The C-Band All-Sky Survey

Charles Judd Copley
Hertford College
University of Oxford

Submitted in fulfillment of the requirements for the award of a DPhil
at the University of Oxford

· June 22, 2014 ·

Acknowledgements

To those that I spent the most significant amount of time with- the level three crew. Thank you for making my time so enjoyable. From kebabylon, to bushdiving the years are littered with golden moments. Griffin, thank you for your rug. And your pens. They meant a lot to me. It was truly a pleasure to spend three years surrounded by the interminable Denys Wilkinson construction. Jack, I find myself looking up whenever I hear the tinkle of keys. Or the ringing of church bells. I think the only thing I regret of my time at Oxford is that you spent so long chopping wood in France—and that you managed to keep the lead on the chess board. And Danny! Oh Danny. From Halloweeen house paintings, to your escapades with early morning Australian drivers. To bringing up the rear with Essex police guards panting down our necks. And your look of confused bewilderment as we arrived back after the Keyhaven incident. I will miss you all.

To the rest of the non-CBASS crew in astrophysics. To Aris, Ian, Kris, and Sascha. I only wish we'd managed to organise a Greek island getaway. To the radio-lab crew.

To the C-BASS crew. Jamie, for your indominatable spirit and adventurousness. One day you will be able to drive on a dirt-road. Angela for your calm thoughtfulness and wildly varying office temperature during telecons.

A huge thanks must go to a number of people at HartRAO. Pieter Stronkhorst in particular made a large part of this project possible, keeping me on my toes throughout. Also Jacques Grobler, Richard Moralo and Andrew Masiteng who were all instrumental in successfully moving and installing both the C-BASS antennas. Mike Gaylard, Jonathan Quick and Keith Jones all provided a great deal of quietly amused observation.

As well as the HartRAO crew, a number of people at the SKA South Africa need mention. In particular Justin Jonas without whom none of this would have been possible. I am saving myself for the next tequila shot at the Rat! I am also deeply indepted to Bernie Fanaroff, Roy Booth, Tracy Cheetham, Carel van der Merwe, Adrian Tiplady and Dawie Fourie for all their help and advice throughout. Significant thanks need to go to Paul Manners who introduced me to the joys of the Karoo, leading me to numerous adventures. I have never looked upon Losberg with the same naivety again.

On to 26 Western road. The magical kitchen and it's ability to transform even the most average evening into a scene of wit and unabashed merriment. Nicole, I have never been able to look at a dirty kitchen in the same way. The same can be said for a pumpkin pie. Or a bobotie pan. Chris, I am still flummoxed by your ability to both remember every National Geographic magazine article ever written, while simultaneously being the most talented pot-drummer in Oxford. You are a truly talented man. Simon and Kayli. I have never enjoyed a Christmas in Oxford as much as the one we spent together, with Simmo quietly muttering about turbogoose in the corner. Elise. I just don't know what to say. I think Chris's pot-drumming said it all. Jo and Jane. I missed out living in Western road during your glory days but I think I had a taster. Arghya. Burns night truly is amazing!

Elizabeth 'Zaltor' Grant. Thank you for everything. For proof reading my thesis when you had your own to write. And for everything else.

Thank you to my parents who have sacrificed a great deal for me over my life. I hope that I have given back at least some of what you have given me.

Finally I would like to thank my supervisor, Professor Mike Jones. Thank you for all your

patience and your willingness to let me experiment myself without dictating any paths. I have learned a huge amount in the past few years and I am truly very grateful for all your help throughout.

Abstract

The C-Band All-Sky Survey (C-BASS) is a 1 GHz bandwidth survey of the radio sky in both intensity and polarization at a frequency of 5 GHz and with a resolution of 0.8°. Northern and Southern sky coverage is provided by antennas located at the Owen's Valley Radio Observatory (OVRO) in California, and the *Meer*KAT support base in South Africa, respectively.

The primary science goal of C-BASS is to provide a highly sensitive C-Band all sky intensity and polarization map to augment the WMAP/Planck surveys. Removal of foregound contamination will place a limit on the success of Cosmic Microwave Background (CMB) experiments that attempt to detect the B-Mode polarization of the CMB. We will provide a HEALPix map (N_{side}=128) with an r.m.s. noise of 0.13 mK/pixel in Stokes Q and Stokes U, and a confusion limited r.m.s. noise of 0.8 mK/pixel in Stokes I. Removal of foregrounds at the higher frequency CMB surveys will be significantly improved by this lower frequency constraint.

