be used across all channels; the optimal value can be found by calculating  $x$  for each detector diode and taking the minimum of these. The backend itself is unable to check that this normalization has been correctly applied; ensuring that it has been properly calculated and applied is the responsibility of the user. That is, the  $N_1$  and  $N_2$  need to be calculated in advance and then passed to the backend `nonlin_XZ_reg` registers. Equation 3.1 is then implemented on the backend as

$$P_{\text{NL}} = \frac{1}{2^{17}} \left[ \frac{N_1 \times (2^4 P_{\text{ADC}})}{2^{17}} + N_2 \right] \times (2^4 P_{\text{ADC}}) \quad (3.5)$$

where the multiplications denoted by “ $\times$ ” are calculated via multiplexing through a single 18x18 multiplier and the output is truncated to 18 b.

Total intensity receiver balance is modeled by two parameters,  $\alpha$  and  $r$ , the first describing cross-talk between the paired, total-intensity channels, the second describing mismatch between sky and reference load temperatures. For a more detailed description, see Chapter 4. Both of these represent linear operations on the paired channels. Owing to the limited resources of the FPGA, they were combined into a single step:

$$P_{\text{sky-load}} = \left[ \frac{1+r}{\alpha} + (1-r) \right] P_{\text{NL,A}} - \left[ \frac{1+r}{\alpha} - (1-r) \right] P_{\text{NL,B}} \quad (3.6)$$

$P_{\text{sky-load}}$  is the balance-corrected difference between sky and reference load power.  $P_{\text{NL,A}}$  and  $P_{\text{NL,B}}$  are the powers measured from the paired channels (with nonlinearity correction optionally applied). The coefficients in square brackets correspond to the `alpha_Y_A_reg` and `alpha_Y_B_reg` registers. The  $\alpha$  correction acts only as a multiplicative offset if the  $r$  correction is not applied (ie,  $r = 1$ ), and the need for an  $r$  correction can be mitigated by carefully matching the sky and load temperatures. Thus, by default, these corrections are disabled.

As with the nonlinearity correction, the coefficients need to be carefully normalized. The constraints are similar to those above, though in this case it is desired that  $P_{\text{sky-load}}$  will be 17 b and so must be less than  $2^{16} - 1$ . This is guaranteed by using

$$R_A = x \left[ \frac{1+r}{\alpha} + (1-r) \right] \quad (3.7)$$

and

$$R_B = x \left[ \frac{1+r}{\alpha} - (1-r) \right] \quad (3.8)$$

where

$$x = \frac{2^{16}}{\frac{1+r}{\alpha} + |1-r|} \quad (3.9)$$

Once again, this calculation must be done before copying to the `alpha_Y_Z_reg` registers. The backend then executes Equation 3.6 through

$$P_{\text{sky-load}} = \frac{1}{2^{17}} (R_A \times P_{\text{NL,A}} - R_B \times P_{\text{NL,B}}) \quad (3.10)$$

where, again, the multiplications are multiplexed through a single multiplier. As hinted above, however, the output is truncated to 17 b.

When the above corrections are disabled, as is the default, the multiplications are still performed. In this case, however, the values  $n_1 = 0$  and  $n_2 = r = 1$  are plugged into the above normalization schemes. Thus, whether nonlinearity and balance corrected or not, the data have now been pairwise subtracted and are ready for demodulation.

### 3.4.6.2 Demodulation

In the absence of modulation, astronomical signal appears in the time series at frequencies determined by the time needed for the telescope to scan across a spatial resolution element. For C-BASS, this is below 40 Hz. With modulation, the astronomical signal is mixed to the phase switch frequency and its harmonics. This is useful, as many sources of systematic noise entering the data post-modulation will tend to affect lower frequencies. So, when demodulation brings the astronomical signal back down below 40 Hz, it also mixes systematic noise up to the modulation frequency. The systematic noise can then be rejected through low-pass filtering.

Demodulation itself is simple. The modulating signal, which is generated in the backend, is delayed by the hardware latency (described in Section 3.4.4) and any firmware latency experienced by the time series data up to this point. Thus synchronized, the pairwise subtracted data have their signs flipped when in modulation state 01 or 10. The six demodulated, 17 b data are then ready to be filtered.

In the unfiltered chain, demodulation is also performed. In this case, however, the individual modulation states are preserved and no pairwise subtraction takes place. Thus, the unfiltered chain maintains 24 data channels through integration.

### 3.4.6.3 Filtering and Integration

Simple integration of the data preserves the astronomical signal, but also allows spurious, out-of-band signals to survive (including, for example, 60 Hz interference from power lines). When the data are integrated and downsampled, these spurious signals are aliased into the resulting time series and can be confused for astronomical structure. Filtering the data as part of the integration and downsampling can suppress these out-of-band signals to give a lower-noise time series.

The final data rate of 100 Hz requires frequencies above 50 Hz to be rejected. Crafting a single-stage filter capable of meeting this goal for 2 MHz data is simply not possible with the hardware at hand. It was therefore necessary to perform this filtering in four stages. After each filter, the data are downsampled, allowing for greater efficiency in multiplier utilization for those filters that follow.

The filters are all of the finite impulse response (FIR) variety, meaning that the filtered data are linear combinations of a finite number of input samples. Three of the filters (the first, third, and fourth) are simple, decimating, low-pass filters with pre-defined filter coefficients. The second filter is a cascaded integrator-comb (CIC) filter with fractional downsampling. The choice of filters was motivated by considerations of the necessary fixed-point precision and the multiplier and memory resources required by various types of filters. A further consideration was aliasing: after each filter and decimation, it is essential that any frequencies aliased between 0–40 Hz be sufficiently attenuated to prevent the survival of systematic noise.

The CIC filter had the largest influence on the overall design. The fractional downsampling required by this filter represents the interface between the digital backend’s 50 MHz clock and the GPS-supplied 1PPS clock. Time series data enter this filter on the former clock, and exit synchronized to the latter. However, this fractional downsampling relies upon downsampling signals originating in the `pps_control` component, which require a custom design for the filter. Such a filter operates by keeping a running mean of the data time series, with newer data points being added and older data points being subtracted. This requires keeping track of the older data points, which is done through FIFO memory. It is here, then, that we arrive at the critical issue: a single block of RAM on this FPGA can only hold 1024 18 b values. BRAM usage is thus optimized by placing this filter early enough in the chain that 18 b are sufficient, late enough that a length of 1000 samples is adequate, and followed by strong enough filtering that a sharp cutoff at 50 Hz is still achievable.

The resulting design is that shown in Figure 3.8, with the filter properties listed in Table 3.12. The filters are characterized by their passband frequencies  $f_{\text{pass}}$ , stopband frequencies  $f_{\text{out}}$ , the passband ripple  $A_{\text{pass}}$ , and the stopband attenuation  $A_{\text{stop}}$ . The input frequencies, filter orders, decimation rates, and output bitwidths are also listed.

Table 3.12: Individual Filter Details

Filter Number	Filter Type	$f_{\text{in}}$ [kHz]	$f_{\text{pass}}$ [Hz]	$f_{\text{stop}}$ [Hz]	$A_{\text{pass}}$ [dB]	$A_{\text{stop}}$ [dB]	Order	Decimation	Output Bitwidth
1	FIR	2 000	40	499 960	$10^{-6}$	100	14	4	18
2	CIC	500	...	...	...	...	1000	25	24
3	FIR	20	40	960	$10^{-4}$	100	118	20	28
4	FIR	1	40	60	$10^{-4}$	80	238	10	30

The first filter was designed under the limitations of using a single multiplier/BRAM pair and providing the CIC filter with 18 b inputs. The optimal filter decimates by a factor of four, has a passband below 40 Hz, and a stopband above 499 960 Hz (with 100 dB of rejection). This is a rather broad filter, only requiring 15 coefficients. The output, 18 b data are output at a 500 kHz rate. The CIC filter follows, with a length of 1 000 samples each and a decimation factor of 25. The CIC filter’s frequency response is a `sinc` function, which falls off only as  $1/\nu$  but has nulls at multiples of 5 kHz. The choice of decimation rate thus places those frequencies aliased to the passband at nulls in the CIC response. Further, any periodic artifacts due to the modulation itself will be nulled, as well. As mentioned above, the downsampling rate is synchronized to the 1PPS timing so that the output data have a rate of exactly 20 kHz. The filter length recommends an output precision of 24 b. For this filter, each channel requires its own BRAM, and so a total of six on the FPGA are used.

The last two filters decimate by factors of 20 and 10 to reach the final data rate of 100 Hz. Achieving the necessary stopbands at 940 and 60 Hz, with attenuations of 100 and 80 dB, required

filter lengths of 119 and 239, respectively. Output precisions of 28 b and 30 b were used. These filters were achieved on the FPGA using 3 multipliers plus 4 blocks of RAM for the first, and 3 plus 6 for the second. Altogether, the filter chain required use of 17 of the FPGA’s 24 multiplier/BRAM pairs.

As mentioned in Section 3.4.4, a small number of modulation frequencies are compatible with this filter chain. There are two constraints. First, the modulation period must be divisible by 16 to accommodate the four phase states and the first decimation by four. Second, the CIC period must be divisible by one quarter the modulation period in order to guarantee removal of modulation artifacts. Eight frequencies meet these requirements. They are listed in Table 3.13 along with the `switch_reg` register values which choose them. The default was chosen as the lowest of these in order to minimize attenuation from the detector diode video bandwidths.

Table 3.13: Allowable Modulation Frequencies

<code>switch_reg</code>	Switch Frequency [kHz]	Modulation Frequency [kHz]
00000001	0.5	1
00000010	1.0	2
00000100	2.5	5
00001000	5.0	10
00010000	12.5	25
00100000	25.0	50
01000000	62.5	125
10000000	125.0	250

Design and simulation of the full filter chain was performed using the `MATLAB` software suite. This software allowed for the easy calculation of filter coefficients given a wide range of desired filter properties, while allowing for the full filter chain (including effects of bit precision) to be directly simulated and evaluated. The full spectrum (in modulated frequency space) is shown in Figure 3.10. The modulated sky signal appears at 1 kHz modulation frequency and its odd harmonics. Any spurious signals outside of this range are attenuated with greater than 80 dB.

The performance was further evaluated with a second simulation, shown in Figure 3.11. Using an input time series with noise properties similar to that of the C-BASS receiver and a 60 Hz signal with an amplitude  $\sim 10\,000$  greater than normal, the spectrum before and after each filter was calculated. The rejection of the interfering signal is dramatic. Effects of intermediate bit precision between the individual filters were tested using this simulation, confirming that the precisions chosen increased the noise of the output time series by less than 1%.

The desired signals do not pass exit the filter chain unscathed. The effective filter length spans 29






Figure 3.10: Frequency response of the full DSP chain. Modulated frequency refers to the frequency prior to demodulation. The passbands are centred upon the modulation frequency and its harmonics, where the sky signal appears.




Figure 3.11: Frequency responses for the four filters. The 60 Hz pickup has been mixed up to kHz frequencies by the demodulation. Grey is the data spectrum prior to filter, black is after.




Figure 3.12: Autocorrelation function of the filter chain, normalized to 1 at  $\Delta t = 0$ . Note that  $\Delta t$  is in units of 10 ms samples.

100 Hz samples, so the filtered data are autocorrelated across this timescale. This autocorrelation is a property of the filter chain and will manifest most strongly as oscillations preceding and following sudden changes in power (such as when the noise diode is turned on or off). Analytical calculation of the autocorrelation function is precluded by the fractional downsampling in the CIC filter, so it is calculated numerically instead. This calculation was performed by injecting white, Gaussian noise into the simulated filter chain and finding the autocorrelation of the result. The function is plotted in Figure 3.12, the values listed in Table 3.14, and an example of the effect in real C-BASS data shown in Figure 3.13.

The unfiltered processing chain uses an integration scheme with fractional triggering, similar in concept to the CIC filter. In this case, the idea is to accumulate  $10^4$  samples for each of the 24 demodulated channels, with the accumulations being triggered by the 1PPS-synchronized counter contained in `pps_control`. The resulting time series are produced at exactly 100 Hz. An output precision of 24b was chosen, which is less than that from the filters. Although concern about truncation noise is reasonable, this precision is constrained by the backend’s output bandwidth. Because these data are used for receiver balance monitoring and other diagnostic tests, the concession in output precision is not expected to adversely affect the final C-BASS data products.

The demodulation of the data into 24 channels leaves these data susceptible to systematic effects

Table 3.14: Filter Chain Autocorrelation Function

$\Delta t$	$\langle P(t), P(t + \Delta t) \rangle$
0	1.0000
1	0.0465
2	-0.0389
3	0.0346
4	-0.0305
5	0.0255
6	-0.0203
7	0.0159
8	-0.0115
9	0.0087
10	-0.0056
11	0.0033
12	-0.0023
13	0.0005
14	-0.0007




Figure 3.13: Effect of the filter autocorrelation function. At sample 0, the noise diode is switched on. Ripples are introduced in the filtered data (black), not present in the unfiltered data (grey).

during the switching of modulation states. The filtered chain is designed to mitigate this problem; the unfiltered chain has no similar, natural recourse. Optional, pre-integration flagging of afflicted data is instead allowed. A constant number of data points can be so flagged and trimmed following each phase state transition. The number of trimmed data points is stored in the `trim_reg` register, and is set to 0 by default.

Although the individual integrations are ostensibly independent, the mismatch between backend and GPS timing implies that a few integrations per second will overlap by a single 2 MHz data point. In principle, this implies cross-correlation between the integrations, though it will be at the 0.01 % level and thus unimportant.

In comparison to the filter chain, the unfiltered integrations are calculated with a much shorter latency, so the integrated time series needs to be stored prior to packaging with the filtered data. The usage of this storage is monitored and made available in Packet #7 of the continuous data (see Table 3.8). This storage requires the use of a single block of RAM. The full DSP chain therefore uses 75 % of the FPGA's multiplier and RAM resources.

The redundancy in the paired DSP chains allows for testing of the filter performance and fidelity. Figure 3.14 compares the pipeline-reduced, smoothed time series of both during a regular observation scan to the 5 GHz sky model of de Oliveira-Costa et al. (2008). The structures in the time series are due to the astronomical signal. Figure 3.15 shows the power spectra of the same data, though without smoothing. The reduction in aliased 60 Hz pickup at 40 Hz is clear.

### 3.4.7 USB Microcontroller

USB access was granted through the USB microcontroller. This chip is built around an industry-standard 8051 microprocessor which can be programmed in the C language. Programming was done using the `Kiel` programming environment, which produced a compiled, hexadecimal file, sent to the microcontroller from the control computer upon startup of the backend. The microcontroller contains no non-volatile memory, so this programming file must be sent to the backend after every power cycle.

Fortunately, it was not necessary to write this firmware from scratch. The firmware provided by Oxford with the hardware needed only superficial modifications to accommodate the backend's registers. A new C driver was written to allow control from a UNIX platform (the Oxford system being run from the Windows operating system) and was incorporated into the telescope control system. This driver worked with the microcontroller firmware to define a number of low-level control commands. These are listed in Table 3.15. The commands are interpreted by the microcontroller firmware and sent to the primary FPGA via its parallel, 8 b `pc` and `pe` inputs.

The microcontroller is clocked at 25 MHz, with the clock signal being provided by the backend's primary FPGA. This different clock speed represents a second clock domain within the primary




Figure 3.14: A single  $I_1$  azimuth scan from January 15, 2012. Data are pipeline-reduced, but not flux calibrated. Shown are filtered data, unfiltered data, and a simulated scan across the sky model of de Oliveira-Costa et al. (2008). The remarkable similarity between the three attests to the fidelity of the filter chain.

Table 3.15: Backend Firmware Commands

C Driver Command	USB Microcontroller Internal Command	Description
<code>cbass_connect</code>		Establishes a USB connection
<code>cbass_disconnect</code>		Ends the USB connection
<code>cbass_load_hex</code>		Sends the USB firmware
<code>cbass_ADC_setup</code>	ADCSETUP	Resets the FPGA memory and registers
<code>cbass_get_data</code>		Requests a bulk transfer of C-BASS data
<code>cbass_FPGA_reset</code>	FPGA.RESET	Toggles the FPGA <code>reset</code> register
<code>cbass_USB_reset</code>		-disabled-
<code>cbass_EP8Fifo_reset</code>	EP8FIFORESET	Resets the USB microcontroller buffers
<code>cbass_bit_set</code>	BIT_SET	Overwrites to a single bit in a register
<code>cbass_store_data</code>	STORE_DATA	Overwrites a register
<code>cbass_nonlin_C</code>	NONLIN_COEFF	Sets a nonlinearity coefficient
<code>cbass_alpha_C</code>	ALPHA_COEFF	Sets a balance coefficient
<code>cbass_r_C</code>		-disabled-
<code>cbass_burst_L</code>	BURST_LENGTH	Sets the burst length
<code>cbass_sw_P</code>	SWITCH_PERIOD	Sets the modulation period
<code>cbass_tr_L</code>	TRIM_LENGTH	Sets the trimming length
<code>cbass_acquire</code>	ACQUIRE	Sets the <code>adc_read</code> register
<code>cbass_continuous</code>	CONTINUOUS	Sets the <code>cont_mode</code> register
<code>cbass_sim</code>	SIM_MODE	Sets the <code>simulate_en</code> register
<code>cbass_trigger</code>	TRIGGER	Triggers a burst of data
<code>cbass_nonlin</code>	NONLIN_EN	Sets the <code>nonlin_en</code> register
<code>cbass_alpha</code>	ALPHA_R_EN	Sets the <code>alpha_r_en</code> register
<code>cbass_noise</code>	NOISE_ON	Sets the <code>noise_on</code> register
<code>cbass_switch</code>	SWITCH_ON	Sets the <code>switch_en</code> register
<code>cbass_switch_alt</code>	SWITCH_ALT	Sets the <code>switch_alt</code> register
<code>cbass_switch_full</code>	SWITCH_FULL	Sets the <code>switch_full</code> register




Figure 3.15: Power spectra of filtered (black) and unfiltered (grey)  $I_1$  data from January 15, 2012. 60 Hz pickup and harmonics appear aliased at 20 Hz and 40 Hz without filtering.

FPGA, which was bridged through a dual-clock FIFO, requiring use of two of the FPGA's blocks of RAM. After crossing this clock domain interface, the data are sent to the microcontroller via 16 b of the FPGA's parallel, 18 b output `usb_in`. Once the microcontroller has received all 512 B for a given bulk transfer, the microcontroller makes those data available to the control system and begins accepting data for the next bulk transfer. In this, the microcontroller's behaviour is the same whether continuous or burst data are being transferred.

### 3.5 Operation in the Field

The C-BASS digital backend was installed on the Northern antenna in December of 2009. At that time, its DSP chain included the unfiltered operation, but not the filter chain described in this chapter. Due to the presence of strong 60 Hz and harmonics in the data, the DSP chain was revised to include the filters. The updated firmware was installed in October of 2011, and the digital backend has been performing well to date.

## Chapter 4

# The C-BASS Pipeline

### 4.1 Context

The C-BASS time series data received from the digital backend need extensive processing and calibration before they are ready for science use. The C-BASS pipeline fills this role. It was implemented primarily in the `MATLAB` software suite and can be called from the command line. It has been designed to be widely configurable and customizable, allowing users to experiment with a variety of data reduction strategies. The final products of the pipeline are calibrated, all-sky maps in the HEALPix format (Górski et al. 2005).

The pipeline was written collaboratively, with various aspects written and tested by different members of the C-BASS collaboration. Careful version control using the `CVS` software package and documentation via an online Wiki greatly assisted in this effort. An overview of the entire pipeline is given in this chapter. Detailed descriptions of the cold cycle and RFI corrections, which were the responsibility of the author, are found in Chapters 5 and 6.

As a software resource, the pipeline is continuing to evolve. This chapter describes the pipeline as it was at the time of writing. Future iterations of the pipeline may diverge significantly from the present description.

The pipeline has four main goals: astronomical calibration, suppression of systematic errors, evaluation of data quality, and map making. These goals are divided into a number of individual tasks, each aimed at a particular aspect of the larger picture. Most of the time-series processing occurs in the `MATLAB` portion of the pipeline, with the remaining processing and map making being handled by the `Descart` software package (Sutton et al. 2009). The `MATLAB` portion is described in Section 4.2, while `Descart` is discussed in Section 4.3.



## 4.2 The MATLAB Pipeline

The MATLAB pipeline is called via a MATLAB function `reduceData()`. This function takes the raw, time-series data and applies the various reduction steps. It can be run interactively or non-interactively, with reduction parameters set via a text file. Each task of the pipeline is called via wrapper scripts, which check that necessary pre-requisites have been satisfied, take care of necessary housekeeping, and produce summary plots (if desired). The reduced data are returned at the MATLAB command line and optionally written to disk in a binary FITS table. Plots are saved as PNG images.

The main input and output of these reductions are in the form of MATLAB data structures. These contain many data vectors and matrices representing time series of astronomical and ancillary data. Most important of these are the times of each observation, the Stokes  $I$ ,  $Q$ , and  $U$  measurements, the antenna pointing records, and the receiver health records. In particular, both the unfiltered 24-channel and the filtered six-channel astronomical data are kept (see Section 3.4.6). During the course of the MATLAB reductions, more vectors and matrices are added as needed.

A text file, typically denoted by suffix `.red`, contains the parameters used by each reduction step. When running the pipeline non-interactively, this file also determines which reduction steps will be run and in what order. In this file one specifies whether to or not to save plots to disk and whether or not to operate interactively (although this latter option can be overridden at the command line).

The pipeline has five inputs: the raw data structure, the `.red` file, the interactive option, the output directory prefix (which can be a full path), and the name of the observing schedule being reduced. Only the first input is strictly necessary, although the `.red` file is needed if one doesn't want to reduce with default parameters. The observing schedule name is only used if one is using the pipeline's ability to produce HTML files. Choosing to produce HTML files is useful as a convenient means for viewing the pipeline results, and is an option often used when sharing pipeline results within the collaboration. Aside from plots and the FITS table saved to disk, the pipeline returns the processed and augmented data structure to the command line.

Two MATLAB functions are useful for running the pipeline on large quantities of data: `day_reduce()` and `period_reduce()`. These run the pipeline on each observing schedule observed by the telescope over a given day or time period, respectively. These can be used to generate HTML files (and prompt `reduceData()` to do so) for the purpose of data quality monitoring.

Several of the pipeline reduction steps rely upon previously reduced data. The results of these previous reductions are stored in text files which can be read by the pipeline as necessary. The affected reduction steps include the  $\alpha$  and 1.2 Hz corrections, the RFI flagging, and the astronomical calibration. In the latter case, it is necessary to run the pipeline itself on calibration data prior to running it on survey data. For the others, dedicated functions must be used. These are described

in the reduction step descriptions that follow.

The reduction steps of the pipeline have been designed to be run in a specific order. This order is shown in Figure 4.1. The additional functions needed to populate the ancillary data files mentioned in the previous paragraph are also shown in this figure. The reduction steps are described briefly below, using the data from January 17, 2012, 11:50:00 UTC to 12:40:00 UTC, which consists of two all-sky scanning observations bracketed by two calibration observations. When plots are produced in the pipeline, they are reproduced here with cosmetic modifications. When time series are plotted, a single point is plotted for every 2, 5, or 10 seconds of data as required for clarity.

### 4.2.1 Initial Flags and Plots

Upon beginning the pipeline, the data are immediately checked for a number of flags. These are: whether the data have been properly received (“FRAME”), the cold receiver is below 20 K (“RX COLD”), the antenna elevation is above  $10^\circ$  (“ELEVATION”), the weather station data have been properly received (“WEATHER”), the backend behaviour is nominal (“CONTINUOUS,” “DATA BACKLOG,” “BACKEND CLOCK,” and “INTEGRATION SHORTFALL”), the servo and cryogenic controller timing are nominal (“SERVO/CRYO 1PPS”), and whether the scans are too close to the Sun, Moon, or the observatory control building (“SUN PROXIMITY,” “MOON PROXIMITY,” and “CONTROL BUILDING”). The flagging results (in terms of number of 100 Hz samples flagged) are displayed to the user and optionally saved to disk. The results for the example data set are shown in Table 4.1.