This thesis describes the development of the C-BASS gain-stabilized receiver capable of making sensitive measurements of both galactic total intensity and polarization. The receiver features a novel digital backend to provide spectral detail across the frequency band of interest. The analog signal conditioning uses a double sideband mixer to mix the RF frequencies to a DC–1000 MHz baseband for digitization. By changing the mixer frequency and/ or duplicating the signal conditioning and digital hardware, the RF frequency coverage can be modified for other projects.

I also describe the process of converting a 7.6 m telecommunications dish to a high performance radio astronomy platform. The discussion includes dish surface measurements, optical design, and the development of an inexpensive telescope servo controller. The antenna conversion process and receiver design can be used to significantly reduce capital costs of future experiments, which is especially useful for short timescale experiments. The African VLBI (Nordling, 2012) is currently following a similar route to repurpose antennas across the African continent.

Contents

List of Figures

List of Tables

Glossary

ADC Analog-to-digital convertor.

AGAA the Astronomical Geographic Advantage Area in the Northern Cape Province, South Africa..

AME Anomalous Microwave Emission.

B.U. Backend Units.

BAO Baryon acoustic oscillations.

BORPH Berkeley Operating System for Re-programmable Hardware.

C-BASS C-Band All-Sky Survey.

C.W. Continuous wave signal.

CASPER Collaboration for Astronomy Signal Processing and Electronics Research.

CMB Cosmic Microwave Background.

COBE Cosmic Background Explorer.

DAC Digital-to-analog convertor.

DESCART The DESCART destriping map maker.

DIRBE Diffuse Infrared Background Experiment.

DRAO Dominion Radio Astrophysical Observatory is an observatory located in British Columbia, Canada.

DSP Digital signal processing.

ESR Equivalent series resistance.

FFT Fast Fourier transform.

f_{knee} Frequency at which the $1/f$ noise begins to dominate over the white noise..

FPGA Field programmable gate array.

GPU General Processing Unit.

HartRAO Hartebeeshoek Radio Observatory.

HEALPix Hierarchical Equal Area isoLatitude Pixelisation.

iADC iADC, or ADC2x1000-8.

ILC Internal Linear Combination- a method employed to remove foreground contamination from the CMB measurements.

IRAS Infrared Astronomical Satellite.

k_B the Boltzman constant=$1.38 \times 10^{-23} m^2 kg s^{-2} K^{-1}$.

LNA Low Noise Amplifiers are used to reduce the noise temperature of the receiver $T_{\rm rx}$, thus improving the sensitivity of the instrument..

LOFAR Low-Frequency Array for radio astronomy.

***Meer*KAT** A SKA pathfinder telescope under construction in South Africa.

MLCC Multilayer ceramic capacitor.

OVRO The Owen's Valley Radio Observatory located in Bishop, California..

PAPER Precision Array for Probing the Epoch of Reionization.

PFB Polyphase filter bank.

Planck The Planck Satellite mission.

r.m.s. Root-mean-square.

RFI Radio Frequency Interference.

Ring5m The Ring5m experiment was a ground based experiment at OVRO designed to detect CMB anisotropies on an arcminute scale..

ROACH ROACH Board from CASPER.

S-PASS S-band All Sky Survey.

SRF Self resonant frequency.

$T_{\rm com}$ The effective system temperature of a continuous comparison radiometer when accounting for the additional noise added to the measurement by the reference load..

$T_{\rm rx}$ The noise temperature of the receiver with no external contribution..

$T_{\rm sys}$ The system temperature is the noise temperature of the system when not looking at a source. This includes contributions from the CMB, atmosphere and the receiver. The system temperature is a metric of the instrument sensitivity. $T_{\rm sys} = T_{CMB} + T_{sky} + T_{optics} + T_{\rm rx}$.

WMAP Wilkinson Microwave Anisotropy Probe.

Chapter 1

The C-Band All Sky Survey

> Science is concerned with what is possible while engineering is concerned with choosing, from among the many possible ways, one that meets a number of often poorly stated economic and practical objectives.
>
> Richard Hamming

The C-Band All Sky Survey, or C-BASS, will provide a 5-GHz full-sky map of unprecedented sensitivity in both total intensity, Stokes I, and linear polarization, Stokes Q and Stokes U. The maps will be essential for Galactic foreground removal from CMB polarization surveys and in addition will provide information about large scale magnetic structure within our own Galaxy. The Northern sky survey is currently underway in Owens Valley Radio Observatory, OVRO, in California. The Southern sky survey will be conducted in South Africa with an early commissioning period at Hartebeeshoek Radio Observatory, HartRAO followed by full implementation in the Astronomical Geographic Advantage Area (AGAA) of the Northern Cape Province. This chapter outlines the C-BASS science case, explains the instrument requirements, and describes the instrument design.

1.1 The Cosmic Microwave Background

Soon after the Big Bang, the Universe existed as a hot, dense, ionized plasma. Baryons and photons were tightly coupled because the highly dense plasma decreased the mean free path of the photons to a very short distance. During the Universe's early expansion the ionized plasma cooled enough to recombine, forming neutral Hydrogen. During this period of recombination, the optical depth decreased, allowing photons to decouple from baryons and begin to travel freely through the Universe. Since recombination occurred over a relatively short period of

time, the energy comes from a well defined surface known as the surface of last scattering. This energy is known as the Cosmic Microwave Background, or CMB.

The CMB was serendipitously discovered in 1965 by Arno Penzias and Robert Wilson while working for AT&T's Bell Labs. Penzias and Wilson were testing a highly sensitive horn antenna and were unable to account for an unexpected additional noise spread evenly across the sky. Dicke, Peebles and Wilkinson at Princeton University had concurently reasoned that the Big Bang should have released not only matter, but detectable energy as well.

The two groups heard of each others' work and concluded together that the excess energy observed by Penzias and Wilson was the same energy predicted by Dicke, Peebles and Wilkinson. They published their predictions and observations simultaneously in Dicke et al. (1965) and Penzias & Wilson (1965). The discovery cemented the Big Bang theory as the preferred explanation for the creation of the Universe.

1.1.1 CMB Temperature Anistropies

The CMB has a near perfect blackbody spectrum with a characteristic temperature (see Fixsen (2009)) of 2.72548 ± 0.00057 K. However, it is not perfectly isotropic. Small density variations in the early Universe, seeds of the large-scale structure we see today, imprinted themselves onto the radiation. Known as CMB anisotropies, these small variations in the temperature of the CMB provide vital information about the early stages of the evolution of our Universe.

Their existence was predicted by Sachs & Wolfe (1967). They expected small temperature variations in the CMB caused by gravitational potential variations during recombination. These were first detected by the Cosmic Background Explorer (COBE) DMR experiment (Smoot et al., 1992). The Sachs-Wolfe Effect, as it has come to be known, dominates CMB fluctuation on large angular scales. Its effect on the spatial distribution of CMB anisotropies is shown in Figure 1.1.

The smaller, sub-degree, angular scale anisotropies in Figure 1.1 were detected in rapid succession by a number of experiments like the balloon-born BOOMERANG (de Bernardis et al., 2000) and the MAXIMA-1 (Hanany et al., 2000), the Microwave Anistropy Telescope (MAT) (Miller et al., 1999), and the Cosmic Anisotropy Telescope (CAT) in Cambridge (Scott et al., 1996) to name but a few. The small, angular-scale peak structure in the anisotropies' spatial distribution results from baryon acoustic oscillations, or BAOs: sound waves in the early baryonic matter caused by the competing forces of gravity and radiative pressure. Gravitational perturbations in the pervasive dark matter created dense regions in the primordial plasma that attracted and compressed standard baryonic matter. As the baryonic matter was compressed, the competing photon radiative pressure increased. These two forces set up oscillations and created acoustic waves that propagated out through the coupled plasma. The

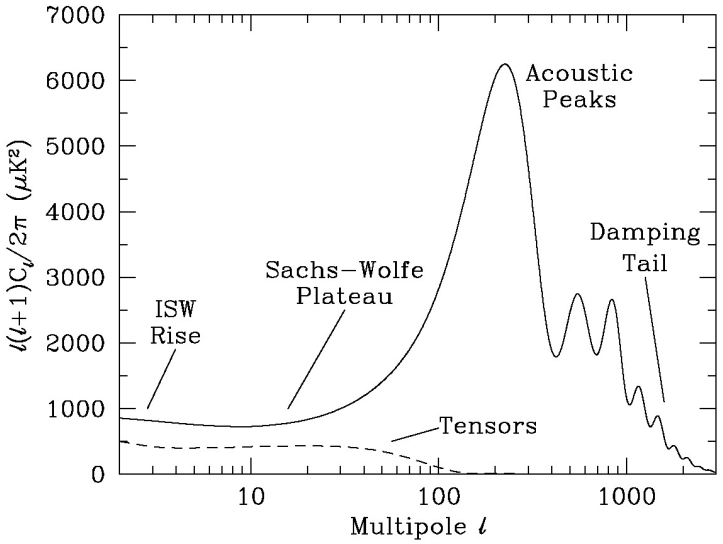

Figure 1.1: The spatial dependence of the CMB temperature anisotropies as a function of the spherical harmonic multipole moment l. The multipole moment, l, is approximately related to the angular scale, θ, of the features by $\theta = \frac{180^\circ}{l}$. (Scott & Smoot, 2004).

resulting compression and rarefaction of the plasma created slightly hotter, overdense regions, and colder, underdense regions. These were imprinted on the CMB at the moment of recombination, creating the acoustic spatial structure.