Table 4.1: Example Initial Flagging Results.

FLAGGING REPORT	
Out of 1739900 frames, the number flagged follows:	
ANTENNA FLAGS	
FRAME:	0
RX COLD:	0
ELEVATION:	0
WEATHER:	0
CONTINUOUS:	0
DATA BACKLOG:	0
BACKEND CLOCK:	0
SERVO/CRYO 1PPS:	600
INTEGRATION SHORTFALL:	0
SUN PROXIMITY:	0
MOON PROXIMITY:	7191
CONTROL BUILDING:	0
TOTAL:	7791

Next, a time series of the backend-filtered channel of Stokes  $I_1$  is plotted, with the line colours

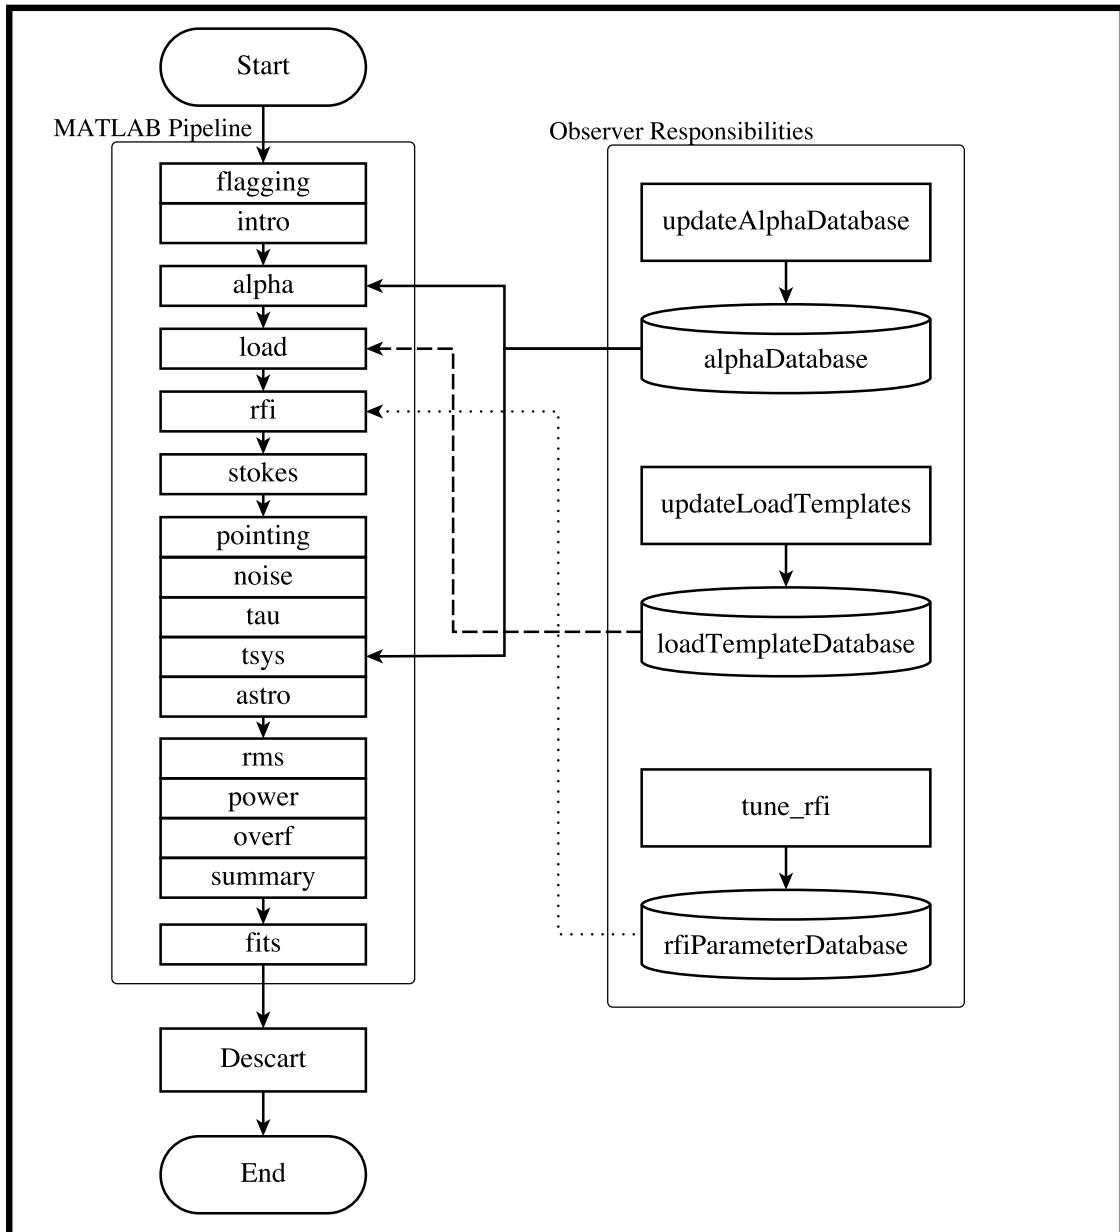


Figure 4.1: Flowchart overview of the C-BASS pipeline. The MATLAB pipeline is outlined on the left, with individual tasks grouped as in the main text's subsections. Outside of the main pipeline are shown the MATLAB scripts and databases that must be run by the observer. Also shown is the `Descart` map maker.

coded for the various observing sources. The elevation and azimuth of the antenna are also plotted. This plot is reproduced in Figure 4.2. A seasoned observer can use this figure to verify that the observing schedules performed as expected. Most of the data are labeled “NCP,” indicating that the antenna is nominally pointed at the North Celestial Pole, with the all-sky azimuth scans being achieved through azimuth offsets from this point. Calibrator sources are also seen (“J1230+123,” “J0927+390,” and “CygA”). Calibration skydips are denoted by “current” and “skydip,” while noise diode firings are labeled as “noise.” The constant elevation during all-sky scans is clearly seen in the elevation plot, with the azimuth scans themselves shown in the azimuth plot. Skydips, in which the antenna is scanned in elevation in order to measure atmospheric opacity, are also easily seen in the elevation plot.

Finally, the time series of all six channels are plotted, with line colours coded for the intent of observations. These are reproduced in Figure 4.3. As with the previous plot, a seasoned observer can verify schedule performance. Data from primary science observations are labeled “source.” Calibrator observations are labeled as such, but also include pointing observations labeled as “radio\_point\_scan.” Times when the noise diode is on are denoted “noise,” while the periods immediately preceding and following are “noise\_event.” Observations of blank sky, used in deprecated corrections for the 1.2 Hz and 60 Hz interfering signals, and skydips are labeled “blank” and “skydip,” respectively. Times when the preceding conditions do not apply are simply called “transition.”

#### 4.2.2 $\alpha$ Corrections

It is possible in the polarimeter that there will be signal leakage between the total intensity sky and load channels, as well as between the polarization channels. The noise diode can be used to measure and correct for these leakage effects. This correction takes place in the `alpha` task of the pipeline.

In the total intensity case, the measured sky and load channels are linear combinations of the isolated sky and load signals. The linear combinations can be parameterized by quantities  $\alpha_i$ , such that the measured power from one of the detector diodes is

$$P_{D,\pm} = G_i [(1 \pm \alpha_i) T_{\text{Sky}} + (1 \mp \alpha_i) T_{\text{Load}} + T_{\text{Noise}}], \quad (4.1)$$

where  $G_i$  is the gain and  $T_{\text{Sky}}$ ,  $T_{\text{Load}}$ , and  $T_{\text{Noise}}$  are the temperatures of the sky (including the atmosphere, the CMB, and any other signals entering the feed horn), the cold reference load, and the instrument noise, respectively. The  $\pm$  and  $\mp$  denotes when a given detector diode is nominally a sky measurement (upper) or a load measurement (lower). In the ideal case,  $\alpha_i = 1$  and the equations reduce to:

$$P_{D,+} = G_i (2T_{\text{Sky}} + T_{\text{Noise}}) \quad (4.2)$$




Figure 4.2: Plot of observed sources. The top panel shows channel  $I_1$ , colour coded for observed sources. For clarity, one data point is plotted per 5 second interval. The middle and lower panels show the antenna elevation and azimuth during these observations. For clarity, one data point is plotted per 10 second interval.




Figure 4.3: Plot of observing intent. Each panel shows one of the six polarization channels. Lines are colour coded to match observing intent. For clarity, one point is plotted for every 10 second interval.

and

$$P_{D,-} = G_i (2T_{\text{Load}} + T_{\text{Noise}}). \quad (4.3)$$

The  $\alpha_i$  can be calculated using the noise diode. The noise diode is coupled to the signal prior to the hybrid that couples the reference load, so when turned on it only adds power to  $T_{\text{Sky}}$ . The differences between noise diode on and off states for the detector diodes are then

$$\Delta P_{D,\pm} = G_i (1 \pm \alpha_i) T_{\text{ND}}, \quad (4.4)$$

in which  $T_{\text{ND}}$  is the noise diode temperature. Taking the ratio between the states and completing the algebra,

$$\alpha_i = \frac{\Delta P_{D,+} - \Delta P_{D,-}}{\Delta P_{D,+} + \Delta P_{D,-}} \quad (4.5)$$

One can then return to equation 4.4 to calculate the gain  $G_i$  relative to  $T_{\text{ND}}$ . Note that  $T_{\text{ND}}$  remains unknown at this point, and will not be measured until Section 4.2.6. With  $\alpha_i$  and  $G_i$  in hand, the sky and load signals can be recovered:

$$\frac{T_{\text{Sky}} + T_{\text{Noise}}}{T_{\text{ND}}} = \frac{1}{G_i} \left[ \left(1 + \frac{1}{\alpha_i}\right) P_{D,+} + \left(1 - \frac{1}{\alpha_i}\right) P_{D,-} \right] \quad (4.6)$$

and

$$\frac{T_{\text{Load}} + T_{\text{Noise}}}{T_{\text{ND}}} = \frac{1}{G_i} \left[ \left(1 - \frac{1}{\alpha_i}\right) P_{D,+} + \left(1 + \frac{1}{\alpha_i}\right) P_{D,-} \right] \quad (4.7)$$

It is important to note two points at this time. First, it is the difference between equations 4.6 and 4.7 that is of astronomical interest, and this difference preserves  $\alpha_i$  as a purely multiplicative factor:

$$\frac{T_{\text{Sky}} T_{\text{Load}}}{T_{\text{ND}}} = \frac{1}{G_i \alpha_i} (P_{D,+} - P_{D,-}). \quad (4.8)$$

Thus, unless an  $r$  factor correction is applied (see Section 4.2.4), the total intensity  $\alpha$  correction is degenerate with gain and is unnecessary. Second, this correction can only be applied prior to taking the difference between sky and load and cannot, therefore, be applied in the pipeline to data filtered in the digital backend. If a total-intensity  $\alpha$  correction is desired for filter-mode data, then it must be calculated in advance from unfiltered data and applied in real time to the filtered data on the digital backend itself (see Section 3.4.6).

For the linear polarization channels, leakage may be described as a Stokes vector modification via a Mueller matrix. The underlying principle is that the noise diode events will ideally only appear in the  $U$  channels and not in the  $Q$  channels. The deviation from this behaviour thus allows measurement of the relevant Mueller matrix elements and the resulting correction for instrumental modifications of the Stokes  $Q$  and  $U$  elements. In the C-BASS pipeline, this correction is also referred to as an  $\alpha$  correction. In contrast to the total intensity  $\alpha$  correction, however, the polarization

correction is important regardless of  $r$  factor correction and can be applied post-detection to data filtered in the backend.

Application of these corrections requires calculation of the  $\alpha$  and  $G$  values. The strategy adopted in the C-BASS pipeline is to calculate these within a few days of observation as part of routine data verification. An automated script calculates the corrections and plots them for the observer, who is then invited to flag any outlying or otherwise suspicious points. The resulting values are written to a set of text files referred to collectively as the  $\alpha$  database. The system temperature is also estimated at this stage. It is assumed to be proportional to the  $I$  channel DC levels, as perfect detector diodes will have output DC levels proportional to input RF power levels. When applying the  $\alpha$  corrections and calculating  $T_{\text{sys}}$ , the values from these database files are used.

Three modes of  $\alpha$  correction are available in the pipeline: filtered, classic, and polonly. The first mode is applied to data filtered in the backend and is necessarily limited to the Mueller matrix correction for linear polarization. The classic mode can only be applied to unfiltered data from the backend. In this case, the corrections are applied to both total intensity and to linear polarization. An  $r$  factor correction can then be applied later, as discussed in Section 4.2.4. The polonly mode is like the classic mode, but takes the difference between sky and load channels after the  $\alpha$  correction is applied. No  $r$  factor correction is applied before subtracting load from sky so, as discussed above, the total intensity  $\alpha$  correction will have no effect after astronomical calibration.

If total intensity  $\alpha$  and  $r$  corrections are desired for the backend-filtered data, then they must be calculated in advance and applied in real time. The ability to do so is programmed into the digital backend, as discussed in Section 3.4.6.

### 4.2.3 Cold Cycle Correction

The C-BASS data are subject to a 1.2 Hz interfering signal. This frequency matches that of the mechanical refrigeration cycle of the cold receiver, and tests have shown it to be consistent with a microphonic origin in receiver components, cables, and/or connections. Astronomical signal exists at this frequency in the detected time series, so suppression is necessary to achieve a high fidelity map. This suppression is implemented in the pipeline's `load` task.

The suppression of this signal and its harmonics is a significant portion of the work in this thesis and is discussed in detail in Chapter 5. The algorithm relies upon the assumptions that the interference is well-fit by a template which is stable over many days and that the phase of the cooling cycle can be precisely measured by a temperature sensor within the cold receiver. The former assumption has appeared to hold for most of the observations to date, but has at times failed and thus limited the effectiveness of this algorithm.

The algorithm works by measuring the phase of the cold cycle across the time series and then interpolating the template over the observing time. The amplitude of the full-length template is fit



to the data via linear regression. The resulting phase- and amplitude-fitted template is subtracted from the time series. The templates are calculated and amplitude-fitted independently for each polarization channel.

As with the  $\alpha$  corrections, the templates must be calculated in advance. This is also done as part of routine data verification and the results saved in text files, this time referred to as the load template database.

After the 1.2 Hz correction has been applied, power spectra plots are generated to show the results. The plots for the data used in this chapter are reproduced in Figure 4.4. These show that the 1.2 Hz signal visible in the raw data (grey curves) is significantly reduced by the correction (black curves).

#### 4.2.4 $r$ Factor Correction

The temperature match between sky and reference load in the receiver is intended to provide suppression of correlated noise. Correlated noise arises as both multiplicative and additive components, both of which will be subtracted optimally if the sky and load temperatures are identical.

A mismatch in sky and load temperature will cause the multiplicative noise suppression to be imperfect. It is possible to improve this suppression by multiplying the load channel by a factor  $r$  prior to subtracting from the sky channel. This is called an  $r$  factor correction (Bersanelli & Mandolesi 2000) and can optionally be applied in the C-BASS pipeline to data unfiltered by the backend. If an  $r$  factor correction is desired for backend filtered data, then it, like the  $\alpha$  correction, must be applied in real time in the backend (see Section 3.4.6). In the former case, this correction is applied in the pipeline’s `stokes` task.

Although an  $r$  factor correction improves suppression of multiplicative noise, it increases additive noise. The greatest reduction of correlated noise requires one to balance these concerns. Further, if  $r > 1$ , then the thermal noise from the reference load is increased. Thus, using an  $r$  factor correction represents a trade-off between different types of noise.

After applying the  $r$  factor correction in the pipeline, a plot is generated comparing the uncorrected to the corrected data. This is reproduced in Figure 4.5. The uncorrected data are raw, unfiltered by the backend, while the corrected data now have both  $\alpha$  and  $r$  factor corrections applied. Thus there are differences in both gain and offset.

#### 4.2.5 RFI Flagging

Radiofrequency interference (RFI) from human activity can be much brighter than astronomical signal and dominate the received data. A number of international and federal agencies allocate spectrum bandwidth for use, with some regions of the spectrum reserved for astronomical use.




Figure 4.4: Power spectra of data before (grey) and after (black) the 1.2 Hz correction.




Figure 4.5: Plot showing the effects of the  $\alpha$  and  $r$  corrections. The corrections for all channels have clearly modified the data in terms of both gain and DC offset. For clarity, the time series show one data point for every 10 seconds.

There is such reserved spectrum at 5 GHz, although it is only 10 MHz wide (Panel on Frequency Allocations and Spectrum Protection for Scientific Uses 2007). In seeking a 1 GHz bandwidth, we therefore have no legal recourse for RFI reduction. Mitigation must come through instrument design and data analysis.

Early tests showed several fixed sources of RFI, including geosynchronous satellites, fixed transmitters in nearby towns, the wireless link to the nearby CARMA Observatory, and local wireless internet transmitters. The first three of these were avoided by augmenting our RF filters (using both notch and bandpass filters). The last was solved by turning off wireless network transmitters in the nearby observatory buildings.

Before the additional notch and bandpass filters were installed, it was necessary to spatially flag those data observed too close to geosynchronous satellites or ground-based transmitters. This capability persists in the pipeline, but no longer flags data during regular observations.

Transient sources of RFI include aircraft and non-geosynchronous satellites. RF filtering of these is less effective, as they may exhibit a wider range of transmission frequencies and, during times when they are not interfering, such filtering would lower the instrument's sensitivity. On the other hand, time domain filtering has the potential to preserve instrument sensitivity while removing the transient interference from the reduced time series. Such time domain filtering occurs in the pipeline during the `rfi` task.

The time domain filtering of RFI forms a significant portion of this thesis and is described in detail in Chapter 6. A brief description is given here. The underlying assumptions are that any transient RFI is highly polarized (linearly or circularly) and that it varies on timescales inaccessible to the antenna's beam crossing time. The former point is consistent with the understanding that the vast majority of human-caused radio transmissions are intrinsically polarized. The latter point arises from the fact that transient transmissions are likely to be modulated and traveling across the sky, which is not true for the astronomical signals of interest.

The algorithm works by finding the standard deviations of the time series on a number of different timescales and across the different polarization channels. If, at any given time, the standard deviation of the data is greater than average, if the standard deviation is significantly greater at shorter timescales than at longer timescales, and if these effects show significant polarization, then the outlying and surrounding data are flagged.

The RFI flagging depends upon a number of parameters, including how great of a deviation in standard deviation is needed to raise an alarm and how polarized the signal must be. The tuning of these parameters determines what fraction of transient RFI is missed and how much non-RFI afflicted data are flagged. There is clearly a trade-off between these undesirable effects: a more aggressive RFI flagging will catch more of the RFI but also flag more good data, while a more conservative RFI flagging will leave good data unharmed but miss more RFI. The proper tuning of

the flagging parameters is thus of critical importance.

The optimal flagging parameters are sensitive to the detailed noise properties of the data, which have shown several changes over the history of C-BASS observations. A software algorithm for optimally tuning the parameters was written and is used to calculate these parameters prior to running the pipeline. See Section 6.4.3 for a detailed discussion.

Once the flagging has been applied, a plot is produced showing which data have been deemed contaminated. This plot is reproduced in Figure 4.6. As is clear, no data were deemed to be contaminated by RFI in the example observations.

#### 4.2.6 Astronomical Calibration

A number of calibration functions are applied, which make use of observations of calibration sources. Sources observed for this purpose are listed in Table 4.2. Calibration observations are not part of the regular survey schedules, but are executed via stand-alone calibration schedules which precede and follow the regular survey schedules. In regular operation, the pipeline’s calibration functions will first be run on the calibration schedules to calculate the calibration values, and then they are run on the survey observations in order to apply the calibrations. The sample data set used in this discussion contains both calibration observations and survey observations, however, so a single application of the calibration functions will calculate and apply the corrections at once.

Table 4.2: List of C-BASS Calibration Sources

Source	$S_{\text{C Band}}$ (Jy)	Reference
Cyg A	367.9	(1)
Tau A	578.6	(2,3)
3C123	16.6	(4)
3C218	13.6	(4)
3C274	75.0	(5)
3C286	7.6	(4)
3C353	23.5	(4)
Cas A	662.0	(1,6)
J0225+621	38.6	(7)
J0927+390	6.9	(7)

- 1: Weiland et al. (2011)
- 2: Vinyaikin (2007)
- 3: Aller & Reynolds (1985)
- 4: Ott et al. (1994)
- 5: Baars et al. (1977)
- 6: Hafez et al. (2008)
- 7: C-BASS estimate

The first calibration step is to calculate pointing offsets in the data (implemented in the `pointing`




Figure 4.6: Plot showing the results of RFI flagging. Data flagged from previous steps are marked as “old flags,” while RFI-flagged data are marked “new flags.” In this case, no data were deemed to be compromised by RFI, and so only previous flags are shown. For clarity, the time series plot one point for every 10 seconds.

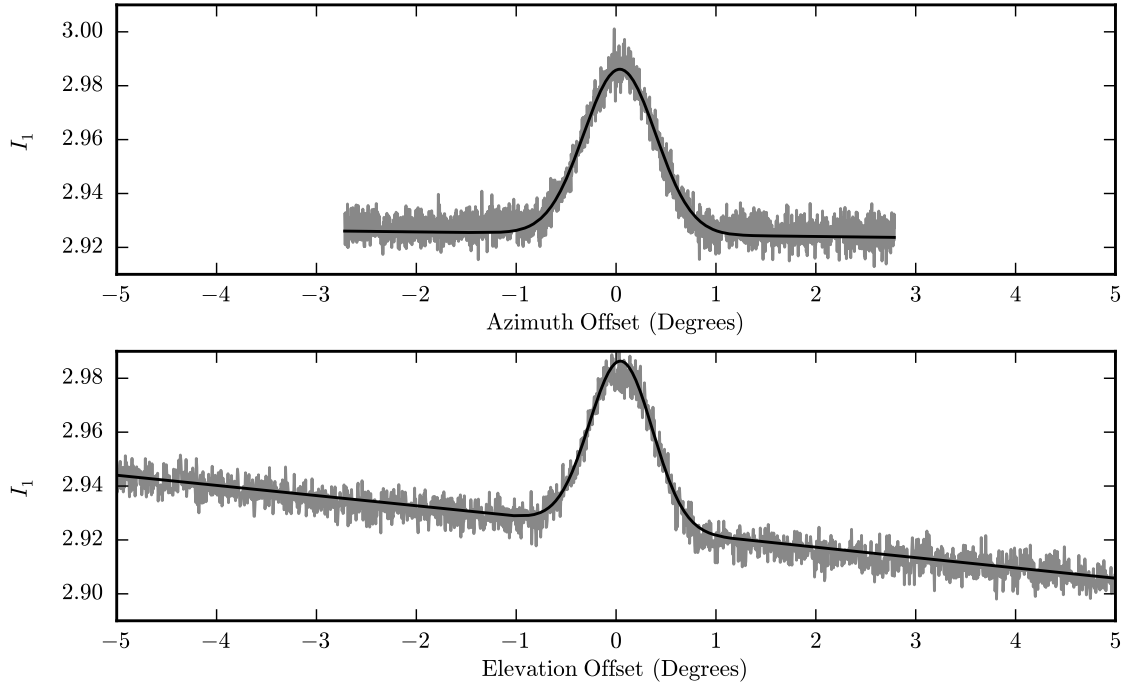


Figure 4.7: Pointing cross plots of calibrator source J1230+123 (aka 3C274). The grey curves are the astronomical data, while the black are the fitted Gaussian source and linear background model. The user may optionally view such plots for all astronomical calibration observations and flag any poor fits.

task). Radio pointing is an important question for the C-BASS experiment: although it is straightforward to develop an optical pointing model to arcsecond accuracy using an optical telescope attached to the antenna, leveraging this towards a precise radio pointing model is challenging. The large radio beam limits sensitivity to only the brightest of point sources while causing any pointing measurements of those sources to have intrinsically large uncertainties. For these reasons, radio pointing precision has been limited to arcminutes, with a typical RMS of  $3'$ .