1.1.2 Polarization Signatures

The polarization of the CMB was first predicted by Rees (1968). Incoming radiation is scattered off of a particle by Thomson scattering, producing polarized radiation. If the incoming radiation is isotropic, the scattered radiation will be equally polarized in all directions, with no net polarization. However, if it is anisotropic, the net scattered radiation will exhibit some degree of polarization. This is illustrated in Figure 1.2. The angular scale and relative strength of the polarization is shown in Figure 1.5.

The polarization state of radiation can be described by the Stokes parameters. However, these are co-ordinate dependent (i.e. they depend on the observer). A more useful description is a co-ordinate independent decomposition of the polarization into two orthogonal components in relation to the direction of travel, known as E-mode and B-mode polarizations. E-mode polarization is symmetric about the direction of travel, while B-mode polarization pattern is chiral, exhibiting handedness. This difference can be seen in the polarization patterns shown in Figure 1.3 and Figure 1.4.

The CMB polarization is caused by quadrupolar anisotropies created by either scalar perturbations, vector perturbations or tensor perturbations. By the period of recombination, vec-

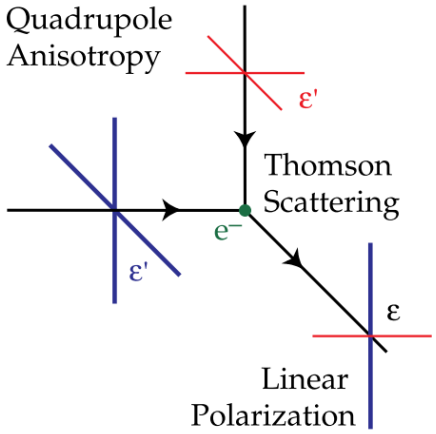

Figure 1.2: How Thomson scattering produces polarized emission (Hu & White, 1997). Quadrupole anisotropies result in slight differences in the Thomson scattering from different directions. The resulting scattered radiation will be polarized.

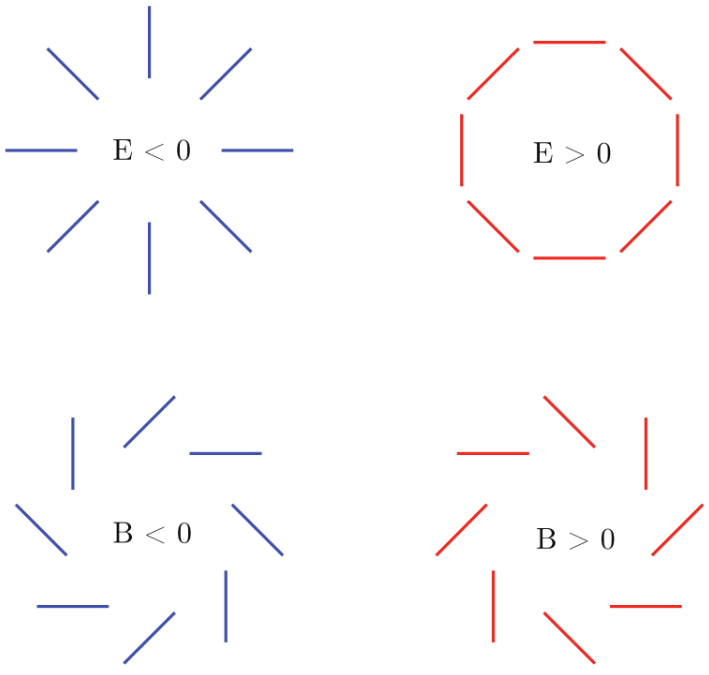

Figure 1.3: Examples of E and B mode polarization patterns with the E-mode often characterized as the curl-free part of the polarization vector field and the B-mode as the divergence-free part. This diagram is from (Baumann et al., 2009).