Calibration observations include elevation and azimuth scans across the calibration source. The pipeline calculates pointing offsets from these scans by fitting a Gaussian function (with FWHM fixed to the known beam width) atop a linear baseline. For each calibration observation, the fits are calculated independently for channels  $I_1$  and  $I_2$ . The user has the option of viewing and approving these fits. A plot showing one such fit is reproduced in Figure 4.7. The calculated pointing offsets are written to a file and will be used when calculating the noise diode temperature.

It is next necessary to calculate the temperature of the noise diode from the calibration observations. This calculation compares the excess of observed temperature when observing the calibrator source to the excess when the noise diode is enabled. Corrections for the atmospheric opacity (calculated from commensurate sky dips) and pointing offsets (calculated previously) are applied. The




Figure 4.8: Noise diode temperatures as measured from astronomical calibrator source observations. The scatter is due to uncertainty in measuring astronomical source fluxes. The grey points have been automatically flagged as suspicious.

expected temperature of the calibration source is calculated from

$$T_C = \frac{\epsilon_a A_p}{2k} S_C, \quad (4.9)$$

in which  $S_C$  is the known flux of the calibration source,  $A_p$  is the geometric area of the aperture,  $\epsilon_a$  is the aperture efficiency of the antenna (theoretically calculated to be 0.55), and  $k$  is the Boltzmann constant. This step is performed in the `noise` task, and applied to each calibration source. A plot of derived noise diode temperatures is produced, reproduced in Figure 4.8. A plot is also generated showing the calculated atmospheric opacity during the observations in the `tau` step, which is reproduced in Figure 4.9.

An estimate for the total system temperature is calculated in the `tsys` task. This is a simple calculation, as it is merely the product of the  $T_{\text{sys}}/T_{\text{ND}}$  ratio from the  $\alpha$  correction multiplied by the recently calculated  $T_{\text{ND}}$ . The resulting plot is reproduced in Figure 4.10.

With the above calculations complete, the astronomical calibration can finally be applied to the data in the `astro` task. The noise diode temperature and atmospheric opacity are interpolated for the survey data of interest and plotted (reproduced in Figures 4.11 and 4.12). The flux scale is






Figure 4.9: Atmospheric opacity, as measured from sky dips and astronomical source observations. The scatter is due to uncertainty in measuring astronomical source flux densities. The grey points have been automatically flagged as suspicious.

converted to units of Kelvin via:

$$I [\text{K}] = I [\alpha] \times 2T_{\text{ND}} \quad (4.10)$$

in which  $[\alpha]$  denotes the data scaled by the  $\alpha$  correction, in which the noise diode events have an amplitude of 0.5. The resulting data are understood to contain several components:

$$I [\text{K}] = (T_{\text{sky}} + T_{\text{CMB}}) e^{-\tau} + T_{\text{atm}} (1 - e^{-\tau}) - T_{\text{Load}} \quad (4.11)$$

where  $\tau$  is the atmospheric opacity,  $T_{\text{atm}}$  is the ambient temperature of the atmosphere, and  $T_{\text{CMB}}$  is the cosmic microwave background temperature. The final step is to use weather station measurements of  $T_{\text{atm}}$  and the cold receiver's internal monitored measurements of  $T_{\text{Load}}$  to construct

$$I_{\text{cal}} = T_{\text{sky}} + T_{\text{CMB}} - T_{\text{Load}} \quad (4.12)$$

which is achieved via

$$I_{\text{cal}} = [I [\text{K}] - (T_{\text{atm}} + T_{\text{Load}}) (1 - e^{-\tau})] e^{\tau} \quad (4.13)$$

$I_{\text{cal}}$  is the final, calibrated time series. It is calculated independently for  $I_1$  and  $I_2$ . The results of




Figure 4.10: System temperature, taken from the  $\alpha$  database and converted to units of kelvins via the interpolated noise diode calibration. It is clear that the scatter in the astronomical calibration of the noise diode temperature has a dominating systematic effect on these values.

this are plotted, and this plot is reproduced in Figure 4.13. The results of the linear polarization calibration are also shown, although this step is still under development and the results are unreliable.

#### 4.2.7 Summary Plots

After the data have been calibrated, the pipeline can provide a number of diagnostic plots for users to judge the overall quality of the data. These include a plots of the data's RMS (in the `rms` task), the data's power spectra (in the `power` and `overf` tasks), and a final plot of data flagging (in the `summary` task).

The RMS plot is shown against data integration time, where individual, adjacent time series points have been integrated together over a range of periods. White noise would show an RMS that decreases with integration time, and so this plot can be used to see if the noise truly does integrate down. This plot is reproduced in Figure 4.14. In this case, the noise properties of the  $Q$  and  $U$  data behave well, while the  $I$  channels have significant, low frequency, correlated noise.

The first of the power spectra plots shows the power spectra for  $I_1$ ,  $I_2$ , and the averaged  $Q$  and  $U$  channels. This allows the user to look for unexpected spectra behaviour, such as slopes and narrow spectral features. See Figure 4.15 for a reproduction of this plot.




Figure 4.11: Linearly interpolated noise diode temperatures. It is clear that the scatter in the astronomical calibration has a dominating systematic effect on these interpolations.

The second of the power spectra plots shows power spectra for the four channels, as above, but in this case calculates and overplots the power spectra on a per-scan basis. In addition, the  $1/f$  noise slope  $\alpha$ , knee frequency  $f_{knee}$ , and white noise standard deviation  $\sigma$  are calculated. The knee frequency is that at which the power from  $1/f$  noise is equal to that from white noise. Figure 4.16 reproduces this plot (with the individual power spectra averaged, for clarity) and shows the fitted noise parameters.

Finally, the time series data are shown once more with flagged data points marked. The flags are colour-coded to show during which stage they were flagged. In addition, the user is presented with a final opportunity to manually flag for which suspicions exist. Because all of the flagging on the example data set was performed during the initial stage, this plot is actually identical to Figure 4.13 and is not reproduced a second time.

#### 4.2.8 FITS Output

Once the pipeline calculations have been completed, the data are written to disk using the `fits` task. The file format is a binary FITS table (Wells et al. 1981; Harten et al. 1988). The table columns and the precision of each are listed in Table 4.3. Based upon these precisions, the approximate file




Figure 4.12: Linearly interpolated atmospheric opacity values. The scatter of the individual values has a dominating systematic effect on these interpolations.

size is 47 MB per hour of data saved.

The user has the option of only saving data taken while observing certain sources. One may only desire the survey observations, for instance. This option allows one to reduce file sizes and avoid passing non-survey data to the map maker program. Advanced users may opt to define new selection functions here, although this does require modifying the pipeline source code.

The pipeline at this point can run the `Descart` map maker program on the data in order to produce a map of the newly-reduced data. `Descart` is discussed further in Section 4.3. After, a `python` script can be called in order to project the HEALPix map onto a convenient format.

This is the concluding step of the `MATLAB` pipeline. A number of external diagnostic scripts exist which collate and present these outputs in `HTML` format for convenient review, but these scripts are beyond the scope of this discussion.

### 4.3 The Descart Map Maker

The `Descart` map maker uses the pipeline-reduced, binary `FITS` data to calculate a map of the astronomical sky, while attempting to optimally reduce any residual correlated noise in the data. The output is in the HEALPix format, which is an equal-area pixelization for maps on spherical




Figure 4.13: Astronomically calibrated time series. Flagged data are marked. It is clear that the scatter in the astronomical calibration has a dominating systematic effect. For clarity, the time series show one point for every ten seconds.




Figure 4.14: RMS of the time series as a function of integration time. The thick, dashed lines show the expected  $N^{-1/2}$  behaviour expected for white noise. For each polarization, a variety of high-pass filters are applied. In the  $I$  channels, the improvement with filtering is due to the presence of low frequency, correlated noise. The filtered data show sharp drops at their cut-off frequencies, which is expected independent of noise properties.

surfaces (Górski et al. 2005).

Subtracting the reference load from the sky data reduces instrumental  $1/f$  noise, though imperfectly, but does nothing to reduce  $1/f$  noise from the atmosphere. Low-frequency, correlated noise therefore persists in the data after pipeline reduction. If left uncorrected, this noise will appear as stripes in the final maps and severely hinder their sensitivity.

The optimal separation of correlated noise from sky signal in scanning data can be performed through linear algebra (Smoot et al. 1992). This approach is computationally expensive, however, scaling as  $N^2$ , where  $N$  is the number of pixels in the map. Approximate algorithms scaling as  $N$  and providing near-optimal performance have been developed to allow for correlated noise subtraction for high spatial resolution experiments. One of these is **Descart**, which stands for Destriping Cartographer. **Descart** is written in the FORTRAN language and described in detail in Sutton et al. (2009).

The strategy employed in **Descart** is to divide the data into a number of short timeseries, alternately called “chunks” and “scans,” and to find a joint fit to constant offsets for each scan and a spatially pixelized map. Scans are required to be of contiguous data, and the user can specify




Figure 4.15: Power spectra of the data, calculated over the entire time series. The  $1/f$  noise is clearly present in all polarizations, but is much stronger in the  $I$  channels than in the linear polarization. Some narrow spectra features can also be seen. For clarity, not all points are shown (the shown points are logarithmically spaced in frequency).

how long each scan should be (with plausible values ranging from one second to several tens of seconds). The net effect of this fitting is to apply a high-pass filter to the data which selectively passes persistent sky structure.

The scans must be identified prior to running `Descart`. This is done by running the python script `split_data.py` on the C-BASS FITS files. A second binary table is then added to the FITS files, denoting the times of the individual scans.

Best results are achieved when each individual pixel has been crossed by scans at a variety of angles. This is achieved by combing many months of survey data, but also implies that maps made from a small number of schedules will have poor fits.

`Descart` takes as input a text file (usually with suffix `.ini`) which sets a large number of options. These options include the name of another text file, itself containing a list of C-BASS FITS files from which to make maps. Also included are the scan length, the convergence criterion of the fitting, and a variety of output options.

After completing the fitting, `Descart` produces a HEALPix file containing maps for the various polarizations, files listing the covariances of the maps and the number of data points per pixel of the maps (also known as a hit map), and (optionally) files containing the fitted offsets for each scan.




Figure 4.16: Power spectra of the data, averaged over the individual astronomical observations. The  $1/f$  noise is clearly present in all channels, as are the 1.2 Hz residuals. The fitted noise properties are overlaid for each polarization.

Assuming all has gone well, the HEALPix maps represent the final form of the reduced data.

The survey data run through the pipeline in the previous section were mapped in this way. A Mollweide projection of the resulting map is shown in Figure 4.17. As these data comprise just two observing schedules, the destriping algorithm behaved suboptimally. Regardless, the Galactic plane is clearly visible in the Stokes  $I$  map.



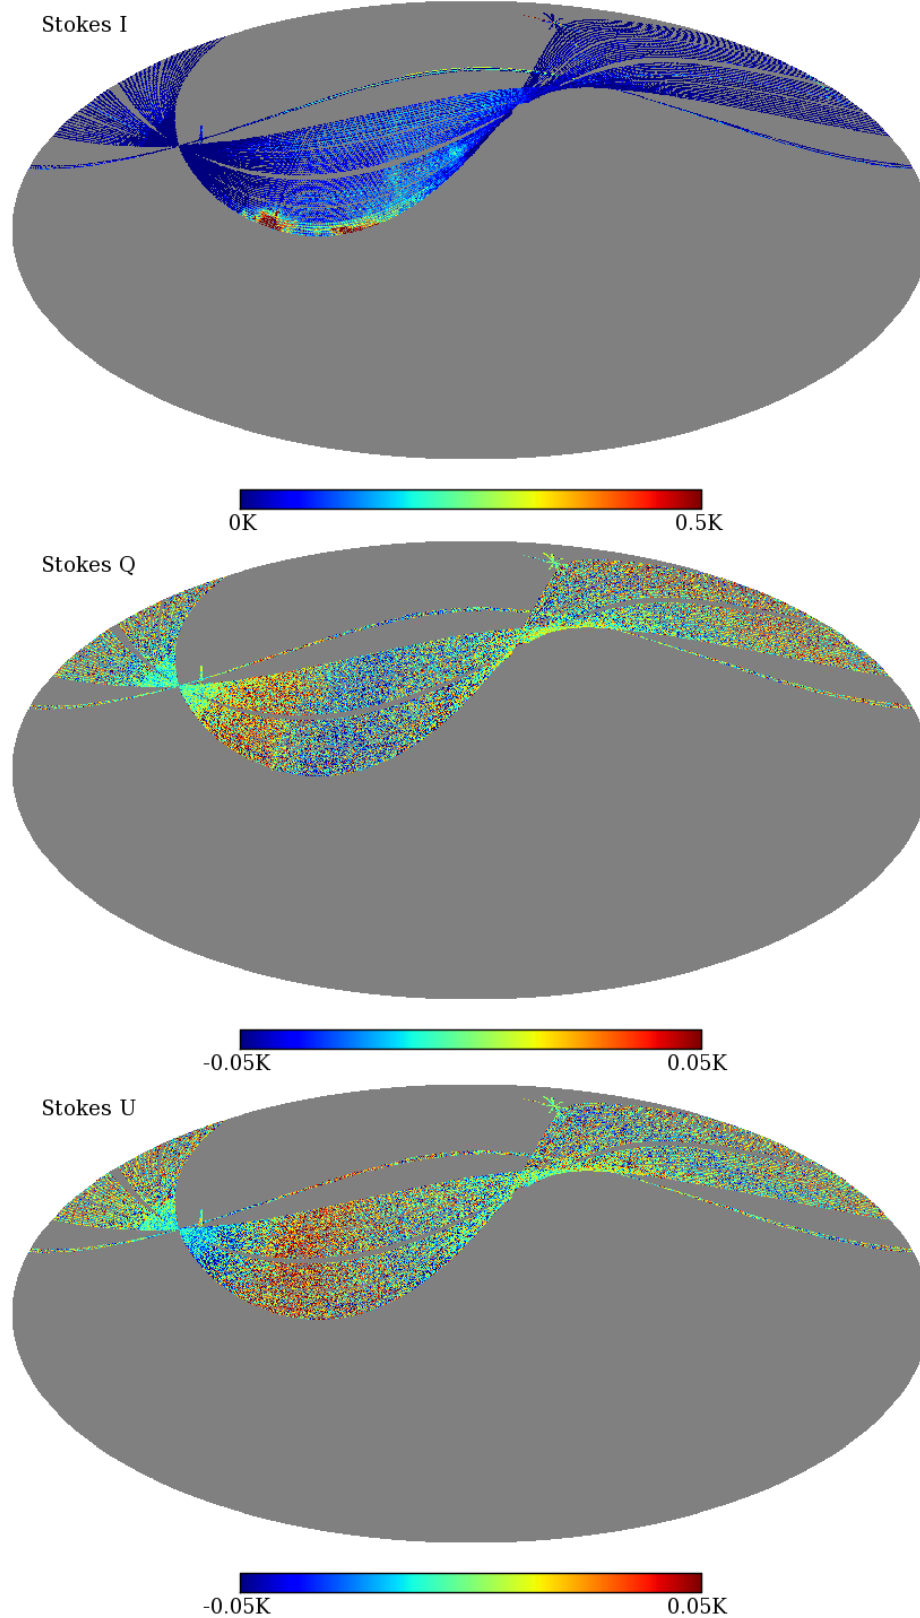


Figure 4.17: Mollweide projected maps of the calibrated and destriped data. Maps are shown in Galactic coordinates, with the Galactic Centre at the centre of the projection. Colourmaps reflect astronomical calibration.

Table 4.3: Columns in the Output FITS Table

Name	Description	Units	Precision
MJD	Modified Julian Date	Days	D <sup>1</sup>
RA	Right Ascension	Degrees	D
DEC	Declination	Degrees	D
AZ	Telescope Azimuth	Degrees	D
EL	Telescope Elevation	Degrees	D
LON	Galactic Longitude	Degrees	D
LAT	Galactic Latitude	Degrees	D
I	Stokes $(I_1 + I_2)/2$	K	D
Q1	Stokes $Q_1$	K	D
U1	Stokes $U_1$	K	D
Q2	Stokes $Q_2$	K	D
U2	Stokes $U_2$	K	D
Q3	Stokes $(Q_1 + Q_2)/2$	K	D
U3	Stokes $(U_1 + U_2)/2$	K	D
V	Stokes $(I_1 - I_2)/2$	K	D
FLAG	Binary Flag	Logical	D
DAYFLAG	Daytime Flag	Logical	D

1: Double float, 8 Bytes.

## Chapter 5

# The C-BASS Cold Cycle Correction

### 5.1 Context

The C-BASS receiver is cooled using a mechanical compressor, which operates at 1.2 Hz. This induces mechanical vibrations and temperature oscillations in the receiver at this frequency and its harmonics. The C-BASS integrated data show the results of these oscillations, with amplitudes comparable to or greater than the thermal noise. The removal of these oscillations from the data is necessary in order to achieve the greatest-possible sensitivity.

The precise mechanism by which these oscillations enter the data has not been determined. An early hypothesis was that the reference load temperatures were being modulated by the cold cycle, and that this resulted in a periodic signal in the detected power. However, upon isolating the reference loads from these temperature oscillations, no improvement was seen. The current hypothesis is that the effect is microphonic: mechanical vibrations in the components, connectors, and cables are causing small changes in their radiofrequency transmission characteristics. This hypothesis has been supported by trials in which the compressor was shut off and the cryostat struck with a mallet. Variations in the output power of similar amplitude were seen during these trials.

Astronomical signal in the integrated C-BASS time series occurs at frequencies below 40 Hz, with most of the astronomical signal at the lowest frequencies. The compressor frequency is coincident with these, so any attempts at notch-filtering or bandpass filtering the signal will cause loss of astronomical signal. This is particularly insidious, as the experiment's scanning strategy will then map such loss onto spurious spatial structure on the sky.

If these power oscillations were perfectly symmetric between the sky and reference load channels, then the demodulation and filtering the digital backend would remove them. They are not symmetric, however, and so the digital backend does not filter them.

An additional complication is that the 1.2 Hz frequency is only approximate. The compressor is phase-locked to the 60 Hz power grid and cycles once for every 50 power grid cycles. The cold cycle frequency is thus subject to frequency and phase fluctuations present in the power grid. Further, due to the uncertainty in the mechanisms involved, amplitude variations in the interfering oscillations should also be considered.

It is therefore necessary to find a method for removing the 1.2 Hz signal (and its harmonics) from the data in such a way that the astronomical signal is unaffected, that accounts for phase and frequency noise on the power grid, and that is capable of handling variations in the 1.2 Hz amplitude. The following sections describe the algorithm created for this purpose.

## 5.2 Assumptions

It is assumed that the oscillations in the data are frequency and phase locked to the compressor cycle. Although the precise cause of the oscillations remains uncertain, this assumption is consistent with the observed behaviour and is borne out by the success of the resulting algorithm.

A simple sinusoidal model for the temperature cycle in the cryostat is assumed:

$$\Delta T(t) = A(t) \sin [2\pi B(t) \cdot t + C(t)] \quad (5.1)$$

where  $\Delta T$  is temperature offset of the cold receiver,  $t$  is time, and  $A$ ,  $B$ , and  $C$  are time-varying parameters to be determined. The phase of the variations is of greatest concern:

$$\theta(t) = 2\pi B(t) \cdot t + C(t) \quad (5.2)$$

It is explicitly assumed that phase of the interfering signal is locked to this.

The three parameters are measured from a 5 Hz temperature sensor on the reference load bobbin (prior to July 30, 2012) or from a 5 Hz temperature sensor on the cold plate (after July 30, 2012). The sampling frequencies of these variables are sufficient to capture the 1.2 Hz frequency, but can still limit the phase accuracy. The sampling frequencies are, however, limited by the monitoring hardware and are not trivially increased.

It is assumed that the  $A$ ,  $B$ , and  $C$  parameters do not vary on timescales faster than several minutes. Although the amplitude of the interference is assumed to be proportional to  $A$ , an additional, constant, scaling factor is also fitted prior to correction.

The interfering signal is eventually subtracted from the astronomical time series via a fitted template. The template is represented by a periodic function  $F(\theta)$ , which is then used to estimate

the interfering signal using the measured phase  $\theta(t)$ :

$$\Delta P(t) = P_0 F[\theta(t)], \quad (5.3)$$

where  $\Delta P(t)$  is the interference estimate to be subtracted and  $P_0$  is the amplitude (fitted using linear regression) which minimizes residual interference. An accurate and precise measurement of the template  $F(\theta)$  is essential, and this can only be achieved under the assumption that the interfering signal is consistent over a sufficiently long timescale. It is possible that the signal is multiplicative rather than additive. If this is true, then the template acts as a 1st order correction.

### 5.3 Parameter Measurement

The temperature monitoring vectors do not show a pure 1.2 Hz signal. They include a DC offset, low frequency noise, and a number of 1.2 Hz sidebands and aliased harmonics. The signal is bandpass-filtered to remove these contaminants, passing frequencies from 1.1 to 1.3 Hz and stopping frequencies below 1.0 and above 1.4 Hz. The model of equation 5.1 is fitted to five minutes (1500 samples) of the filtered temperature series. This is repeated, offsetting by 30 s (150 samples), until the entire time series has been fitted. Note that each individual, five minute fit will overlap with many others.

The resulting  $A$ ,  $B$ , and  $C$  time series have 120 samples per hour but are auto-correlated on a five minute timescale. One should ask whether this auto-correlation timescale under-samples variations in these parameters. This has been tested by shortening the fitting timescale, and no changes were seen. However, shortening the fitting length did increase the noise, resulting in a less precise measurement. Thus the five minute timescale is justified.

Next, the parameter time series are interpolated over the full 100 Hz sampling of the astronomical data and smoothed with a 15 minute kernel. This smoothing serves two purposes: eliminating sharp changes in the parameters and lessening the impacts of isolated, poor fits.

The results of this fitting are shown in Figure 5.1. The sufficiency of the fitting timescale is clear. The variations in  $B$  and  $C$  are particularly smooth. These parameters are degenerate, and their variabilities are clearly correlated. They are consistent with actual variations in the power grid frequency and phase. Although measuring such frequency and phase variations in the power grid is not a challenging task, it is rather pleasing that C-BASS can measure them synchronous to its own data acquisition.

Of course, what is germane to this adventure is the phase  $\theta$ , given in equation 5.2. This is calculated directly from  $B$  and  $C$  and is plotted in Figure 5.2. Clearly, assuming a pure 1.2 Hz frequency tone for this with no phase noise would give a resoundingly poor fit. A more robust fitting is needed, and the procedure outlined here performs well.




Figure 5.1: Cold cycle model parameters fitted to an hour of data from January 17, 2012. Black points are from individual, five-minute fits and the grey curves are smoothed interpolations. The top panel shows the amplitude parameter  $A$ , which varies at the percent level. The middle panel shows the frequency  $B$ , and the bottom the phase parameter  $C$ . The correlation between the two lower parameters is clear.




Figure 5.2: Deviation in phase of the fitted cold cycle model to that of a perfect, 1.2 Hz signal, as measured from an hour of data from January 17, 2013. If the cold cycle were a perfect 1.2 Hz signal with no phase noise, then this value would be a constant zero. The gentle slope for most of the hour and sudden structure at the end reflect true phase and frequency noise on the power grid.

## 5.4 Template Estimation

Once the 1.2 Hz phase has been measured, it is possible to estimate the interference template. While the phase  $\theta(t)$  is measured from the temperature monitoring of the cryostat, the templates  $F(\theta)$  are measured directly from the Stokes  $I$ ,  $Q$ , and  $U$  time series. The calculation of  $F(\theta)$  must be performed in advance of data reduction, as it relies upon estimating the template from many days worth of data. The C-BASS observer is required to periodically run the `MATLAB` script `updateLoadTemplates`, which then writes the calculated templates to a template database.

The 1.2 Hz templates are calculated as the mean of a number of template estimates  $F_i(\theta)$ :

$$F(\theta) = \langle F_i(\theta) \rangle \quad (5.4)$$

Each template estimate is calculated independently from 2 hours of data. The template update script calculates these template estimates and writes them to the C-BASS data archive. The C-BASS observer is then required to view the template estimates and determine if and when a qualitative change in the template estimates has occurred. Each contiguous period of template estimate stability is averaged to give the templates for that time period.




Figure 5.3: Two hours of Stokes  $I_1$  data from January 17, 2012, folded over the measured cold cycle phase. The grey dots are the individual data points, mean-subtracted and normalized by standard deviation. The black line is the mean of these over 1024 phase bins. This represents an estimate of the interference template  $F_i(\theta)$  for the  $I_1$  channel.

For each 2-hour period, the  $A$ ,  $B$ , and  $C$  parameters are calculated as in Section 5.3. The pipeline’s  $\alpha$  correction is applied in the filtered, classic, and polonly cases, giving 20 unique data vectors (8 for filtered data, 8 for classic-mode data, and an additional 4 for polonly data). The noise diode amplitudes are subtracted from the noise diode events to give continuous time series. A 1.1 Hz high-pass filter is applied to each channel, removing a large fraction of the astronomical signal and  $1/f$  noise. At this point, a simple flagging stage is performed, rejecting portions of the time series with anomalously high standard deviations. This flagging is intended to remove RFI, poorly filtered data, and any anomalous discontinuities that may be present in the data.

Next, each channel is folded over the cold cycle phase  $\theta$ , as measured above, with mean subtracted and normalized by standard deviation. The template estimates are then reached after averaging in phase bins from 0 to  $2\pi$ . An illustration of this step is shown in Figure 5.3. As alluded to above, this is performed for all 20 unique outputs of the  $\alpha$  correction’s various modes.

These template estimates are calculated for every non-overlapping, two-hour period for which data exist. The estimates are most easily visualized in bit-map images, where each column represents an individual estimate and phase is constant along rows. Four such images are shown in Figure 5.4 for the backend-filtered, Stokes  $I_1$ ,  $I_2$ ,  $Q$ , and  $U$  channels. As is clear, the template estimates were



quite stable through most of early 2012, but then became less stable due to failures of two of the low noise amplifiers.

Every time the C-BASS observer updates the template estimates, the observer is requested to judge where qualitative changes in the estimates' behaviour lay. The observer can add new boundaries, delete erroneous ones, and tweak those in need of modification. The `MATLAB` script takes note of this input and, once the observer has finished, re-calculates the templates for any modified time periods.

The actual templates are calculated as the weighted averages of the template estimates. The estimates are weighted by the number of individual data samples averaged for each phase bin. Any estimates which disagree significantly with the overall averages are discarded. Finally, the templates are repeated for three periods and low-pass filtered with a 60 Hz cutoff. This final filter ensures that there is no discontinuity at a phase of 0 and that no hints of 60 Hz pickup from the power grid itself will pollute the template, while the padding prior to filtering guarantees that transient filter artefacts will not be a problem. Figure 5.5 shows some sample templates as calculated with this method. The templates are then written to the template database for use with the pipeline.

## 5.5 Template Fitting

Within the C-BASS pipeline, the cold cycle correction is invoked in the `load` routine, as described in Section 4.2.3. This begins in a similar fashion to the template calculation of the previous section, with the cold cycle parameters  $A$ ,  $B$ , and  $C$  being calculated from the cryostat temperature monitoring. The relevant templates are read from the template database and interpolated over the full time range of the data. The amplitude is scaled by  $A(t)$  and then scaled again by a constant factor fitted to the astronomical data time series. The scaled template is then subtracted.

Figure 5.6 demonstrates the final steps of this procedure. The uncorrected time series, the fitted template, and the corrected time series are shown for comparison. Figure 5.7 shows the power spectra of the data before and after correction, showing the template's success at reducing not just the primary frequency, but also the harmonics.

## 5.6 Discussion

The success of this method depends on the quality of the underlying assumptions. The phase stability of the interference with the cold cycle, the stability of the templates, and the stability of the interference amplitude are particularly critical. None of the assumptions are perfect, particularly the latter two.

Looking at Figure 5.4 it is clear that the template estimates have, at times, been quite unstable.




Figure 5.4: Template estimates for the Stokes  $I_1$ ,  $I_2$ ,  $Q$ , and  $U$  channels. Greyscale shows amplitude (lighter represents higher values), with each column of each polarization normalized to an RMS of 1. These were estimated every two hours for six months from late 2011 through early 2012. Time increases along the x axes, while phase is along the y axes. The stability of the estimates for the first two months of 2012 is remarkable. The instabilities that followed were symptoms of unhealthy low-noise amplifiers in the cold receiver.




Figure 5.5: Templates for the Stokes  $I_1$ ,  $I_2$ ,  $Q$ , and  $U$  channels for the middle of January, 2012. These were calculated by averaging nearly two weeks of template estimates.

It has also been observed during times of poor receiver health that the interference amplitude can change over minute timescales. Neither of these are expected behaviour when the receiver is in good health, and so the presence of these effects can then be used as diagnostics of receiver health.

When the assumptions fail, the correction performs poorly. This causes increased noise in the reduced data and maps, as the interference cannot be filtered without introducing artefacts, as described in Section 5.1. This shows the need for both this algorithm and for maintaining the good health of the receiver.




Figure 5.6: Demonstration of the template fitting and removal on 10 seconds of data from January 17, 2012. The top and bottom panels show the  $\alpha$ -corrected, Stokes  $I_1$  channel before and after the template subtraction. The centre panel shows the template that was fitted and subtracted. The high qualities of the fit and correction are clear to the eye.




Figure 5.7: Demonstration of the template fitting and removal on 30 minutes of data from January 17, 2012. The power spectra for the  $\alpha$ -corrected Stokes  $I_1$  channel are shown before (grey) and after (black) the cold cycle removal. The drop in interfering power is dramatic, not only at 1.2 Hz, but even up to the 5th harmonic.

## Chapter 6

# The C-BASS RFI Rejection Strategy

### 6.1 Context

The contamination of radio astronomical data by man-made radiofrequency interference (RFI) is a widespread problem. A variety of mitigation techniques have been suggested and implemented on different experiments (e.g., Fridman & Baan 2001; Nita & Gary 2010), typically relying on high spectral and temporal resolution. That RFI is also a problem for C-BASS is no surprise, but the lack of spectral resolution and the low time resolution of this experiment have made such widely-used techniques inapplicable.

International and federal agencies allocate and define the use of radiofrequency spectrum, whether for private, commercial, government, or scientific use (Panel on Frequency Allocations and Spectrum Protection for Scientific Uses 2007). A 20 MHz band of spectrum is reserved for radio astronomy at 5 GHz, but this is much smaller than the 1 GHz bandwidth of the C-BASS experiment. As such, the C-BASS band overlaps a range of man-made transmitters, including fixed transmitters on the ground, aircraft communications and radar, wireless internet communications, and satellite transmitters in orbit (geosynchronous and otherwise) about the Earth.

Mitigation of the resulting interference is best achieved through a variety of complementary methods. The C-BASS strategy is therefore a multi-pronged approach, beginning with passive, analog filtering for the most active frequencies, making use of spatial filtering for fixed transmitters, and flagging transient sources in the time series via statistical tests, with the latter two steps being performed in the `rfi` task of the C-BASS pipeline. These various approaches are discussed in this chapter. The statistical flagging is emphasized, as it formed a substantial portion of the work of this thesis.

## 6.2 Analog Filtering

When the Northern C-BASS instrument first began to observe, the contamination by RFI was clearly apparent. In particular, by scanning the telescope across all azimuths and elevations, it was possible to construct an RFI map of the sky. Many fixed transmitters were present, including an arc of geosynchronous satellites and several ground-based transmitters located in nearby towns. This is shown in the top panel of Figure 6.1. As these transmitters are fixed in position, any observations taken in these directions would always been contaminated. This would make certain parts of the sky unobservable, and so analog filtering of these signals was needed.

A spectrum analyzer was used with a hand-held feedhorn to measure the frequencies of these transmitters. It was found that the geosynchronous satellites were transmitting outside of the C-BASS nominal band, below 4.2 GHz, while the ground-based sources were transmitting in-band at 4.9 and 5.2 GHz. Notch filters were designed and built in Oxford for the in-band RFI, while an additional stage of bandpass filters were installed in order to reject the out-of-band interference. The use of additional filters increased the loss of the analog receiver, thereby increasing the instrument's system temperature, but did succeed in eliminating these sources of RFI. This is demonstrated in the lower panel of Figure 6.1.

## 6.3 Spatial Flagging

Prior to the installation of the additional analog filters, it was necessary to flag data based upon position on the sky. This flagging was applied purely as azimuth and elevation cuts in the data. The positions of the transmitters, whether satellite or ground-based, were identified. Any data points taken within a certain angle of these positions (with the angles depending on the strengths of the individual transmitters) were automatically flagged.

After the filters were installed, the majority of fixed transmitters ceased to be a problem. A single transmitter remains, but it is not seen at the typical C-BASS observing elevations of  $37^\circ$  and above. So, although this flagging step remains in the pipeline, it no longer flags any data during regular observations.

## 6.4 Statistical Flagging

Transient sources of RFI are not so easily flagged. These move about the sky and transmit at a variety of frequencies. Analog filtering all of these sources would reduce our bandwidth by too much, while flagging by position is clearly impossible. Time series flagging is the only viable approach, but owing to the lack of spectral resolution and low temporal resolution, it is very difficult to distinguish between RFI and astronomical signal. An algorithm is required that is automated and quantifiable,




Figure 6.1: All-sky maps, plotted in azimuth-elevation co-ordinates, taken over 12 hours each. Greyscale shows received power (lighter is brighter). Upper map is from September 29, 2010, before the installation of analog filters; lower is from November 18, 2011, after the installation of filters. RFI from geosynchronous satellites appears in the upper map as a dark arc peaking in elevation at  $180^\circ$ . The satellite RFI shows up as negative because the emission is outside of the band of the polarimeter's  $180^\circ$  hybrids: the hybrids' out-of-band phase response places most of this RFI in the cold load channels, rather than the sky channels. RFI from ground transmitters appear at azimuths of  $100^\circ$ ,  $170^\circ$ , and  $320^\circ$ . Ground pickup from mountains and Galactic plane emission appear in both maps. The upper map also shows the Moon at azimuth  $160^\circ$  and elevation  $75^\circ$ , elongated due to the elapsed time during mapping.






Figure 6.2: Example RFI event from 18:58:00 UTC on January 17, 2012. Data have been  $\alpha$  and cold cycle corrected. A noise diode event appears at 20 s. The strong polarization of the RFI event is clear.

that optimizes RFI rejection against the erroneous flagging of uncontaminated data, and that runs quickly.

Through an agonizing process of trial-and-error, an algorithm which satisfies these requirements has been developed. The method is a standard deviation rejection: the time series data are divided into short segments, the standard deviation of each is calculated, and the results are compared to those of other segments in the time series. This is done over a variety of timescales and across polarizations. Time periods with significantly outlying statistics are rejected.

An example of this RFI is shown in Figure 6.2. As can be seen, the RFI varies rapidly and has strong polarization. Astronomical signal, on the other hand, varies only with the scan rate and is significantly less polarized.

The following subsections describe the assumptions that go into this model, the procedure itself, and the method by which the flagging parameters are tuned.

### 6.4.1 Assumptions

Flagging RFI based upon time-series statistics requires pre-defining what RFI statistical behaviour is expected. This is not a trivial question; in this work, the assumed properties of RFI are based

upon a qualitative inspection of the C-BASS time-series data. This is necessarily subjective, but subjectivity cannot be avoided without independent verification of every RFI event. Subjectivity is therefore accepted, so long as the results are quantifiable.

The assumptions of this algorithm are as follows:

- RFI is assumed to be strongly polarized. Man-made transmitters are usually linearly polarized, but are also sometimes circularly polarized. The former case may be identified by seeing strong structure in the  $Q$  or  $U$  when compared to the  $I$  channels; the latter case manifests as discrepant response between the two  $I$  channels.
- RFI may be smoothly varying or highly structured. In the former case, it may even mimic a point source.
- Rapidly varying RFI will excite time series frequencies inaccessible to the C-BASS antenna and scan strategy.
- The non-interference noise in the data is assumed to be largely uncorrelated on short timescales.

A variety of RFI events are shown in Figure 6.3. Transient RFI events typically match the qualitative behaviour of these examples. A comparison to the assumptions listed above bears them out.

It is also necessary to specify what are the assumed properties of astronomical signal. Astronomical structure will only be able to contribute to time series frequencies below the ratio of scan rate to beam size. As regular observations occur at scan rates of a few degrees per second, and the beam size is a little less than a degree, astronomical structure will be confined to time series frequencies of less than a few Hz. The noise will have two parts: a  $1/f$  component that will also be below a few Hz and a white noise component that will affect all frequencies. Polarization of astronomical structure is predicted to be substantially less than that of RFI. A polarization fraction is not assumed; as will be seen in Section 6.4.3, the algorithm will seek the optimal value for this.

### 6.4.2 Flagging Procedure

Data are flagged for RFI through their standard deviations. In practice, this occurs in two steps: flagging on standard deviations of short time segments and flagging on standard deviations over longer time segments. These two steps complement each other. The short timescale is well suited to rapidly-varying RFI and is largely insensitive to sky structure. The long timescale is sensitive to slowly-varying RFI, but needs to be carefully tuned to avoid flagging astronomical structure. The timescales are illustrated in Figure 6.4.

The short-timescale flagging works by calculating the normalized, time-varying standard deviation  $\sigma(t)$  of the Stokes  $I_1$ ,  $I_2$ ,  $Q$ , and  $U$  polarization channels over 0.25 s and 0.5 s timescales. The  $Q$  and  $U$  channels are the means of the individually detected signals. If  $X$  represents polarization




Figure 6.3: Example RFI events from early 2012. The UTC dates and times are shown in the upper left. Data have been  $\alpha$  and cold cycle corrected. The upper two events appear strongly linearly polarized, while the lowest panel shows a brief event which is completely circularly polarized. The centre panel shows two noise diode events at  $-20$ s and  $40$ s.




Figure 6.4: Illustration of  $\sigma$  timescales against the RFI event of 22:45:00 UTC on January 17, 2012.

and  $T$  represents timescale, then eight  $\sigma_{X,T}(t)$  time series are calculated, which are then reduced to four by combining  $I_1$  with  $I_2$  and  $Q$  with  $U$ . These four  $\sigma_{X,T}(t)$  time series are then compared to each other in a fashion found to preferentially flag RFI over non-RFI data.

The calculations of the RMS time series  $\sigma_{X,T}(t)$  are now described. For a given polarization channel  $X$ , a preliminary  $\sigma_{X,T}^\circ(t)$  is calculated on a given timescale  $T$ . The running RMS of this time series is then calculated, denoted  $\sigma_{\sigma,X,T}(t)$ . Care is taken to avoid outlier points in  $\sigma_{X,T}^\circ(t)$  when calculating  $\sigma_{\sigma,X,T}(t)$ . The eight time series are then calculated via:

$$\sigma_{X,T}(t) = \frac{\sigma_{X,T}^\circ(t)}{\sigma_{\sigma,X,T}(t)} - 1 \quad (6.1)$$

These time series have an RMS equal to 1 and a mean of 0. This calculation is illustrated in Figure 6.5. Once calculated, the means of the total intensity time series are taken to give  $\sigma_{I,0.25}(t)$  and  $\sigma_{I,0.5}(t)$ , and the maxima of the linear polarization series are taken to give  $\sigma_{P,0.25}(t)$  and  $\sigma_{P,0.5}(t)$ . Thus there are four  $\sigma_{X,T}(t)$  time series for the short-timescale flagging. An example set of these is shown in Figure 6.6.

With the  $\sigma_{X,T}(t)$  time series calculated, two comparisons are made:

$$C_1(t) = [\sigma_{P,0.25} > S_1] \wedge [\sigma_{P,0.25} > S_2 \sigma_{I,0.25}] \wedge [\sigma_{P,0.5} > S_2 \sigma_{I,0.5}] \wedge [\sigma_{P,0.5} < S_3 \sigma_{P,0.25}] \quad (6.2)$$




Figure 6.5: Calculation of  $\sigma_{I_1, 0.25}(t)$  during the RFI event of 22:45:00 UTC on January 17, 2012. The upper panel shows the  $\alpha$  and cold cycle corrected data from channel  $I_1$ . The central panel shows the calculated  $\sigma_{I_1, 0.25}^\circ$  (grey) and the fitted running median (dashed). The bottom panel shows the final  $\sigma_{I_1, 0.25}$ , with RMS of 1 and mean of 0.

and

$$C_2(t) = \sigma_{P, 0.25} > S_4, \quad (6.3)$$

where  $S_1$ ,  $S_2$ ,  $S_3$ , and  $S_4$  are parameters. Whenever  $C_1(t)$  or  $C_2(t)$  are true, data are flagged.

These conditions require some explanation. The comparisons in  $C_1$  check that the standard deviation in linear polarization is significantly high at the 0.25 s timescale, that the standard deviation in linear polarization is significantly higher than that for total intensity for both the 0.25 s and 0.5 s timescales, and that the standard deviation for linear polarization is less deviant at 0.5 s than for 0.25 s. The net effect is that any flagged data will have stronger structure in linear polarization at short timescales than in total intensity, and that structure will get weaker at longer timescales. This last point helps to guard against astronomical signal, such as from polarized point sources. The comparison in  $C_2$  is much more simple: if the polarization structure at short timescales is much greater than ever expected from astronomical signal, then the data are flagged.

The long-timescale flagging works in a somewhat similar fashion, but using 2.5 s and 10 s timescales.




Figure 6.6: Calculation of the  $\sigma_{I_1, T}(t)$  during the RFI event of 22:45:00 UTC on January 17, 2012. The upper panel shows the  $\alpha$  and cold cycle corrected data from channel  $I_1$ . The lower panels show the calculated  $\sigma_{I_1, T}^\circ$  for the four timescales.

In this case, however, the  $\sigma_{I_1, 2.5}(t)$  time series is not used, but the  $\sigma_{I_1, 10}(t)$  and  $\sigma_{I_2, 10}(t)$  are. These latter two are compared, with significant discrepancy being flagged as circular polarization. Examples of these time series are shown in Figure 6.6.

The exact conditions for the long-timescale flagging are:

$$C_3(t) = [\sigma_{P, 2.5} > L_1] \wedge [\sigma_{P, 10} > L_2] \wedge [\sigma_{P, 10} > L_3 \sigma_{I_1, 10}] \quad (6.4)$$

and

$$C_4(t) = [\max(\sigma_{I_1, 10}, \sigma_{I_2, 10}) > L_4] \wedge [\max(\sigma_{I_1, 10}, \sigma_{I_2, 10}) > L_5 \min(\sigma_{I_1, 10}, \sigma_{I_2, 10})] \quad (6.5)$$

where  $L_1$ ,  $L_2$ ,  $L_3$ ,  $L_4$ , and  $L_5$  are parameters. Whenever  $C_3(t)$  or  $C_4(t)$  are true, data are flagged.

The first of these conditions flags data that show significant polarized structure at both the 2.5 s and 10 s timescales exceeding that expected from the total intensity structure. At these timescales,

astronomical signal is significant, so the comparison to total intensity is critical to avoid flagging of legitimate signal. The second condition assumes that structure in  $I_1$  and  $I_2$  are similar, as they should be for astronomical signal. Circularly polarized RFI will not satisfy this assumption and will therefore cause  $C_4(t)$  to flag.

After the above flagging has been completed, the algorithm proceeds to flag several seconds of data preceding and following the statistically flagged segment. This is done under the theory that an entire RFI event will not be flagged completely by this method and that some significant RFI may be buried in the noise surrounding the event.

Taken as a whole, the four comparisons cover the RFI assumptions listed in Section 6.4.1. They do, however, require no less than nine parameters. This is a large number of parameters, and it is unreasonable to expect good results if the user must set them manually. Rather, a facility was written for pre-calculating the parameters and saving them to a text file. Five parameter sets are calculated, ranging from extremely conservative (very few false-positive flags) to extremely aggressive (very little missed RFI). The user then simply selects a parameter set when running the RFI flagging.

### 6.4.3 Parameter Tuning

Identifying a set of useful statistical flagging parameters requires careful tuning. The multi-dimensional parameter space represented by the parameters is difficult to navigate manually, and so this tuning must be performed in a systematic fashion. Further, the tuning must be repeated periodically as changes to the receiver have caused changes in the noise levels of the various Stokes channels and changes in the receiver's polarization sensitivity.

A MATLAB program has been written to calculate useful parameter sets for the statistical flagging in advance of running the C-BASS pipeline. The strategy is to compare three independent flagging methods to identify true RFI events. The statistical flagging parameters are then chosen to maximize the flagging of true events while minimizing the flagging of non-RFI events. The three independent flagging methods are the statistical flagging above, a method based upon folding survey scans over azimuth, and an interactive, user identification. The program is called through the `tune_rfi` MATLAB script. There are five steps: data folding, parameter estimation, parameter testing, manual verification, and results. The first three steps are fully automatic. After running these, the program halts and waits for the user begin the manual verification stage. After this, the program halts again and waits for the user to request the results. The five steps are now described in detail.

Data folding refers to an RFI candidate identification scheme designed to work independently of the statistical scheme. This scheme uses regular survey scans and constructs an average power level versus azimuth. Prior to averaging, the data are high-pass filtered to remove DC offsets, correlated noise, and most astronomical signal. The averaged scans will contain residual features due to the remaining terrestrial pickup and astronomical signal. The method is illustrated in Figure 6.7.




Figure 6.7: RFI folding procedure illustrated on data from January 19, 2012. The upper panel shows the regular, all-sky scans (with the azimuth shown in grey and on the right axis). The data have been  $\alpha$  and cold cycle corrected, as well as high-pass filtered. The lower panel shows a close-up of a single azimuth scan (grey), as well as the mean intensity across azimuth and the  $3\sigma$  level above this. The RFI event at azimuth  $90^\circ$  is clear.

A similar average value is produced for noise-diode events, though in this case the telescope position is ignored and median values are subtracted in lieu of high-pass filtering. The noise diode events thus give a single value for comparison, rather than an averaged scan.

Any individual time samples inconsistent with the averaged scans or averaged value (as appropriate) are identified as candidate RFI. This comparison is performed on the total intensity and the linear polarization channels. This method works well for finding RFI during regular observations, but it is susceptible to falsely flagging astronomical point sources and performs poorly during skydips and calibration observations.

The  $\sigma_{X,T}(t)$  from the statistical flagging are calculated for the data that have not been marked as RFI candidates by the folding algorithm. These are then compiled into  $N$ -dimensional histograms corresponding to the four comparisons  $C_i$  described in section 6.4.2.

The  $N$ -dimensional  $C_i$  histograms of non-candidates are examined in order to suggest statistical flagging parameter sets. These suggestions are based upon a list of threshold false-positive rates. A



set of suggested statistical flagging parameters is suggested for each false-positive threshold.

Each false-positive threshold is applied separately to each of the four  $C_i$ , meaning that the true false-positive rate may be higher for a given suggested set. Furthermore, because the RFI candidates will contain some false-positive events, a further underestimate by the thresholds should be expected.

With the suggested parameter sets in hand, the full, raw time series are run through the statistical flagging to produce the various  $\sigma_{X,T}(t)$  time series for the time series. The  $\sigma_{X,T}(t)$  values are run through the various suggested parameter sets, producing an RFI flagging vector for each false-positive threshold. These flags comprise a set of RFI candidates to be compared to the flags from folding.

The statistical flagging is also produced for parameter sets in which the parameters have been varied by 10% from the suggested values. These additional flagging runs will give an estimate to the stability of the suggested parameter sets.

The statistical flags from the suggested parameter sets are then consolidated with the folded candidates to produce a master RFI candidates list. For each candidate, it is noted whether it was flagged via folding or statistical flagging (or both) and, if statistical, under which flagging parameter sets. After this step, the `MATLAB` program halts.

The user must now continue the `MATLAB` program in order to begin the manual verification. Each RFI candidate is plotted in turn, and the user is asked to judge whether or not it is truly RFI. Example manual verification plots are shown in Figures 6.8 and 6.9. The former plot shows an event judged to be true RFI, while the latter shows an event judged not to be RFI. The Sun and Moon, point sources, and the Galactic plane are denoted. A simulated scan over the 5 GHz sky model of de Oliveira-Costa et al. (2008) is also shown. Whether or not RFI is present is often obvious, but many candidates are difficult to judge or are affected by unusual circumstances. The user is able to decline classification in such cases.

This process will take a few hours per week of data to perform. For this reason, this process has not been performed on all data, but rather on a representative set. The product of this step is a list of RFI candidates, as identified by each parameter set, with a manual classification of each.

The candidates and classifications for each parameter set are compiled into a number of useful plots and statistics. For each parameter set, these include: a plot of each RFI candidate, the measured rates at which RFI events are caught, and the measured false-positive rates. From these results, one can view the overall performance of the parameter sets and make recommendations for parameter use in the pipeline. The parameter sets measured from a week of data beginning January 13, 2012 are shown in Table 6.1.




Figure 6.8: Example manual verification plot for an RFI candidate (which was judged to be RFI) at 02:15:28 UTC on January 13, 2012. The  $\alpha$  and cold cycle corrected  $I_1$ ,  $Q$ ,  $U$ , and  $I_2$  data are shown and labeled. They have been automatically scaled and offset for clarity. A simulated scan across the sky model of de Oliveira-Costa et al. (2008) is shown in grey, itself scaled for clarity. Hatched regions denote known astronomical objects. The RFI candidate is bordered by the heavy, vertical lines.

## 6.5 Discussion

When judging the success of the C-BASS RFI mitigation strategy, one must consider the expected deleterious effects of RFI on the final data product and the measurable benefits of mitigation. RFI of fixed position will introduce coherent features into the all-sky maps, while transient RFI will lead mainly to an increased noise level. It is reasonable to suspect that transient RFI could lead to biases in the maps, but this can be controlled for via null tests. Thus, as the analog filters and positional flagging have shown success in preventing coherent features, the main concern that remains is signal-to-noise ratio.

A detailed examination of the noise properties with and without statistical RFI flagging shows that, indeed, the noise properties of the final maps are improved (see Section 7.3). This justifies the assumptions underlying this algorithm and argues for its use in regular observations.

It must be noted, however, that the tuning of flagging parameters depends critically upon the noise properties of the data. If the noise has significant correlation or includes large, non-Gaussian




Figure 6.9: Example manual verification plot for an RFI candidate (which was judged to not be RFI) at 02:51:31 UTC on January 13, 2012. See description for Figure 6.9.

features, then the effectiveness of this flagging method will decrease. Further, if the noise properties change, then the tuning procedure must be repeated.

A possibility exists that the standard deviation diagnostics and comparisons used in the statistical flagging are not the optimal set. An exhaustive search for optimal statistical tests was beyond the scope of this work, as only an adequate set were required. It is therefore suggested that a search for an improved set of statistical tests may be of great use for future experiments facing similar RFI issues.

Table 6.1: RFI Flagging Parameters

Set	RFI Catch Rate	False Positive Rate	$S_1$	$S_2$	$S_3$	$S_4$	$L_1$	$L_2$	$L_3$	$L_4$	$L_5$
1	$67.2 \pm 3.4\%$	$0.23 \pm 0.08\%$	3.4	0.1	14.7	3.4	0.1	0.2	0.1	2.2	1.5
2	$54.5 \pm 3.0\%$	$0.04 \pm 0.06\%$	3.9	0.1	14.7	3.9	1.1	0.6	0.1	9.8	1.6
3	$46.3 \pm 1.9\%$	$0.002 \pm 0.0002\%$	6.5	0.1	14.7	6.5	1.6	7.4	0.1	7.8	6.6
4	$35.1 \pm 2.6\%$	$< 0.0001\%$	9.6	0.1	1.5	18.7	0.2	0.6	0.5	-1.0	-1.0
5	$21.6 \pm 4.5\%$	$< 0.0001\%$	7.8	0.5	1.2	45.7	1.4	0.1	0.5	-1.0	-1.0

## Chapter 7

# The North Celestial Pole with C-BASS

### 7.1 Context

The anomalous microwave emission (AME) was first discovered in observations near the North Celestial Pole (NCP) by Leitch et al. (1997). The discovery took the form of a puzzling spatial correlation between microwave observations and the 100  $\mu\text{m}$  map of the IRAS experiment (Beichman et al. 1988). The correlated emissivity at 14.5 GHz was greater than at 32 GHz, which was suggestive of free-free emission. No associated correlation was seen with  $\text{H}\alpha$  emission, however. Leitch et al. (1997) therefore suggested that this emission could be due to a new Galactic component of flat-spectrum synchrotron emission or free-free emission from exceptionally hot ( $> 10^6$  K) gas.

It was quickly realized that the microwave excesses seen by Kogut et al. (1996) and de Oliveira-Costa et al. (1997) may be similar examples of AME. A number of observational studies followed to seek a better understanding of the emission spectrum and survey the variety of astronomical environments and scales to see where it might occur. These showed that the emission generally turns over below 15 GHz and tends to trace dusty, highly radiated environments (Finkbeiner et al. 2004; Watson et al. 2005; Casassus et al. 2008). At the same time, a series of theoretical papers argued that the emission could be explained by a population of very small, very rapidly spinning dust grains (Draine & Lazarian 1998b; Ali-Haïmoud et al. 2009; Hoang et al. 2010; Silsbee et al. 2011; Hoang et al. 2011). In addition, Draine & Lazarian (1998a) argued that the free-free hypothesis of Leitch et al. (1997) was implausible due to the energy injection rate required to keep the gas so hot.

A notable mystery is whether the AME emission seen at large scales, apparently originating in diffuse gas, is due to the same mechanism as the AME seen in compact clouds. Although the spinning dust model would support a unified interpretation, the lower peak frequencies in diffuse regions have hampered observational support (Bennett et al. 2012). Davies et al. (2006) attempted to resolve this issue through a cross-correlation analysis of the WMAP data with a number of emission templates

across a variety of large-area, Galactic regions (including the NCP). While they found that the AME was indeed dust-correlated, it was impossible to rule out a flat spectrum origin. Ghosh et al. (2012) repeated this analysis over a wider region, finding a similar result, while Peel et al. (2012) used the 2.3 GHz Southern hemisphere map of Jonas et al. (1998) to show that, at least in the Southern hemisphere, the diffuse AME was inconsistent with flat-spectrum emission. Most recently, Planck Collaboration (2013b) again showed the presence of anomalous emission, consistent with spinning dust models but lacking the necessary low-frequency confirmation.

The NCP region remains exceptional. To date, no studies have ruled out flat-spectrum AME in this region. Further, Meyerdierks et al. (1991) showed that this region may be part of an unusual, cylindrical shock in the Galactic ISM, possibly due to the collision of a high velocity cloud with the Galactic plane. This raises the possibility that the energetics argument against free-free emission may not apply. If the region has received a transient injection of energy, then the average energy injection rate need not be high enough to sustain the emission.

Further study of the NCP region is therefore needed. A 5 GHz map of the region will allow for similar spatial correlation analysis. This will be able to show whether the AME in this region is consistent with a flat spectrum or can be explained by the spinning dust hypothesis. This is further convenient for the C-BASS experiment, as the NCP region is of particular importance for the Northern survey strategy. A dedicated, preliminary mapping of this region allows a better understanding of this important region and a useful test of the experiment’s data analysis pipeline. This chapter describes such an NCP map, its observation with the C-BASS antenna, and its processing with the C-BASS pipeline.

## 7.2 C-BASS Data

The C-BASS data were collected from late 2011 through early 2012, during which time the receiver’s  $I_1$  channel was working as expected. The  $I_2$ ,  $Q$ , and  $U$  channels were not producing good data during this time period.

Data were collected during the day, as the all-sky survey requires night time observations. Azimuth scans were performed, similar to the all-sky scans, but at  $1^\circ/\text{s}$  instead of the all-sky survey’s  $4^\circ/\text{s}$ . This slower scan speed was chosen to decrease the fraction of time spent between scans. In order to achieve even coverage over the field, scans were performed at a number of Nyquist-sampled elevations.

The total area covered is  $150^\circ$  in the form of a  $7^\circ$  radius circle centred at the NCP. This area completely covers the region observed by Leitch et al. (1997) and covers both hot and cold areas in infrared maps.

The observing strategy is to begin at an elevation offset  $7^\circ$  above the NCP, and scan in azimuth




Figure 7.1: Pointing centre of the antenna during a single NCP schedule observed on January 17, 2012. Co-ordinates are azimuth and elevation. The broad, grey line illustrates a  $7^\circ$  radius about the NCP. The track of the antenna's pointing centre shows the scans across the field and the skydips.

Table 7.1: List of NCP Observing Outages

Dates		Reason
2011-12-13	– 2011-12-24	Incorrect pointing
2011-12-27	– 2012-01-04	Power outage
2012-01-07	– 2012-01-11	Incorrect pointing
2012-01-21	– 2012-01-28	LNA bias issue
2012-02-02	– 2012-02-09	Movie filming
2012-02-16	– 2012-02-17	Compressor failure
2012-03-02	– 2012-03-08	Compressor failure

for 130 s. The noise diode is fired before and after each scan. Once done, the antenna takes an offset  $0.35^\circ$  below the first and repeats the azimuth scans. This continues until the antenna reaches an elevation offset  $7^\circ$  below the NCP. For each elevation offset, the extent of the azimuth scan is large enough to fully cross the circular field. Thus the scans at greater elevation offsets are of smaller extent. A sky dip is performed before, after, and halfway through this process. The noise diode is fired every two minutes. The observing schedule takes approximately 2.5 hours to run, with a total integration time within the observing area of 60 minutes. Observing overheads are dominated by time spent decelerating and accelerating the antenna between scans. The schedule was run three times per day, with astronomical calibration observations before the first schedule, between the second and third, and after the third schedule.

This scanning strategy is illustrated in Figure 7.1, in which the antenna’s azimuth and elevation positions are plotted for an entire observing schedule from January 17, 2012. The same scans are shown in celestial coordinates in Figure 7.2; the distortion due to the Earth’s rotation is clear.

Data were collected for six months, beginning in November, 2011. Issues with antenna pointing lowered the quality of data from November and early December, 2011, while issues with amplifier performance significantly degraded system performance after the middle of March, 2012. There are, therefore, three months of science-quality data, though these months were interspersed with occasional telescope outages due to commissioning tests, compressor failures, and site-wide power failures. These outages are listed in Table 7.1. During these three months, the NCP schedule was run 130 times.

During the times that these data were collected, only channel  $I_1$  was producing science quality data. The other total intensity channel  $I_2$  was adversely affected by poor amplifier performance and a system temperature significantly higher than that of the  $I_1$  channel. The linear polarization  $Q$  and  $U$  channels had extremely low sensitivities during these observations due to a phase imbalance in the polarimeter. Thus, only the  $I_1$  data are considered for this analysis.

These data were reduced using the C-BASS pipeline. As the full calibration strategy has yet to






Figure 7.2: Pointing centre of the antenna during a single NCP schedule observed on January 17, 2012, in celestial coordinates. The broad, grey line illustrates a  $7^\circ$  radius about the NCP. Curvature of the scans is because constant-elevation scans are not along great circles. Compression and rarefaction of the scans at 12 h and 0 h, respectively, are due to the Earth's rotation during the observation.




Figure 7.3:  $\alpha$ , gain, and  $T_{\text{sys}}$  values for the Stokes  $I_1$  channel, as taken from the  $\alpha$  database. The observing outages are seen as blank segments. The values are remarkably stable over this time.

be finalized, a modified version was used. The reduction and analysis of these data are described in Section 7.3.

### 7.3 C-BASS Analysis

The C-BASS data were reduced using the experiment’s MATLAB pipeline for the  $\alpha$ , cold cycle, and RFI steps (see Chapter 4). The  $\alpha$ , gain, and  $T_{\text{sys}}$  values, as calculated for the  $\alpha$  database are plotted in Figure 7.3. RFI flagging was performed using a variety of flagging levels (see Section 6.4.3). Backend-filtered data were used, so no  $r$ -factor correction was applied. The full time series were pared down to those times of constant elevation and constant scan speed, so to avoid issues arising from the turnaround between scans. The data were calibrated using a custom approach, as described below. `Descart` was then used to remove  $1/f$  noise and produce the astronomical map (see Section 4.3).

The data were calibrated outside of the pipeline. The astronomical calibration observations of the bright radio sources Cygnus A, Taurus A, and Cassiopeia A are plotted in Figure 7.4. Not




Figure 7.4: Astronomically calibrated values for the noise diode temperature  $T_{\text{ND}}$ , as measured from the sources Cygnus A, Taurus A, and Cassiopeia A. The values are remarkably stable, except for two step changes. The averages of the  $T_{\text{ND}}$  values for the three stable periods are shown as thick, grey lines.

shown are observations flagged for lying outside of  $3\sigma$  of the mean for each source. Observations of other calibrations sources are not included, as those sources are much fainter and showed much lower signal-to-noise ratios. Two points are evident. First, the calibrations on Taurus A and Cassiopeia A are in good agreement with each other, but are in some tension with Cygnus A. This is likely due to an error in the assumed flux density of the latter source. Second, the calibration values have several step changes. The first, in early February, is likely due to phase testing in the receiver, while the second occurred after a compressor failure and likely represents a change in amplifier gain.

The flux density of a calibration source is measured as excess over the received power when the antenna is pointed off-source. It is possible that astronomical signal will be present in the off-source measurement, causing an error in the resulting calibration. The practice of measuring the source flux density from a cross-scan on the source is meant to mitigate this issue, however it remains possible that the presence of nearby astronomical signal will contaminate a calibration observation. This can be tested by searching for trends in the calibration values against the parallactic angle of the source when it was observed, as is done in Figure 7.5. The lack of clear trends indicates that systematic effects arising from the parallactic angle of observation do not dominate the scatter in the measurements. The calibrations are therefore taken as the average  $T_{\text{ND}}$  values for the three time




Figure 7.5: Astronomically calibrated values for the noise diode temperature  $T_{\text{ND}}$  plotted against the parallactic angles of the calibration observations. The three panels represent the three stable periods of Figure 7.4. The averaged  $T_{\text{ND}}$  values are shown by the thick, grey lines. The lack of clear structure indicates that effects stemming from parallactic angle do not contribute significantly to the  $T_{\text{ND}}$  scatter.

periods. The averaged values are listed in Table 7.2. The calibration factor applied to the NCP data was the mean of those derived from all three sources.

The noise of the data was examined through several tests. First, the time series data had `Descart`'s fitted  $1/f$  and sky signal subtracted, leaving `Descart`'s estimates for the noise time series. These time series were plotted as histograms and power spectra in order to illustrate the noise properties. A sample histogram from a schedule run on January 17, 2012 is shown in Figure 7.6. A Gaussian function, with variance equal to that of the histogram, is plotted for comparison. The power spectrum for the same data is shown in Figure 7.7. The noise is white over most of the bandwidth, though residual 1.2 Hz interference is present. The decrease in power above 40 Hz is due to the filtering in the digital backend. The complete lack of aliased 60 Hz is striking. From these figures, the noise matches a Gaussian histogram and is white over most of the time series bandwidth, consistent with expectations for thermal noise.

Fitting Gaussian functions to the noise histograms allows comparison between the noise proper-

Table 7.2: NCP Calibration Measurements

Source	2011-12-5 – 2012-2-3			2012-2-9 – 2012-2-16			2012-2-18 – 2012-3-12		
	$N$	$T_{\text{ND}}$ (K)	$\sigma_T$ (K)	$N$	$T_{\text{ND}}$ (K)	$\sigma_T$ (K)	$N$	$T_{\text{ND}}$ (K)	$\sigma_T$ (K)
Cyg A	113	3.402	0.035	24	3.344	0.033	64	3.478	0.045
Tau A	91	3.336	0.018	19	3.285	0.021	39	3.382	0.022
Cas A	100	3.353	0.024	28	3.283	0.021	64	3.383	0.018
Mean	304	3.366	0.039	71	3.304	0.038	167	3.419	0.056

ties during the observing schedules. Figure 7.8 shows the fitted variances and averages of the noise for all NCP schedules that were run during this time. The data have been reduced with level 2 RFI flagging and `Descart` offset lengths of 100 samples. As with the  $T_{\text{ND}}$  calibration, the noise variances undergo a number of steps. When steady, however, the scatter is at the 2% level with only occasional outliers. The averages of the noise are nicely centred at zero, as expected.

A further test of the noise is performed by filtering the reduced time series with the kernel  $[-0.5, 1.0, -0.5]$ . This is a simple high-pass filter which strongly suppresses astronomical signal and




Figure 7.6: Histogram of the offset- and sky-subtracted time series noise from a schedule run on January 17, 2012. The histogram is shown in black; a Gaussian function with variance equal to that of the histogram is plotted in grey.




Figure 7.7: Power spectra of the time series noise from a schedule run on January 17, 2012. The spectra are plotted before (grey) and after (black) the subtraction of `Descart`-estimated offsets and sky signal. The noise is white over most of the bandwidth.  $1/f$  noise appears below the 1.2 Hz residual feature. The drop in power at high frequencies is due to filtering in the digital backend.

$1/f$  noise, but introduces correlation into the time series. Applying this filter and constructing noise histograms provides a much cleaner picture of noise variance, as can be seen in Figure 7.9. The variances are higher than in Figure 7.8, which is expected due to the filtering, but the scatter in the variances is far less. Also shown in this figure are the system temperature values from the  $\alpha$  database with astronomical calibrations applied. The clear correlation between noise variance and system temperature is an expected property of thermal noise.

The RMS noise level is related to the system temperature by the radiometer equation:

$$\sigma = \frac{T_{\text{sys}}}{\sqrt{B\tau}} \quad (7.1)$$

in which  $B$  is the receiver bandwidth and  $\tau$  is the integration time. The filtering in the digital backend causes an autocorrelation in the time series (see Section 3.4.6.3), and so a multiplicative correction of 0.96 is required when applying the radiometer equation. Based on the average system temperature of 67.1 K, a noise variance of 29.2 mK is expected. This is 5% higher than the average variance of 27.7 mK, which may indicate an inaccuracy in the method of estimating  $T_{\text{sys}}$  from the DC levels of the unfiltered data (see Section 4.2.2). The actual variance implies a system temperature




Figure 7.8: Fitted variances and means of the noise histograms for the NCP observing schedules.

of 63.7 K.

The averages of the noise histograms reveal the effects of the RFI flagging (Chapter 6). The noise averages are plotted in Figure 7.10 for data without RFI flagging and for the five RFI flagging levels. Without flagging, the noise is biased to negative values because `Descart` erroneously fits offsets to the spurious RFI events. These erroneous offsets are subtracted from the surrounding data, biasing the estimated noise to negative values. When RFI flagging is enabled, this negative bias vanishes. However, the most aggressive RFI flagging levels show increased scatter because they are flagging too much data. These comparisons show that level 2 RFI flagging strikes a fair balance.

The resulting map is shown in Figure 7.11. A `Descart` offset length of 600 samples was used for this map. The map uses a HEALPix  $N_{\text{side}}$  of 256, corresponding to 13.7' pixels. Point sources and extended structure are seen. A map of the integration time per pixel is shown in Figure 7.12. The average integration time per pixel is 134 s. A gradient in the integration time across the field is clear, which arises from the asymmetric scanning seen in Figure 7.2. This gradient would not be present if the observations spanned a full year, rather than three months.

A key parameter when running `Descart` is the length of the fitted offsets. The algorithm's ability to distinguish between correlated noise and astronomical structure at scales larger than this length is weaker than for smaller scales. Thus, a greater offset length in `Descart` will lead to higher fidelity in extended structure. On the other hand, the offsets must be contiguous, and increasing this length




Figure 7.9: Noise variances for the kernel-filtered data and astronomically calibrated  $T_{\text{sys}}$  values. The correlation between noise variance and  $T_{\text{sys}}$  is clear.

beyond the lengths of the shorter azimuth scans will reduce the amount of data available. As can be seen in Figure 7.1, the scans at the highest and lowest elevations are shorter in azimuth, so setting too great of an offset length will reduce the integration times at large distances from the NCP. The resulting differences in integration time are illustrated in Figure 7.13, which compares the total integration time per pixel between `Descart` runs with offset lengths of 100 and 1000 samples. This effect is most dramatic when offset lengths exceed 600 samples.

The fidelity of the map can be tested through null tests, in which the time series data are divided into two sets. Two maps are produced, and the difference between them analyzed. The presence of systematic errors will appear as features in the differenced map. Applying this approach to the present data set is difficult, however, because the correlated noise removal depends on having the individual pixels crossed by multiple scans at multiple angles. Splitting the data into subsets reduces the number of scan angles per pixel, reducing the effectiveness of correlated noise removal. The results of the null tests should therefore be viewed as upper limits on the systematic errors in the final map.

A null test is performed by splitting the data into two subsets, consisting of alternating weeks of observations. The weeks of observation are divided into sequential weeks: week 1 (Dec. 5 to Dec. 11, 2011), week 2 (Dec. 12 to Dec. 16, 2011), and so on. The first subset consists of odd weeks 1,






Figure 7.10: Fitted means of the noise histograms for the variety of RFI flagging options. A higher level of RFI flagging indicates more data being flagged.




Figure 7.11: Map of the NCP region. Graticules show equatorial coordinates, with declination graticules spaced by  $3^{\circ}$ . Greyscale indicates sky temperature; darker represents high sky temperature. Several point sources are seen, including 3C61.1 at  $02^{\text{h}}22^{\text{m}}48^{\text{s}} + 86^{\circ}18'59''$ , as well as significant extended structure.




Figure 7.12: Map of the integration time across the NCP region. Graticules show equatorial coordinates, with declination graticules spaced by  $3^\circ$ . Greyscale indicates the number of 10 ms samples per pixel, with a logarithmic colourmap. Darker greys represents greater integration times. A gradient across the region is clearly seen.




Figure 7.13: Difference between the integration times per pixel when using `Descart` offset lengths of 100 samples and 1000 samples. Graticules show equatorial coordinates, with declination graticules spaced by  $3^\circ$ . Greyscale indicates the number of 10 ms samples per pixel, with a logarithmic colourmap. Darker greys represents greater integration times. The radial steps are due to scans at large elevation offsets being excluded by the 1000 sample offset length.




Figure 7.14: Null test of the NCP map for a `Descart` offset length of 600. Graticules show equatorial coordinates, with declination graticules spaced by  $3^\circ$ . Greyscale indicates the difference in sky temperature. Darker greys represents excess temperature in the first subset. A gradient across the map is present at the mK level.




Figure 7.15: RMS of the null maps calculated in  $1.5^\circ$ -radius regions at Declination  $85.5^\circ$ . Null maps were calculated for `Descart` offset lengths of 200, 400, 600, and 1000.

3, 5, etc., while the second subset includes the even weeks 2, 4, etc. When divided in this way, the first subset includes 71 schedules and the second 59. The difference between the maps is shown in Figure 7.14. A gradient is present in this map at the mK level, indicating that the correlated noise removal was imperfect. However, the features in this map are not correlated with the structure in Figure 7.11, indicating that the extended structure in the map is reliable to this level.

The null test can be used to calculate the noise level in the map. In order to minimize the influence of the map's gradient, the RMS of the null map was calculated in 8  $1.5^\circ$ -radius regions. These regions were centred at Right Ascensions of  $0^h$ ,  $3^h$ ,  $6^h$ , etc, all with a Declination of  $85.5^\circ$ . The calculated RMS values are plotted in Figure 7.15 for `Descart` offset lengths of 200, 400, 600, and 1000 samples. There is a clear correlation between the RMS and the offset length used, which is indicative of longer offsets removing less of the highest-frequency  $1/f$  noise. A system temperature of 63.7 K (as required by the noise variances of Figure 7.8), with an average integration time of 134 s per pixel, predicts an RMS in the null map of 0.25 mK. This is significantly lower than the values calculated by the null test, indicating again the presence of residual systematic errors.

Using the 5 GHz source counts of (Gregory et al. 1996), the expected confusion noise for this map was calculated to be 0.64 mK. The actual noise in the map is measured to be 0.75 mK, which is consistent with the quadrature sum of the predicted confusion noise and measured system noise.

Thus, while increased observing time would decrease the correlated and thermal noise of the map, confusion noise will hinder any dramatic increases in sensitivity.

A test of the map pointing and calibration was performed by fitting a Gaussian function to the point source 3C61.1. A four-parameter model was used: two for the Gaussian centroid, one for the Gaussian amplitude, and one for a constant background level. The width of the Gaussian function was fixed at  $44'$ . The fitting was performed by least squares minimization within a radius of  $100'$  of the source's nominal position. The resulting fit is shown in Figure 7.16. The centroid was offset from the source's nominal coordinates by  $16''$  in right ascension and  $13''$  in declination, and the measured source flux density was  $1.77$  Jy. Attempting the fit using maps with different `Descart` offset lengths showed a scatter in the fitted flux density of 10%, making this measurement consistent with the published value of  $1.90$  Jy (Kühr et al. 1981).




Figure 7.16: Gaussian model fitted to the point source 3C61.1. The black points are the pixel values of the map. The grey curve shows the best-fit model.

## 7.4 Multiwavelength Comparison

Understanding the microwave emissions from a given region requires multiwavelength data. A wide range of radio and microwave frequencies are needed to measure the spectral energy distribution, while infrared and spectral line maps can be used to constrain the emission components. For example,

the approach of Davies et al. (2006) was to use a radio map as a synchrotron emission template, an infrared map as a dust template, and an  $H\alpha$  map as a free-free template. Though a detailed cross-correlation analysis is beyond the scope of this work, the new C-BASS map is briefly compared to existing maps in order to show the usefulness of the C-BASS data.

Table 7.3: Relevant All-Sky Maps

Description	Resolution	Reference
408 MHz	51.0'	(1)
1.4 GHz	35.0'	(2)
<i>WMAP</i> 23 GHz	50.7'	(3)
<i>Planck</i> 30 GHz	32.2'	(4)
<i>WMAP</i> 33 GHz	38.8'	(3)
<i>WMAP</i> 41 GHz	30.6'	(3)
<i>Planck</i> 44 GHz	27.0'	(4)
<i>WMAP</i> 61 GHz	20.9'	(3)
<i>Planck</i> 70 GHz	13.2'	(4)
<i>WMAP</i> 95 GHz	14.8'	(3)
<i>Planck</i> 100 GHz	9.7'	(4)
<i>Planck</i> 143 GHz	7.2'	(4)
<i>Planck</i> 217 GHz	5.0'	(4)
<i>Planck</i> 353 GHz	4.8'	(4)
<i>Planck</i> 545 GHz	4.7'	(4)
<i>Planck</i> 857 GHz	4.3'	(4)
<i>IRAS</i> 100 $\mu\text{m}$	4.3'	(5,6)
<i>IRAS</i> 60 $\mu\text{m}$	4.0'	(5,6)
<i>IRAS</i> 25 $\mu\text{m}$	3.8'	(5,6)
<i>IRAS</i> 12 $\mu\text{m}$	3.8'	(5,6)
<i>Planck</i> CO	5.5'	(4)
FDS8 Dust	6.1'	(7)
WHAM $H\alpha$	60.0'	(8)
<i>Planck</i> CMB	5.0'	(4)

(1): Haslam et al. (1982)

(2): Reich (1982)

(3): Bennett et al. (2012)

(4): Planck Collaboration (2013a)

(5): Beichman et al. (1988)

(6): Miville-Deschênes & Lagache (2005)

(7): Finkbeiner et al. (1999)

(8): Haffner et al. (2003)

A number of relevant all-sky maps are available that span the radio, sub-mm, and infrared spectrum. A list of these maps is given in Table 7.3. Figure 7.17 presents the NCP as seen in these maps. The  $H\alpha$  map has the lowest native resolution at  $1^\circ$ , so all of the other maps have been convolved to this resolution to better facilitate comparison. The colourmaps are scaled independently for each map. The extended structure in these maps evolves with wavelength. The synchrotron emission, at






Figure 7.17: Maps of the NCP region from the surveys of Table 7.3. The maps are  $18^\circ$  on a side and are convolved to  $1^\circ$  resolution. The colourmaps have been scaled independently for each map.




Figure 7.18: Comparison of the new C-BASS map to existing low-frequency maps, whose extended emission is dominated by synchrotron emission. The maps are shown at native resolution. There is considerable agreement in the extended emission and point sources appearing in the three maps.

the lowest frequencies, has a different morphology than the anomalous microwave emission at 23 and 33 GHz. Between the frequencies of 33 GHz and 143 GHz, the CMB temperature anisotropy dominates the field. At higher frequencies, thermal dust emission becomes significant. The morphology of the thermal dust is steady through the far infrared wavelengths. Also shown are the CO map from *Planck*, (derived from the 100, 217, and 353 GHz maps and tracing molecular gas), the FDS8 94 GHz dust template of Finkbeiner et al. (1999), the WHAM  $H\alpha$  map (tracing ionized gas), and the *Planck* estimate of the CMB temperature anisotropy.

As the C-BASS observing frequency was chosen to target synchrotron emission, a closer comparison between the existing synchrotron maps and C-BASS is justified. Figure 7.18 shows the 408 MHz and 1.4 GHz maps alongside the new C-BASS data at native resolution. The presence of systematic errors in the older maps is striking, as is the comparatively low level of systematic errors in the new C-BASS map. Some features that may have been mistaken for systematic errors in the older maps are clearly shown to be astronomical in the new C-BASS map.

The lowest frequency bands of the *WMAP* and *Planck* experiments are dominated by the anomalous microwave emission at the NCP (Bennett et al. 2012; Planck Collaboration 2013b). Figure 7.19 shows the C-BASS NCP map next to the *WMAP* and *Planck* maps, allowing a qualitative comparison. There is a hint of correlated structure between C-BASS and the other maps, though it is fainter by an order-of-magnitude than would be predicted by flat-spectrum anomalous microwave emission.

The C-BASS map is compared to the  $H\alpha$  and FDS8 dust maps in Figure 7.20. In this case, the




Figure 7.19: Comparison of the new C-BASS map to the lowest frequency *WMAP* and *Planck* maps, whose extended emission is dominated by the anomalous microwave emission. The maps are shown at native resolution. There is a hint of correlation between the maps.

maps do not appear to be strongly correlated.

It is clear that a detailed, joint analysis of these multifrequency data will yield new insights into the anomalous microwave emission in this field by constraining the emission’s spectral behaviour and peak frequency. The C-BASS data are critical this due to their improvement in systematic errors over the other low frequency maps and due to the convenient placement of the C-BASS observing frequency.

## 7.5 Discussion

The North Celestial Pole region was mapped using the C-BASS Stokes  $I_1$  channel as a full systems test of the experiment. In collecting these data, the C-BASS antenna operation and scanning strategy were demonstrated. Reduction of the data allowed testing of the pipeline, the calibration strategy, and the *Descart* map maker. The noise properties were found to be consistent with thermal noise predictions, although there is clear evidence of systematic errors in the reduced map. The flux density of the point source 3C61.1 was measured and compared to the published value, and the same was done for the map’s confusion noise. In both cases, the measured values were consistent with expectations, corroborating the calibration.

The noise in the map was found to be near the confusion limit, in spite of the receiver’s very high system temperature at the time of observation. This is encouraging, as it suggests that the full survey, with much lower system temperature, will be limited by confusion rather than systematic




Figure 7.20: Comparison of the new C-BASS map to the WHAM H $\alpha$  map (with units of Rayleighs, R) and the FDS8 94 GHz dust template. The maps are shown at native resolution. There does not appear to be a strong correlation between the maps.

errors.

A cursory comparison of this map with a variety of multwavelength data demonstrated the scientific promise of this experiment.

## Chapter 8

# The Spinning Dust Emission Spectrum

### 8.1 Context

The existence of an anomalous component of diffuse microwave emission is well established, though it has yet to mature as an astrophysical probe. This emission was first detected as a cosmological foreground by Kogut et al. (1996) and de Oliveira-Costa et al. (1997), and was first discovered to be anomalous by Leitch et al. (1997) in observations near the North Celestial Pole. It was quickly demonstrated by Draine & Lazarian (1998b, DL98 hereafter) to be consistent with electric dipole radiation from very small dust grains, a process now commonly referred to as spinning dust emission. This explanation has gained wide favor, though perhaps prematurely, as it remains to be proven that this is the cause of the North Celestial Pole emission. Free-free emission from very hot gas and magnetic dipole emission from dust grains (Draine & Lazarian 1999; Draine & Hensley 2013) may plausibly explain the anomalous emission in this region and in some others.

Anomalous microwave emission has now been observed by many authors in a variety of Galactic and extragalactic environments (Finkbeiner et al. 2002; Finkbeiner 2004; Murphy et al. 2010; Lu et al. 2012; Murphy et al. 2012). The emission is characterized by a broad peak around 20 – 40 GHz, spatial correlation with dust on degree scales, peak brightness roughly four orders of magnitude less than that of thermal dust emission, and little polarization (Rubiño-Martín et al. 2012). Puzzlingly, a strong correlation with infrared tracers of small grains at arcminute scales has not been observed (Tibbs et al. 2011, 2012).

The original model of DL98 was derived under the key assumptions of a Maxwellian distribution of grain angular velocity, grain rotation about the axis of maximum moment of inertia, simple grain geometries (spherical, disk-like, and rod-like), and electric dipole moments of the grains based on random walks over chemical bonds. This model found wide success in fitting anomalous microwave emission measurements (Finkbeiner et al. 2004; Gold et al. 2009).

Lazarian & Draine (2000) explored the theory of polarized spinning dust radiation, finding that the radiation could not be polarized by more than 10%, and even then only below 10 GHz. Observational studies have consistently found upper limits of anomalous microwave emission polarization at the percent level (Dickinson et al. 2007; Mason et al. 2009; Macellari et al. 2011; Rubiño-Martín et al. 2012). These upper limits have been interpreted as supporting the spinning dust model. Most recently, Hoang et al. (2013) used the 2175 Å polarization feature, as observed for two stars, to argue that the spinning dust polarization should peak at 3% at 5 GHz, and decrease rapidly above 20 GHz.

Since DL98, spinning dust theory has advanced in both precision and scope. Rafikov (2006) applied the theory to protoplanetary disks, while Ysard & Verstraete (2010) showed that a quantum treatment gave the same results as the classical approach of DL98. Ali-Haimoud et al. (2009, AHD09 hereafter) advanced the theory by allowing for non-Maxwellian distributions of grain rotation via the Fokker-Planck equation, through refined treatments of the excitation and damping processes, and by producing the SpDust IDL package, which allowed users to calculate custom spectra given astrophysical parameters. Hoang et al. (2010, HDL10 hereafter) considered the dramatic effects of irregular rotation about non-principal axes and used the Langevin equation instead of the Fokker-Planck equation in order to capture the transient effects due to collisions with individual ions. Silsbee et al. (2011, SAH11 hereafter) updated SpDust to include irregular rotation and improved calculations of the rotational damping and excitation. Hoang et al. (2011, HLD11 hereafter) extended the theory to irregularly shaped grains and further explored the distribution of rotational energies arising from vibrational-rotational energy coupling. These refinements and extensions have been accompanied by increases in complexity: the latest models depend on upwards of 30 parameters.

SpDust has had a large impact on this field. It allows for quick calculation of models using nine physical parameters and has shown great utility in fitting observations (Planck Collaboration 2011). Its use has been limited, however, in cosmological foreground separation. Nine parameters is more than the shape of the spectrum justifies, and the code runs too slowly to allow rapid exploration of multi-dimensional parameter space.

Foreground separation efforts have instead resorted to simple, analytical models with three parameters or less. No single function has found wide use. Bonaldi et al. (2007), following Tegmark (1998) and de Oliveira-Costa et al. (2004), suggested a parabola in  $\log S - \log \nu$  space, Tegmark et al. (2000) proposed a modified greybody, and Gold et al. (2009) simply shifted the numerical models of DL98. Although expedient for fitting, these approaches do not easily lead to astrophysical interpretation. An analytical function that is easily relatable to the physics would offer an advantage: it would be well suited to fitting and to interpretation. It is the aim of this work to provide such a function by analytically deriving the spinning dust spectrum through use of careful approximations.

The approach in this study is to follow the derivation of SAH11, but using analytical approximations where numerical calculations would otherwise be required. The approximations are tested

against the results of SpDust to demonstrate where they succeed in capturing the numerical model. As this approach uses the Fokker-Planck equation instead of the Langevin equation, it is not possible to reproduce the transient spin-up effects of HDL10. These effects were shown to be contained in the high-frequency fall-off the spectrum. Their exclusion does represent a limitation of this new approach, although uncertainties in the dust grain size distribution may be more significant. The present treatment of irregular rotation can be applied to triaxial grains and the range of vibrational-rotational energy coupling considered in HLD11. However, as these effects are not calculated in SpDust, their calculation is beyond the scope of this study. These omissions should be considered when applying this model to data. Polarization of the spinning dust radiation is not considered.

This chapter is structured as follows. An overview of the derivation steps is provided in Section 8.2. Section 8.3 describes the assumed dust properties. The rotational distribution function, and its dependence on environment, is discussed in Section 8.4. Section 8.5 presents the emissivity itself, including a treatment of irregular rotation. Finally, Section 8.6 provides a discussion of the derived function, its use, and various caveats.

## 8.2 Overview

The total emissivity of an ensemble of rotating grains,  $j_\nu/n_H$ , is the integral of the emissivity of grains of a given size,  $j_\nu^a$ , weighted by the grain size distribution,  $1/n_H dn_{\text{gr}}/da$ . This is written as

$$\frac{j_\nu}{n_H} = \frac{1}{n_H} \int_{a_{\text{min}}}^{a_{\text{max}}} da \frac{dn_{\text{gr}}}{da} j_\nu^a. \quad (8.1)$$

The grain emissivity is calculated by integrating emitted power over the angular momentum and electric dipole moment distributions,

$$j_\nu^a = 2\pi \int_0^\infty 4\pi J^2 dJ \int_0^\infty d\mu f_a(J, \mu) P(\mu) \frac{P_{\text{ed},\omega}(J, \mu)}{4\pi}, \quad (8.2)$$

in which  $f_a(J, \mu)$  is the angular momentum distribution function for grains of size  $a$  and electric dipole moment  $\mu$ ,  $P(\mu)$  is the electric dipole moment distribution, and  $P_{\text{ed},\omega}(J, \mu)$  is the power emitted at frequency  $\omega$ . This last function accounts for the complex, torque-free motion of aspherical grains (called “wobbling” in HDL10 and “tumbling” in SAH11).

The integrals benefit from two changes of variable. The first is to calculate rotation using the ratio of angular momentum to maximum moment of inertia

$$\Omega \equiv \frac{J}{I_M} \quad (8.3)$$

rather than the angular momentum itself.  $\Omega$  is henceforth referred to as the rotation rate, though

it is understood that this label is only truly accurate in the non-tumbling case. The second is to separate the electric dipole moment from the grain size using a new variable  $b$  (as discussed in Section 8.3.5). Equation 8.2 then becomes

$$j_\nu^a = 2\pi \int_0^\infty \Omega^2 d\Omega \int_0^\infty db f_a(\Omega, b) P(b) P_{\text{ed},\omega}(\Omega, b). \quad (8.4)$$

The strategy adopted in this work is to make a number of judicious simplifications aimed at approximating Equation 8.4 as a log-normal distribution function. Equation 8.1 is then evaluated analytically to give the desired result.

A large number of symbols are used in this chapter. The most important of these are compiled in Table 8.1.

## 8.3 Dust Grains

Spinning dust emission is sensitive to fundamental properties of the grains. The grain sizes and permanent electric dipole moments are the most important; geometry and charge are of lesser concern. A simple, thermal calculation shows that rotation at tens of GHz requires sub-nm grains. Such a population is consistent with the polycyclic aromatic hydrocarbon population described in Weingartner & Draine (2001, WD01 hereafter) and Draine & Li (2007), although debate persists regarding the relative importance of aliphatic and aromatic structures in these grains (Kwok & Zhang 2011). This population is needed to explain the observed infrared emission and its properties can be constrained by observations of ultraviolet extinction.

### 8.3.1 Size

If  $a$  is the spherical-equivalent radius, then the number of atoms per grain is roughly

$$N_{\text{at}} \approx 600 \left( \frac{a}{1 \text{ nm}} \right)^3, \quad (8.5)$$

consistent with the prescription of Li & Draine (2001) if there is one hydrogen atom for every three carbons. The grains of interest thus contain fewer than 600 atoms. The smallest may be plausibly described as large molecules.

A log-normal size distribution is conventionally assumed for these grains (WD01; Compiègne et al. 2011). As noted by WD01, this form is not motivated by physics, but by mathematical convenience. This distribution is accompanied by a second log-normal distribution peaking at 3 nm and a power-law extending beyond 0.1  $\mu\text{m}$ , though these additional components are insignificant below 1 nm. Inspired by photolytic considerations (Guhathakurta & Draine 1989), the distribution is assumed



Table 8.1: Important variables used in this chapter

Variable	Description	Equation
$a$	Grain size	(8.5)
$N_{\text{at}}$	Number of atoms	(8.5, 8.6)
$\frac{1}{n_H} \frac{dn_{\text{gr}}}{da}$	Distribution of $a$	(8.6)
$B_1$	Normalization of $a$ distribution	(8.6)
$a_m$	Minimum grain size	(8.6)
$a_0$	Peak grain size	(8.6)
$\sigma$	Width of $a$ distribution	(8.6)
$b$	Normalized electric dipole moment	(8.8)
$\epsilon_{\text{ip}}$	In-plane $b$ fraction	(8.11)
$P(b)$	Distribution of $b$	(8.9)
$\beta$	Width of $b$ distribution	(8.9)
$N_b$	Dimension of $b$ distribution	(8.9)
$\Omega$	Grain rotation rate	(8.3)
$\omega = 2\pi\nu$	Radiation frequency	(8.22)
$q_r$	Ratio of $\omega$ to $\Omega$	(8.22)
$f_a(\Omega, b)$	$\Omega$ distribution function	(8.12)
$A_\Omega$	Exp. coefficient of $\Omega$ distribution	(8.15)
$\alpha_a$	Power law on $a$	(8.15)
$\alpha_b$	Power law on $b$	(8.15)
$\alpha_\nu$	Power law on $\Omega$	(8.15)
$\Omega_{p,a}$	Peak $\Omega$ for size $a$	(8.17)
$\mathcal{I}_a(\Omega)$	$\Omega$ distribution, integrated over $\beta$	(8.32, 8.33)
$\mathcal{I}_0$	Normalization of $\mathcal{I}_a$	(8.34)
$\sigma_\Omega$	Width of $\mathcal{I}_a$	(8.35)
$P_{\text{ed},\omega}(\Omega, b)$	Emission from single grain	(8.20)
$R(\omega, \Omega)$	Dimensionless emission spectrum	(8.20, 8.27)
$R_0$	Normalization of $R$	(8.27)
$\sigma_r$	Width of $R$	(8.27)
$j_\nu^a$	Emissivity for grains of size $a$	(8.4, 8.31, 8.36)
$\sigma_\nu$	Width of $j_\nu^a$	(8.37)
$j_\nu/n_H$	Total emissivity per H	(8.1, 8.39)
$\alpha_s$	Power law of $j_\nu/n_H$	(8.41)
$\sigma_s$	Log-normal width of $j_\nu/n_H$	(8.42)
$\nu_0$	Characteristic frequency of $j_\nu/n_H$	(8.40)
$\eta_\nu$	Error function slope for $\nu$	(8.45)
$\eta_a$	Error function slope for $a_m$	(8.46)

to truncate sharply at a smallest grain size. The size distribution is therefore approximated as

$$\frac{1}{n_H} \frac{dn_{\text{gr}}}{da} = \begin{cases} 0 & a < a_m \\ \frac{B_1}{a} \exp \left\{ -\frac{1}{2} \left[ \frac{\log(a/a_0)}{\sigma} \right]^2 \right\} & a \geq a_m. \end{cases} \quad (8.6)$$

Following WD01 and AHD09, the values  $B_1 = 1.2 \times 10^{-6}$ ,  $\sigma = 0.4$ , and  $a_0 = 3.5 \text{ \AA}$  are used when calculating model parameters, though in practice these can be varied if the data require. In particular,  $B_1$  represents the abundance of the small grains and there is no reason to expect it to be fixed by nature. Departing from previous approaches, it is not assumed that  $a_0$  and  $a_m$  are equal.

The log-normal form of Equation 8.6 heavily influences the mathematics of this work. Different size distributions would require different approximations to be made in Sections 8.5.1 and 8.5.2, resulting in a qualitatively different analytical forms for  $j_\nu/n_H$ .

### 8.3.2 Shape

It is unlikely that these grains have simple shapes. DL98 assumed rod-like and disk-like geometries for the smallest grains, inspired by aliphatic and aromatic molecules. Spherical shapes were assumed for larger grains. Sharp transitions between these occur at sizes  $a_1$  and  $a_2$ , with the grains smaller than  $a_1$  being rod-like, grains larger than  $a_1$  but smaller than  $a_2$  being disk-like, and grains larger than  $a_2$  being spherical. Later models followed DL98 in setting  $a_2 = 6 \text{ \AA}$ , but set  $a_1 = 0$ . These precedents are followed here, though it is shown in Section 8.5.3 that  $a_2$  has little effect on the final result.

Grain shape influences the rotational distribution functions by way of grain cross sections, charge distributions, and electric dipole moment geometry. These effects are important when numerically calculating the parameters of Equation 8.15 and when considering the irregular rotation of grains. In Section 8.5.1, the effects of geometry and irregular rotation are parameterized as part of the full derivation. This parameterization is applicable to the range of plausible geometries, although only disk-like and spherical grains are explicitly considered.

### 8.3.3 Temperature

The internal temperatures of the grains are not constant. The grains are transiently heated by UV photon absorption and cool to near ground state before the next UV photon is absorbed. The result is a grain temperature distribution (Guhathakurta & Draine 1989). This is important for the rotational distribution functions, as it will affect IR photon emission rate and the atom desorption rates and evaporation temperatures.

The temperatures of the grains will be coupled to their rotational energy. HDL10 and HLD11

showed that the strength of this coupling, or the rate of internal relaxation, has a significant effect on the grain tumbling. In the case of strong coupling, there will be a minimum vibrational temperature at which coupling can occur (due to the sparsity of the vibrational mode spectrum at low temperatures). If this temperature is much greater than the rotational energy of the grain, then there will be a uniform distribution of  $\sin \theta$ , where  $\theta$  is the rotation angle: the angle between the grain's angular momentum and axis of maximum moment of inertia. This is the case considered by SAH11. Conversely, if the decoupling temperature is much less than rotational energy, then the rotation angle will be zero. When coupling is weak, the rotation angle is governed by a Maxwellian distribution. Section 8.5.1 explicitly calculates the grain tumbling in the case of strong coupling with a high decoupling temperature, though the suggested parameterization can also be applied to the other cases.

### 8.3.4 Charge

Collisional and photoelectric charging of grains has implications for electric dipole moments and interaction cross-sections. DL98 considered this and presented the charge distribution functions for a variety of grain sizes and environments. Sub-nm grains had typical charges between  $-1$  and  $3$ . Such small charges are unlikely to dominate the electric dipole moments. They are, however, important when calculating grain rotation rates (see Section 8.4).

### 8.3.5 Dipole Moment

The intrinsic, electric dipole moments of the grains are poorly constrained observationally, and attempts to derive them theoretically are subject to uncertainty in the specific chemical compositions of the grains. DL98 instead assumed a typical moment per molecular bond  $b$  and used a random walk over all bonds to get the total dipole moment  $\mu$ . AHD09 extended this by having  $\mu$  normally distributed. The variance is then

$$\langle \mu^2 \rangle = N_{\text{at}} \beta^2. \quad (8.7)$$

Note that this distribution is a function of grain size.

With the aim of separating the integrals over grain size and dipole moment cleanly, a different approach was taken. Defining the normalized dipole moment,  $b$ , via

$$b^2 \equiv \frac{\mu^2}{N_{\text{at}}} \quad (8.8)$$

allows use of the normal distribution

$$P(b) = \frac{2 \left(\frac{N_b}{2}\right)^{\frac{N_b}{2}}}{\Gamma\left(\frac{N_b}{2}\right)} \frac{1}{b} \left(\frac{b}{\beta}\right)^{N_b} \exp\left[-\frac{N_b}{2} \left(\frac{b}{\beta}\right)^2\right]. \quad (8.9)$$

$N_b$  is the dimensionality of the distribution and is 1, 2, or 3 for linearly, cylindrically, and spherically distributed dipole moments.  $\Gamma(x)$  is the Gamma function. When calculating numerical parameters,  $\beta$  is taken as 0.4D.

SAH11 considered the case of disk-like grains with three-dimensional electric dipole moment distributions, which could be due to disk warping from pentacyclene structures. They parameterize this possibility via the in-plane fraction of the dipole moment:

$$\epsilon_{\text{ip}} = \frac{\langle \mu_{\text{ip}}^2 \rangle}{\langle \mu^2 \rangle}. \quad (8.10)$$

The out-of-plane fraction is similarly defined, and

$$\epsilon_{\text{ip}} + \epsilon_{\text{op}} = 1. \quad (8.11)$$

## 8.4 Distribution Function

The physics of dust grain rotation is nontrivial. Desired is a rotational distribution function, which will be a function of grain size and astrophysical environment. Smaller grains tend to rotate faster due to smaller moments of inertia, while for a given grain size, the preferred rotation rate is the result of a variety of excitation and damping mechanisms.

It is useful to think of rotation rate using an intuitive picture, in which the various excitation and damping mechanisms are competing to thermalize grain rotation to their respective temperatures. Torques from the emission of infrared photons (which follow thermal spikes due to UV photon absorption) push the grain rotation towards the average IR radiation temperatures, which depend upon the grain heat capacities and emission spectra and can reach  $\sim 10^3$  K. Desorption of atoms (adsorbed via gas collisions) pushes to the evaporation temperature, of order  $10^2$  K. Plasma interactions cause the rotation to tend to the gas temperature, which varies widely with interstellar phase. At the same time, drag from the electric dipole emission itself can limit grain rotation, causing the distribution function to fall off non-thermally at high rotation rates.

Detailed treatments of these effects need to be done numerically, and no attempt to reproduce or improve upon these efforts are made here. See DL98, AHD09, HDL10, and SAH11 for careful discussions and calculations. Rather, it will be shown how a simple parameterization of the distribution function can encompass the important effects.

AHD09 and SAH11 used the Fokker-Planck equation to calculate the distribution function. This differential equation allows one to account for damping and excitation of a stationary system from small impulses. Adapted from SAH11,

$$\frac{df_a(\Omega, b)}{d\Omega} + \frac{I_M \Omega^2 F}{kT} \frac{f_a(\Omega, b)}{G \Omega} = 0, \quad (8.12)$$

where

$$F = \sum_j F_j \quad (8.13)$$

and

$$G = \sum_j G_j \quad (8.14)$$

are the sums of the dimensionless damping and excitation coefficients. Note that, while SAH11 treated the electric dipole damping as a separate term, it is included here as one of the  $F_j$ . In the current work, the above is simplified further by assuming that the actions of dipole moment, grain size, and rotation frequency are separable and are described by power laws (with influences of ISM environment, grain charge, and grain temperature being folded into the parameterization):

$$\frac{I_M \Omega^2}{kT} \frac{F}{G} \equiv \alpha_\nu A_\Omega \left(\frac{b}{\beta}\right)^{\alpha_b} \left(\frac{a}{a_0}\right)^{\alpha_a} \left(\frac{\Omega}{\Omega_{p,a_0}}\right)^{\alpha_\nu}. \quad (8.15)$$

In quantifying these assumptions, this equation serves as the definition of the power-law indices and the peak rotation frequency  $\Omega_{p,a_0}$  for grains of size  $a_0$ .  $\Omega_{p,a_0}$  is guaranteed to be the peak frequency of the rotational distribution function by the definition of  $A_\Omega$ ,

$$A_\Omega^2 \equiv \left(\frac{N_b}{2}\right)^{\alpha_b} \frac{[\Gamma(8/\alpha_\nu)\Gamma(N_b/2 + 1 - 5\alpha_b/2\alpha_\nu)]^{8\alpha_\nu}}{[\Gamma(7/\alpha_\nu)\Gamma(N_b/2 + 1 - 2\alpha_b/\alpha_\nu)]^{5\alpha_\nu} [\Gamma(9/\alpha_\nu)\Gamma(N_b/2 + 1 - 3\alpha_b/\alpha_\nu)]^{3\alpha_\nu}}. \quad (8.16)$$

Although quite useful, it should be clear that reducing the Fokker-Planck equation to this form may introduce degeneracy amongst the astrophysical parameters and ultimately limit the physics one can infer when fitting this model.

The power law indices and rotational peak can be acquired directly from Equation 8.15 given numerically calculated tables of the  $F_j$  and  $G_j$ . SpDust was used to do this for the idealized environments of DL98: cold neutral medium (CNM), dark cloud (DC), molecular cloud (MC), photodissociation region (PDR), reflection nebula (RN), warm ionized medium (WIM), and warm neutral medium (WNM). The results of these calculations are presented in Table 8.2. For each environment, the parameters are calculated for disk-like (1) and spherical (2) grains at sizes of 4.5 Å and at 6.3 Å. The peak rotation frequencies are extrapolated to  $a_0$  and  $a_2$ .

These parameters are plotted in Figure 8.1. If the astrophysical parameters were not already degenerate prior to imposing Equation 8.15, then the correlations between these parameters ought to be weak. As can be seen, however, the parameters are highly correlated. The dichotomy seen in  $\alpha_b$  and  $\alpha_\nu$  reflects the cases of whether or not electric dipole damping is dominant, with  $\alpha_b \approx 0$  and  $\alpha_\nu \approx 2$  in the latter case. The correlation between  $\alpha_a$  and  $\Omega_{p,a_0}$  is more subtle, as the various  $F_j$  and  $G_j$  depend on  $a$  to different degrees. The result of these correlations is the implication that the astrophysical parameters are themselves highly degenerate, and that inferring environmental physics

Table 8.2: Rotational distribution function parameters.

Environment		Parameter									
		Calculated at 4.5 Å					Calculated at 6.3 Å				
		$\alpha_a$	$\alpha_b$	$\alpha_\nu$	$\Omega_{p,a}$	$\Omega_{p,a_0}$	$\alpha_a$	$\alpha_b$	$\alpha_\nu$	$\Omega_{p,a}$	$\Omega_{p,a_2}$
CNM	(1)	5.01	1.63	3.86	98.6	135.3	5.49	1.32	3.36	60.9	66.1
	(2)	5.09	1.54	3.93	153.3	210.3	5.43	1.29	3.58	94.8	102.3
DC	(1)	5.69	0.12	2.10	70.9	137.1	6.10	0.06	2.01	26.8	31.3
	(2)	4.56	0.14	2.13	115.4	194.4	4.60	0.11	2.04	54.4	61.0
MC	(1)	5.05	1.59	3.32	126.5	183.3	5.93	1.47	2.58	67.2	75.4
	(2)	3.53	1.57	3.36	200.9	259.7	5.63	1.60	2.83	124.8	137.9
PDR	(1)	6.66	0.13	2.13	367.6	787.4	6.67	0.01	2.01	120.6	142.5
	(2)	5.95	0.25	2.25	653.2	1243.2	5.79	0.03	2.03	253.8	292.9
RN	(1)	6.19	0.31	2.33	200.5	383.3	6.23	0.04	2.03	74.0	86.4
	(2)	5.20	0.52	2.54	355.5	584.9	5.31	0.10	2.10	161.3	183.1
WIM	(1)	4.77	1.69	3.97	93.4	125.2	4.39	1.60	3.79	62.6	66.4
	(2)	4.92	1.57	4.01	146.6	197.8	4.99	1.44	3.88	95.1	101.5
WNM	(1)	4.28	1.73	3.95	81.6	106.3	4.21	1.58	3.71	56.2	59.4
	(2)	4.62	1.64	3.99	127.9	169.7	4.80	1.50	3.82	84.4	89.9

(1): Disk-like grains

(2): Spherical grains

from this emission would be challenging even without the power law approximation.

It is useful to define,  $\Omega_{p,a}$ , the peak rotation frequency for grains of size  $a$ ,

$$\left(\frac{a}{a_0}\right)^{\alpha_a} \left(\frac{\Omega_{p,a}}{\Omega_{p,a_0}}\right)^{\alpha_\nu} \equiv 1. \quad (8.17)$$

Equation 8.15 becomes

$$\frac{I\Omega^2}{kT} \frac{F}{G} = \alpha_\nu A_\Omega \left(\frac{b}{\beta}\right)^{\alpha_b} \left(\frac{\Omega}{\Omega_{p,a}}\right)^{\alpha_\nu} \quad (8.18)$$

and the Fokker-Planck equation is integrated to give

$$f_a(\Omega, b) = \frac{\alpha_\nu A_\Omega^{3/\alpha_\nu}}{4\pi\Gamma(3/\alpha_\nu)\Omega_{p,a}^3} \left(\frac{b}{\beta}\right)^{3\alpha_b/\alpha_\nu} \exp\left[-A_\Omega \left(\frac{b}{\beta}\right)^{\alpha_b} \left(\frac{\Omega}{\Omega_{p,a}}\right)^{\alpha_\nu}\right]. \quad (8.19)$$

The validity of this approach is demonstrated in Figure 8.2. Distribution functions for a variety of grain sizes are plotted for the CNM environment. Curves calculated from SpDust are compared to the results of the power law approximation (extrapolated from 4.5 Å grains). Agreement is satisfactory for grains smaller than 6 Å, but then deteriorates rapidly. This disagreement is due to variations in  $\alpha_a$  and  $\alpha_\nu$  (indicating a failure of the power law approximation), also shown in this figure. This figure demonstrates that the power law approximation is reasonable over small ranges




Figure 8.1: Power law parameters from Equation 8.15 plotted against peak rotation frequency for the ideal interstellar environments. Parameters are calculated numerically from SpDust. Data are listed in Table 8.2.




Figure 8.2: Rotational distribution functions for grains of various sizes in the CNM environment with  $b = \beta$  are shown in the upper panel. Solid, black curves show the analytical, power law approximation. Grey, dashed curves show the numerically calculated functions from SpDust. Lower panel shows the variation of power law parameters with grain size.

in grain size, but becomes a significant source of error when used over wider ranges of sizes and frequencies. In particular, one should be careful when using parameters derived at  $a_0 = 3.5 \text{ \AA}$  when  $a_m > 6 \text{ \AA}$ .

## 8.5 Emissivity

It is now possible to seek a solution to the integrals in Equations 8.1 and 8.4. Further approximations will be needed in order to achieve an analytical result, and a log-normal form is suggested by the grain size distribution. These integrals are approached with this goal in mind. The effects of grain tumbling are considered first, leading to a log-normal approximation for  $j_\nu^a$ , which then allows the completion of the derivation of  $j_\nu/n_H$ .



In the following, tumbling is only considered in the case of axisymmetric grains with strong vibrational-rotational coupling and high decoupling temperatures, resulting in a uniform distribution in the Sine of the rotation angle. Although these techniques may also be applied to triaxial grains and the broader variety of vibrational-rotational coupling, these cases are not treated here.

### 8.5.1 Grain Tumbling

The effects of irregular grain rotation manifest themselves in the power emitted by a single grain. In general, this can be written as

$$P_{\text{ed},\omega}(\Omega, b) = \mathcal{P}_{p,a} \frac{b^2}{\beta^2} \frac{\Omega^4}{\Omega_{p,a}^4} R(\omega, \Omega) \quad (8.20)$$

where

$$\mathcal{P}_{p,a} \equiv \frac{2}{3} \frac{\beta^2 N_{\text{at}} \omega_{p,a}^4}{c^3}. \quad (8.21)$$

This form naturally allows for the various permutations of grain geometry and rotation dynamics, with the emission spectrum itself being contained in the  $R$  function. The emission frequency is related to rotation frequency by

$$\omega \equiv q_r \Omega \quad (8.22)$$

and the emission spectrum is described by  $R(\omega, \Omega)$ .

In the non-tumbling case,  $q_r = 1$ . The spectrum is a delta function:

$$R(\omega, \Omega) = \epsilon_{\text{ip}} \delta(\omega - \Omega) \quad (8.23)$$

or

$$R(\omega, \Omega) = \frac{2}{3} \delta(\omega - \Omega) \quad (8.24)$$

for disk-like and spherical grains, respectively. In the case of tumbling, there is emission due to the in-plane and out-of-plane electric dipole moments. From SAH11, the out-of-plane emission has  $q_r = 2$  and

$$R(\omega, \Omega) = \frac{2\epsilon_{\text{op}}}{3} \delta(\omega - 2\Omega) \quad (8.25)$$

while in-plane has

$$R(\omega, \Omega) = \begin{cases} \frac{\epsilon_{\text{ip}}}{4q_r^4} \frac{\omega^4}{\Omega^5} \left(3 - \frac{\omega}{\Omega}\right)^2 & \Omega < \omega < 3\Omega \\ \frac{\epsilon_{\text{ip}}}{2q_r^4} \frac{\omega^4}{\Omega^5} \left(1 - \frac{\omega^2}{\Omega^2}\right) & \omega < \Omega \end{cases}. \quad (8.26)$$

The latter case does not lend itself to analytical progress, so it is approximated with a log-normal function having the same first and second moments. This approximation is shown in Figure 8.3.




Figure 8.3: Exact vs. approximate forms of the tumbling spectrum for in-plane emission from disk-like grains.

The fit is clearly not perfect, yet it deviates by less than 10 % of the peak across most of the range.

The approximated function is

$$R(\omega, \Omega) \approx \frac{R_0}{\sqrt{2\pi}\sigma_r\omega} \exp \left\{ -\frac{1}{2} \left[ \frac{\log(\omega/q_r\Omega) - \sigma_r^2}{\sigma_r} \right]^2 \right\} \quad (8.27)$$

with integral

$$R_0 = \frac{5\epsilon_{\text{ip}}}{q_r^4} \quad (8.28)$$

width

$$\sigma_r^2 \approx 0.0518 \quad (8.29)$$

and peak

$$q_r \approx 1.775. \quad (8.30)$$

Triaxial grains, lower vibrational-rotational coupling temperatures, and weak vibrational-rotational coupling, as described in HDL10 and HLD11, are not explicitly considered here. The above approximation can be applied to these cases, resulting in different values for  $q_r$ ,  $\sigma_r$ , and  $R_0$ .

It is convenient to continue the derivation using the log-normal form of  $R(\omega, \Omega)$ . The results can be applied to non-tumbling cases by setting  $\sigma_r = 0$  and  $R_0$  equal to the coefficients in Equations 8.23, 8.24, and 8.25.

### 8.5.2 Grain Emissivity

The integrals of Equation 8.4 can be rewritten as

$$j_\nu^a = \frac{1}{2} \mathcal{P}_{p,a} \int_0^\infty d\Omega R(\omega, \Omega) \mathcal{I}_a(\Omega) \quad (8.31)$$

in which

$$\mathcal{I}_a(\Omega) = 4\pi \int_0^\infty db \frac{b^2}{\beta^2} P(b) \frac{\Omega^6 f_a(\Omega, b)}{\Omega_{p,a}^4}. \quad (8.32)$$

In this, the functions  $\mathcal{I}_a d\Omega$  represents the dimensionless rotation spectrum and  $R$  the dimensionless emission spectrum for a given rotation rate. Equation 8.31 is a simple convolution of these.  $R$  has already been approximated as a log-normal function, so if a similar approximation to  $\mathcal{I}_a$  can be found, then  $j_\nu^a$  will have the desired form.

The main concern is whether the assumed power law  $\alpha_b$  is constant over the  $b$  integral. A nonzero  $\alpha_b$  indicates that grains of different dipole moments will rotate at different frequencies, so deviation from the power-law assumption leads to errors in the width and peak frequency of this integral. Figure 8.4 shows  $\alpha_b$  for the ideal environments and demonstrates that the power-law assumption is reasonable, as the  $\alpha_b$  do not change greatly over the peak of  $P(b)$ .

A log-normal approximation for  $\mathcal{I}_a(\Omega)$  is achieved by calculating the first and second moments of  $\Omega$  over Equation 8.32. The integral over  $b$  becomes analytical once  $\Omega$  has been integrated, giving the result

$$\mathcal{I}_a(\Omega) \approx \frac{\mathcal{I}_0}{\sqrt{2\pi}\sigma_\Omega} \frac{1}{\Omega} \exp \left\{ -\frac{1}{2} \left[ \frac{\log(\Omega/\Omega_{p,a}) - \sigma_\Omega^2}{\sigma_\Omega} \right]^2 \right\} \quad (8.33)$$

with

$$\mathcal{I}_0 = \left( \frac{N_b}{2} \right)^{2\alpha_b/\alpha_\nu} \frac{\Gamma(7/\alpha_\nu) \Gamma(N_b/2 + 1 - 2\alpha_b/\alpha_\nu)}{\Gamma(N_b/2 + 1) \Gamma(3/\alpha_\nu)} A_\Omega^{-4/\alpha_\nu} \quad (8.34)$$

and

$$\sigma_\Omega^2 = \log \left[ \frac{\Gamma(7/\alpha_\nu) \Gamma(N_b/2 + 1 - 2\alpha_b/\alpha_\nu) \Gamma(9/\alpha_\nu) \Gamma(N_b/2 + 1 - 3\alpha_b/\alpha_\nu)}{\Gamma(8/\alpha_\nu)^2 \Gamma(N_b/2 + 1 - 5\alpha_b/2\alpha_\nu)^2} \right]. \quad (8.35)$$

The emissivity for grains of size  $a$  then follows immediately:

$$j_\nu^a \approx \frac{\mathcal{P}_{t,a}}{4\pi} \frac{1}{\sqrt{2\pi}\sigma_\nu} \exp \left\{ -\frac{1}{2} \left[ \frac{\log(\nu/\nu_{p,a}) - \sigma_\nu^2}{\sigma_\nu} \right]^2 \right\} \quad (8.36)$$

in which

$$\sigma_\nu^2 = \sigma_r^2 + \sigma_\Omega^2 \quad (8.37)$$




Figure 8.4:  $\alpha_b$  vs  $b$  for  $3.5 \text{ \AA}$  grains at peak emission frequency. Curves are calculated via SpDust. Light grey curves show  $bP(b)$  for  $N_b$  of 2 and 3.

and

$$\mathcal{P}_{t,a} = \frac{2}{3} \frac{\beta^2 N_{\text{at}} \omega_{p,a}^4}{c^3} R_0 \mathcal{I}_0. \quad (8.38)$$

This is plotted in Figure 8.5 for disk-like grains in both the tumbling and non-tumbling cases. The analytical functions continue to show satisfactory agreement with the numerically derived curves.

### 8.5.3 Total Emissivity

Given the above approximations, the integration over grain size follows analytically,

$$\frac{j_\nu}{n_H} = \frac{j_\nu}{n_H} \Big|_{\nu_0} \left( \frac{\nu}{\nu_0} \right)^{\alpha_s} \exp \left\{ -\frac{1}{2} \left[ \frac{\log(\nu/\nu_0)}{\sigma_s} \right]^2 \right\} \text{erfc} \left[ \eta_\nu \log \frac{\nu}{\nu_0} + \eta_a \log \frac{a_m}{a_0} \right]. \quad (8.39)$$

The characteristic frequency

$$\nu_0 \equiv \nu_{p,a_0} \exp(-\alpha_s \sigma_\nu^2) \quad (8.40)$$

is that at which grains of size  $a_0$  make their greatest fractional contribution to the total emissivity (assuming a flat grain size distribution). The power law and log-normal width are

$$\alpha_s \equiv 3 - 3 \frac{\alpha_\nu}{\alpha_a} \quad (8.41)$$




Figure 8.5: Emissivity for an ensemble of grains of size  $3.5 \text{ \AA}$  in the CNM environment. Disk-like grains are assumed in the tumbling and non-tumbling cases. Black, solid curves are analytically approximated. Grey, dashed curves are numerically calculated from SpDust.

and

$$\sigma_s^2 \equiv \frac{\alpha_a^2}{\alpha_\nu^2} \sigma^2 + \sigma_\nu^2. \quad (8.42)$$

The emissivity at  $\nu_0$  is

$$\left. \frac{j_\nu}{n_H} \right|_{\nu_0} = B_1 A_s \frac{\beta^2 N_{\text{at}} \omega_0^3}{6c^3} \quad (8.43)$$

where

$$A_s \equiv \frac{\sigma}{\sigma_s} R_0 \mathcal{I}_0 \exp \left[ -\frac{\sigma_\nu^2}{2} \left( 9 \frac{\alpha_\nu^2}{\alpha_a^2} - 8 \right) \right]. \quad (8.44)$$

The complementary error function provides the high-frequency fall-off with parameters

$$\eta_\nu \equiv \frac{1}{\sqrt{2}} \frac{\alpha_a \sigma}{\alpha_\nu \sigma_\nu \sigma_s} \quad (8.45)$$

and

$$\eta_a \equiv \frac{1}{\sqrt{2}} \frac{\sigma_s}{\sigma_\nu \sigma}. \quad (8.46)$$

Equation 8.39 is the chief result of this chapter. It should be regarded as the natural functional




Figure 8.6: Total spinning dust emissivity for the CNM environment. Solid, black curve is the analytical function of Equation 8.39, while the dashed, grey curve is that from SpDust. Components of the analytical curve are also shown: *A* is the power-law term and *B* is the power-law multiplied by the log-normal distribution. The high-frequency fall-off is provided by the complementary error function, shown in grey using the right *y*-axis.

form for the spinning dust emission. The components of this function are plotted in Figure 8.6, in which the power-law and log-normal components are shown in turn, as is the complementary error function. The analytical curve is plotted alongside the numerically calculated emissivity from SpDust. The agreement is excellent. This is particularly noteworthy as the parameter  $a_2$  has been ignored: the grain geometry is thus of only secondary importance. The same is shown for the rest of the ideal environments in Figure 8.7.

The function contains six parameters, which are heavily degenerate both in derivation and effect. These are the amplitude coefficient  $A_s$ , the characteristic frequency  $\nu_0$ , the power-law slope  $\alpha_s$ , the log-normal width  $\sigma_s$ , and the exponential slopes  $\eta_\nu$  and  $\eta_a$ . The parameters have been calculated in Table 8.3 for the idealized environments for disk-like (1) and spherical (2) grains. The same are plotted in Figure 8.8 for case (1). Also shown are a generic set of parameters recommended for use when, for example, the data are not able to break degeneracies between the parameters. Indeed, the strong correlations between these parameters are clear.

The parameters in Equation 8.39 are not independent. They depend on the excitation and damping power laws ( $\alpha_a$ ,  $\alpha_b$ , and  $\alpha_\nu$ ), the rotational peak  $\Omega_{p,a_0}$ , and the tumbling parameters ( $R_0$ ,




Figure 8.7: Analytical estimates (solid, black) of  $j_\nu/n_H$  compared to that of SpDust (grey, dashed) for idealized interstellar environments.

$q_r$ , and  $\sigma_r$ ), which are themselves dependent on the environment and grain properties. Allowing the parameters of Equation 8.39 to vary independently will complicate physical interpretation. However, despite the ranges of these parameters, the  $j_\nu/n_H$  curves show remarkably little diversity. This can be seen by plotting the analytical curves for the various environments, but with  $\nu_0$  set to some constant value. This is done in Figure 8.9, with  $\nu_0 = 30$  GHz. It is apparent that the choice of generic values for  $\alpha_s$ ,  $\sigma_s$ ,  $\eta_\nu$  and  $\eta_a$  are likely to provide a satisfying fit in any environment.

The presence of the  $a_m$  to  $a_0$  ratio allows probing of the smallest grain size. This requires care, though, as it will be heavily degenerate with  $\nu_0$ , itself depending on environment. If the latter can be constrained independently, then measuring  $a_m$  with this method will provide a new window on grain formation and destruction.

Judging from published data, breaking the degeneracies in this model will be challenging. I suggested that one set  $a_m = a_0$ , use the generic values for  $\alpha_s$ ,  $\sigma_s$ ,  $\eta_\nu$  and  $\eta_a$ , and allow only  $j_\nu/n_H|_{\nu_0}$  and  $\nu_0$  to vary. This is comparable to the analysis of Bennett et al. (2012) and Planck Collaboration (2013d). As data quality improves, varying  $\sigma_s$  may allow an improved fit. If the fit is still unsatisfactory, then fitting the  $\alpha_a$ ,  $\alpha_b$ , and  $\alpha_\nu$  directly may be best, as the higher level




Figure 8.8: Parameters of the analytical  $j_\nu/n_H$  function for the idealized interstellar environments. The suggested generic parameters are shown by the grey, horizontal lines.



Table 8.3: Derived parameters for the analytical emissivity.

Environment		Parameter						
		$\frac{j_\nu}{n_H} \Big _{\nu_0}$ (Jy sr <sup>-1</sup> cm <sup>-2</sup> H <sup>-1</sup> )	$\nu_0$ (GHz)	$A_s$ (10 <sup>-2</sup> )	$\alpha_s$	$\sigma_s$	$\eta_\nu$	$\eta_a$
Generic		...	...	5.00	1.200	0.900	1.50	4.10
CNM	(1)	$7.48 \times 10^{-18}$	33.9	6.40	0.686	0.666	1.32	2.82
	(2)	$7.77 \times 10^{-18}$	31.8	8.08	0.685	0.587	2.26	3.76
DC	(1)	$3.24 \times 10^{-18}$	30.8	3.68	1.892	1.137	1.94	5.79
	(2)	$3.26 \times 10^{-18}$	27.8	5.08	1.599	0.895	2.59	6.07
MC	(1)	$1.53 \times 10^{-17}$	40.7	7.56	1.028	0.777	1.14	2.84
	(2)	$1.56 \times 10^{-17}$	40.7	7.71	0.148	0.533	1.71	2.88
PDR	(1)	$5.23 \times 10^{-16}$	174.2	3.30	2.040	1.297	1.97	6.63
	(2)	$6.98 \times 10^{-16}$	175.0	4.34	1.864	1.087	2.68	7.48
RN	(1)	$7.53 \times 10^{-17}$	86.8	3.83	1.871	1.117	1.96	5.75
	(2)	$9.92 \times 10^{-17}$	84.3	5.52	1.532	0.856	2.65	5.94
WIM	(1)	$6.36 \times 10^{-18}$	32.4	6.25	0.505	0.638	1.27	2.69
	(2)	$6.81 \times 10^{-18}$	30.2	8.26	0.559	0.563	2.24	3.62
WNM	(1)	$4.17 \times 10^{-18}$	28.7	5.85	0.236	0.611	1.17	2.51
	(2)	$4.44 \times 10^{-18}$	26.1	8.32	0.413	0.545	2.11	3.37

(1): Disk-like grains

(2): Spherical grains

parameters are ultimately functions of these. Caution should be exercised, though, as inaccuracies in the model due to approximations may become significant at this point.

The model is demonstrated through comparison to the Perseus Molecular Cloud data of Planck Collaboration (2011) in Figure 8.10. The free-free and thermal dust curves are taken directly from their fit. The spinning dust model is that of this work, with the generic parameters assumed and  $\nu_0 = 30$  GHz. The amplitude of the emission was scaled by 1.15 to improve agreement. The model clearly provides an adequate fit to the data.

## 8.6 Discussion

The analytical derivation presented in this work allows one to understand spinning dust emission intuitively. Emission from a given grain size is spread over a broad peak. The breadth is greater if rotation is not limited by electric dipole damping, but is thermal. The emission is also broadened if rotation is a function of the electric dipole moment, as in the case of plasma drag and electric dipole damping, though not enough to overcome non-thermal rotation. There is a further broadening due to tumbling rotation. Larger grains rotate more slowly, so integrating over grain size leads to a




Figure 8.9:  $j_\nu/n_H$  curves for various environments, but with  $\nu_0 = 30$  GHz, demonstrating the similar curve shapes. The grey, solid curve shows the function using the suggested generic parameters.

gently sloped, low frequency tail. The existence of a smallest grain size leads to an exponential cut-off at high frequency. The log-normal shape is largely due to the log-normal grain size distribution, though it was shown that emission from a given grain size is also well-approximated by this form.

As this derivation involved taking the products many independent functions, a log-normal shape is not surprising. Indeed, simple algebra would allow factoring the power-law component of Equation 8.39 into the log-normal component. A pure log-normal spectrum is not justified, though, as the high frequency cut-off gives asymmetry to the spectrum. This asymmetry is important theoretically, as it contains information on the grain size distribution, and observationally, because anomalous microwave emission measurements from *WMAP* and *Planck* are at frequencies above the peak (Bennett et al. 2012; Planck Collaboration 2013d).

Different interstellar environments lead to different power laws in the distribution function and thus to different combinations of low-frequency slope, log-normal width, and high-frequency fall-off. As seen in Figure 8.9, however, these effects do not lead to large deviations in spectral shape. The characteristic frequency  $\nu_0$  does vary with environment as it is closely related to the grain rotation temperatures. There is degeneracy between  $\nu_0$  and the other parameters of Equation 8.39, implying that the spectral shape ought to change as  $\nu_0$  is shifted. However, the similarity between curves in




Figure 8.10: Spectral energy distribution of the Perseus Molecular Cloud. The data are as published in Planck Collaboration (2011), as are the free-free and thermal dust models. The spinning dust model is from this work, with generic parameters,  $\nu_0 = 30$  GHz, and scaled by a factor of 1.15.

Figure 8.9 suggests that such variations in shape will only become important when the precision of anomalous microwave emission observations has greatly improved.

Measurements of the characteristic frequency  $\nu_0$  will face strong degeneracy with the smallest grain size. Decreasing  $a_m$  gives smaller, faster rotating grains which extend the radiation to higher frequencies. This will present a significant challenge to any attempts to constrain environment or smallest grain size with spinning dust radiation.

A number of approximations were needed in this work. These are the power-law dependencies on  $a$ ,  $b$ , and  $\Omega$ , the log-normal distribution function, and the log-normal spectrum for a tumbling grain. Disagreement with the SpDust model below  $\nu_0$  is mainly due to the first, while the latter two are to blame above  $\nu_0$ . Transient effects due to individual gas collisions and the rotational consequences of triaxial grains and vibrational-rotational coupling were disregarded. These omissions cause inaccuracy above  $\nu_0$  and possible frequency shifts and broadening of the spectrum, respectively.

The caveats of this work extend beyond the analytical approximations. Whether fitting numerical or analytical models, one must bear these in mind. The most important are the following:

**Grain size distribution.** The log-normal form of this distribution is inspired by convenience, not

astrophysics. It is a four parameter model, with parameter values consistent with but not required by infrared and extinction data WD01. These parameters may vary with local conditions, as may the form of the distribution itself. Such variations would be degenerate with variations in rotational excitation and damping.

**Smallest grain size.** Sublimation of small grains is a runaway process (Guhathakurta & Draine 1989). Below a given size, grains will have a very short lifetime. However, a sudden cutoff is not predicted: a range of grains sizes will be undergoing sublimation. Further, sublimation is likely to dehydrogenate the grains before destroying the carbon skeleton, which will undoubtedly affect the electric dipole moments of these grains. Spinning dust models ignore these complications, which may lead to structure above  $\nu_0$ .

**Grain geometry.** A sharp transition between disk-like and spherical grains is unlikely to be physical. Indeed, the very existence of purely disk-like and spherical grains is itself an approximation. The true geometries of will be more complex and could conspire to have grains of different sizes radiating at the same frequency. This would cause structure in the spectrum below  $\nu_0$ .

**Electric dipole moments.** DL98 noted that the permanent electric dipole moments of these grains are extremely uncertain. The random-walk prescription laid out therein and adopted thereafter has allowed for agreement between models and observations, but is not physically motivated. Possible effects can be probed by varying  $\alpha_b$  and  $\beta$ . An increase in the former broadens the spectrum via  $\sigma_s$ , but weakens it via  $A_s$ . An increase in the latter brightens the emission, but boosts electric dipole damping to decrease  $\nu_0$ . A breakdown in the random-walk prescription would introduce  $a$  dependencies in  $\alpha_b$  and  $\beta$ , potentially leading to structure in the spectrum.

**Rotational distribution function.** The Fokker-Planck equation assumes that the damping and excitation mechanisms are stationary processes and that the impulses are small compared to the overall motion. HDL10 showed, however, that impulsive torques lead to a non-thermal tail at high angular momenta. This translates to additional radiation above  $\nu_0$  which the Fokker-Planck approach cannot capture.

These caveats may ultimately limit the precision of spinning dust models, as they can only be resolved through detailed knowledge of grain chemistry. On the other hand, if future observations improve enough to require such precision, then constraining this chemistry directly may become possible.

A new, analytical derivation of the radiation from spinning dust grains has been shown. This work bypasses the lengthy, numerical calculations of previous models while encouraging an intuitive picture of the radiation. Accuracy is not significantly compromised by this approach. Indeed, strong caveats are present in even the numerical models. These approximations are clearly described and

their applicability is demonstrated by comparison with numerical calculations. The final result is a simple, analytical function, well-suited for fitting to astronomical data.

## Chapter 9

# Conclusion

The purpose of this thesis was to make advances in low-frequency foreground separation of CMB experiments through instrumentation, observation, and theory. Diffuse, Galactic emission at these frequencies presents challenges to such efforts, such as the need for separation of the various emission components and the need for high sensitivity. C-BASS is an experiment intended to address the need for data in this area. The experiment's challenges in reaching the required sensitivity motivated the novel data processing algorithms in this thesis. The deep observations of the NCP allowed a practical demonstration of foreground separation techniques, while providing a demonstration of the instrument's performance. Physical separation of synchrotron from spinning dust emission inspired the theoretical aspect of this work.

The C-BASS digital backend is an FPGA-based, digital signal processing platform which provides powerful filtering of the C-BASS time series to remove interfering signals. The limited resources of the digital backend's FPGA required a novel approach to achieve this performance. A chain consisting of demodulation and four low-pass filters was designed, the second filter performing the dual roles of removing periodic phase-switching effects and synchronizing the data rate to the GPS-provided clock. The full performance of the demodulation and filter chain is demonstrated in Figures 3.10, 3.11, and 3.15. The GPS timing required a new algorithm for tracking the arrival of the GPS 1PPS signal and filtering the analog noise on the line, the performance of which is demonstrated in Figure 3.5.

C-BASS required a custom data analysis pipeline. The removal of systematic noise was a key contribution from this thesis. A matched filter with variable frequency and phase was designed for the removal of cold cycle temperature variations, with its performance shown in Figures 5.6 and 5.7. An automated, statistical algorithm for the flagging of RFI was produced. Sample performance for this algorithm is listed in Table 6.1.

A deep, targeted survey of the NCP was performed. The resulting map provided a full test of the C-BASS system and pipeline, including noise characterization and calibration verification. Comparisons to existing data sets showed a marked reduction in systematic noise and intriguing

hints of correlation with other diffuse emission tracers, as shown in Figure 7.17. The potential of the full C-BASS data set was neatly demonstrated.

A new, analytical approximation to the spectrum of spinning dust emission was derived. Each step of the derivation was tested against numerical calculations in order to demonstrate accuracy. The result is a physically-motivated analytical form which, it is hoped, will greatly aid in the spectral separation of low-frequency, cosmological foregrounds and form a framework for the characterization of spinning dust emission. The success of this analytical function in reproducing numerical models is shown in Figures 8.6 and 8.7 for a variety of idealized interstellar environments.

This work represents a significant step in ongoing efforts to characterize the diffuse Galactic radio emission over the whole sky and hence toward understanding the cosmology of the Universe. Foreground separation remains a challenging aspect in this field, and the data set that C-BASS is producing and the theoretical framework for spinning dust emission are crucial steps.

# Bibliography

- Ackermann, M., Ajello, M., Atwood, W. B., et al. 2012, *ApJ*, 750, 3
- Ali-Haïmoud, Y., Hirata, C. M., & Dickinson, C. 2009, *MNRAS*, 395, 1055
- Aller, H. D., & Reynolds, S. P. 1985, *ApJ*, 293, L73
- Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, *A&A*, 61, 99
- Beichman, C. A., Neugebauer, G., Habing, H. J., Clegg, P. E., & Chester, T. J. 1988, *Infrared astronomical satellite (IRAS) catalogs and atlases. Volume 1*
- Bennett, C. L., Larson, D., Weiland, J. L., et al. 2012, *arXiv:1212.5225v2*
- Bersanelli, M., & Mandolesi, N. 2000, *APL&C*, 37, 171
- Bock, J., Church, S., Devlin, M., et al. 2005, *Task Force On Cosmic Microwave Background Research*
- Bonaldi, A., Ricciardi, S., Leach, S., et al. 2007, *MNRAS*, 382, 1791
- Casassus, S., Dickinson, C., Cleary, K., et al. 2008, *MNRAS*, 391, 1075
- Compiègne, M., Verstraete, L., Jones, A., et al. 2011, *A&A*, 525, A103
- Davies, R. D., Dickinson, C., Banday, A. J., et al. 2006, *MNRAS*, 370, 1125
- de Oliveira-Costa, A., Kogut, A., Devlin, M. J., et al. 1997, *ApJ*, 482, L17
- de Oliveira-Costa, A., Tegmark, M., Davies, R. D., et al. 2004, *ApJ*, 606, L89
- de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, *MNRAS*, 388, 247
- Dickinson, C., Davies, R. D., Bronfman, L., et al. 2007, *MNRAS*, 379, 297
- Draine, B. T., & Hensley, B. 2013, *ApJ*, 765, 159
- Draine, B. T., & Lazarian, A. 1998a, *ApJ*, 494, L19
- . 1998b, *ApJ*, 508, 157



- . 1999, *ApJ*, 194, 890
- Draine, B. T., & Li, A. 2007, *ApJ*, 657, 810
- Finkbeiner, D. P. 2004, *ApJ*, 614, 186
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, 524, 867
- Finkbeiner, D. P., Langston, G. I., & Minter, A. H. 2004, *ApJ*, 617, 350
- Finkbeiner, D. P., Schlegel, D. J., Frank, C., & Heiles, C. 2002, *ApJ*, 566, 898
- Fridman, P. A., & Baan, W. A. 2001, *A&A*, 378, 327
- Ghosh, T., Banday, A. J., Jaffe, T., et al. 2012, *MNRAS*, 422, 3617
- Gold, B., Bennett, C. L., Hill, R. S., et al. 2009, *ApJS*, 180, 265
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
- Gregory, P. C., Scott, W. K., Douglas, K., & Condon, J. J. 1996, *ApJS*, 103, 427
- Guhathakurta, P., & Draine, B. T. 1989, *ApJ*, 345, 230
- Hafez, Y. A., Davies, R. D., Davis, R. J., et al. 2008, *MNRAS*, 388, 1775
- Haffner, L. M., Reynolds, R. J., Tufte, S. L., et al. 2003, *ApJS*, 149, 405
- Harten, R. H., Grosbol, P., Greisen, E. W., & Wells, D. C. 1988, *A&AS*, 73, 365
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, 47, 1
- Hinshaw, G., Larson, D., Komatsu, E., et al. 2012, [arXiv:1212.5226v2](https://arxiv.org/abs/1212.5226v2)
- Hoang, T., Draine, B. T., & Lazarian, A. 2010, *ApJ*, 715, 1462
- Hoang, T., Lazarian, A., & Draine, B. T. 2011, *ApJ*, 741, 87
- Hoang, T., Lazarian, A., & Martin, P. G. 2013, [arXiv:1305.0276v1](https://arxiv.org/abs/1305.0276v1)
- Holler, C. M., Taylor, A. C., Jones, M. E., et al. 2013, *IEEE Transactions on Antennas and Propagation*, 61, 117
- Jaffe, T. R., Ferriere, K. M., Banday, A. J., et al. 2013, *MNRAS*, 1
- Jonas, J. L., Baart, E. E., & Nicolson, G. D. 1998, *MNRAS*, 297, 977
- Kogut, A., Banday, A. J., Bennett, C. L., et al. 1996, *ApJ*, 460, 1
- Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, *Nature*, 420, 772

- Kühr, H., Witzel, A., Pauliny-Toth, I. I. K., & Nauber, U. 1981, *A&A*, 45, 367
- Kwok, S., & Zhang, Y. 2011, *Nature*, 479, 80
- Lange, A. E., Ade, P. A., Bock, J. J., et al. 2001, *Phys. Rev. D*, 63, 42001
- Lazarian, A., & Draine, B. T. 2000, *ApJ*, 536, L15
- Leitch, E. M., Readhead, A. C. S., Pearson, T. J., & Myers, S. T. 1997, *ApJ*, 486, L23
- Li, A., & Draine, B. T. 2001, *ApJ*, 554, 778
- Lu, M., Dunkley, J., & Page, L. 2012, *ApJ*, 749, 165
- Macellari, N., Pierpaoli, E., Dickinson, C., & Vaillancourt, J. E. 2011, *MNRAS*, 418, 888
- Mason, B. S., Robishaw, T., Heiles, C., Finkbeiner, D., & Dickinson, C. 2009, *ApJ*, 697, 1187
- Mather, J. C., Cheng, E. S., Eplee, R. E., et al. 1990, *ApJ*, 354, L37
- Meyerdierks, H., Heithausen, A., & Reif, K. 1991, *A&A*, 245, 247
- Miville-Deschênes, M.-A., & Lagache, G. 2005, *ApJS*, 157, 302
- Murphy, E. J., Helou, G., Condon, J. J., et al. 2010, *ApJ*, 709, L108
- Murphy, E. J., Bremseth, J., Mason, B. S., et al. 2012, *ApJ*, 761, 97
- Nita, G. M., & Gary, D. E. 2010, *PASP*, 122, 595
- Ott, M., Witzel, A., Quirrenbach, A., et al. 1994, *A&A*, 284, 331
- Panel on Frequency Allocations and Spectrum Protection for Scientific Uses, Committee on Radio Frequencies, N. 2007, *Handbook of Frequency Allocations and Spectrum Protection for Scientific Uses* (The National Academies Press)
- Parsons, A., Werthimer, D., Backer, D., et al. 2009, arXiv:0904.1181v1
- Peel, M. W., Dickinson, C., Davies, R. D., et al. 2012, *MNRAS*, 424, 2676
- Penzias, A. A., & Wilson, R. W. 1965, *ApJ*, 142, 419
- Planck Collaboration. 2011, *A&A*, 536, 20
- . 2013a, arXiv:1303.5062v1
- . 2013b, arXiv:1303.5072v1
- . 2013c, arXiv:1303.5076v1

—. 2013d, arXiv:1301.5839v1

Rabiner, L. R., & Gold, B. 1975, Theory and application of digital signal processing (Englewood Cliffs, NJ: Prentice-Hall)

Rafikov, R. R. 2006, ApJ, 646, 288

Readhead, A. C. S., Myers, S. T., Pearson, T. J., et al. 2004, Science, 306, 836

Reich, P., & Reich, W. 1986, A&AS, 63, 205

Reich, P., Testori, J. C., & Reich, W. 2001, A&A, 376, 861

Reich, W. 1982, A&AS, 48, 219

Reichold, A., Dawson, M., Green, J., et al. 2006, Proc. of EPAC06, Edinburgh, Scotland, 520

Rubiño-Martín, J. A., López-Caraballo, C. H., Génova-Santos, R., & Rebolo, R. 2012, Advances in Astronomy, 2012, 1

Silsbee, K., Ali-Haïmoud, Y., & Hirata, C. M. 2011, MNRAS, 411, 2750

Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, ApJ, 396, L1

Stevenson, M. A. 2013, ApJ, submitted

Sutton, D., Johnson, B. R., Brown, M. L., et al. 2009, MNRAS, 393, 894

Tegmark, M. 1998, ApJ, 502, 1

Tegmark, M., Eisenstein, D. J., Hu, W., & de Oliveira-Costa, A. 2000, ApJ, 530, 133

Testori, J. C., Reich, P., & Reich, W. 2008, A&A, 484, 733

Tibbs, C. T., Flagey, N., Paladini, R., et al. 2011, MNRAS, 418, 1889

Tibbs, C. T., Paladini, R., Compiègne, M., et al. 2012, ApJ, 754, 94

Van der Ziel, A. 1988, Proceedings of the IEEE, 76, 233

Vinyaikin, E. N. 2007, Astron. Rep., 51, 570

Watson, R. A., Rebolo, R., Rubiño-Martín, J. A., et al. 2005, ApJ, 624, L89

Weiland, J. L., Odegard, N., Hill, R. S., et al. 2011, ApJS, 192, 19

Weingartner, J. C., & Draine, B. T. 2001, ApJ, 548, 296

Weinreb, S., & Schlee, J. 2013, RSI, submitted

Wells, D. C., Greisen, E. W., & Harten, R. H. 1981, A&AS, 44, 363

Wolleben, M., Landecker, T. L., Reich, W., & Wielebinski, R. 2006, A&A, 448, 411

Ysard, N., & Verstraete, L. 2010, A&A, 509, 12