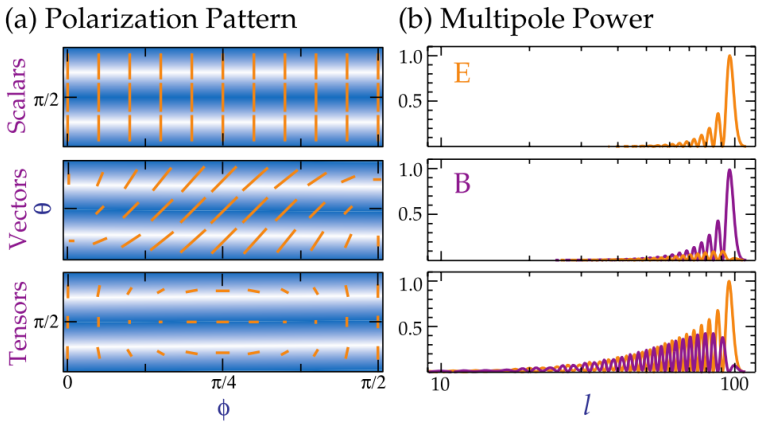

Figure 1.4: Scalar, vector and tensor quadrupolar anistropies and what sort of polarization they are predicted to create. Vector perturbations are predicted to be negligible, so B-mode detection provides information on tensor anisotropies created by gravitational waves (Hu & White, 1997).

tor perturbations are predicted to have been negligible, leaving only scalar perturbations and tensor perturbations. Symmetric scalar perturbations create only E-mode polarization, while asymmetric tensor perturbations create both E-mode and B-mode polarization. This is shown in Figure 1.4 from Hu & White (1997), who provide an excellent primer to the topic of CMB polarization. Measurement of the two polarizations will allow a direct measurement of the gravitational conditions during the period of recombination.

1.2 Galactic Emission Mechanisms

On large angular scales, Galactic emission is dominated by the four mechanisms illustrated in Figure 1.6 (Bennett et al., 2003, Dunkley et al., 2009b, Tegmark et al., 2000, Wilson et al., 2009) namely,

1. Diffuse Galactic synchrotron emission, the dominant mechanism at low frequencies. It is described in Section 1.2.1.

2. Anomalous Microwave Emission (AME), that peaks at around 20 GHz with the peak emission less than the corresponding synchrotron component. The emission falls off rapidly at higher and lower frequencies. AME is described in Section 1.2.2.

3. Thermal Dust emission, the dominant high frequency component of the foreground, that is described in Section 1.2.3.

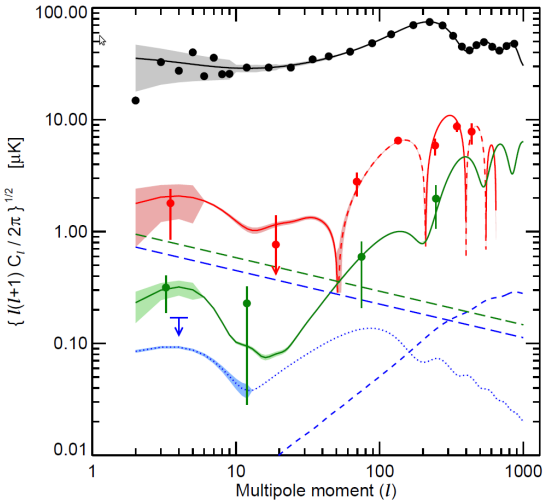

Figure 1.5: The best fit detected strength of TT (in black), EE (in green) and TE (in red) in the 2007 WMAP data release (Page et al., 2007). Also included are predicted levels of BB from the CMB (dotted-blue). The level of the polarized foreground EE signal (synchrotron and Bremmsstralung) is shown as a dashed green line while the corresponding BB mode signal is given by the dashed blue line.

4. Free-free emission, produced by collisions between charged particles. These collissions produce radiation at low frequencies which is typically an order of magnitude lower intensity than synchrotron emission. We explore this in more detail in Section 1.2.4.

1.2.1 Diffuse Galactic Synchrotron Emission

Synchrotron radiation arises as charged electrons and positrons travelling at relativistic speeds spiral around magnetic fields as a result of the Lorentz force. The braking of these charged particles in the magnetic field results in polarized emission. George A. Schott won the Adams Prize at the University of Cambridge in 1908 (subsequently published in Schott (1912)) for an essay that detailed the theoretical existence of this radiation. Despite this early prediction, synchrotron emission itself was only observed 35 years later at the General Electric Synchrotron Accelerator by Elder et al. (1947) — although the observation did not appear to have impressed a young Ronald Reagan, who was doing publicity for General Electric at the time. Astronomical synchrotron radiation was then detected from the astronomical source M87 in 1956 (Burbidge, 1956). For the interested reader, a colourful account of the early history of synchrotron radiation in synchrotron accelerators is given in citeasnounBlewett1998.

Synchrotron radiation is responsible for most of the observed Galactic radiation at low frequencies, as seen in Figure 1.6. Given the polarization of the radiation, measuring the distribution of synchrotron emission is of great importance for CMB polarization studies.

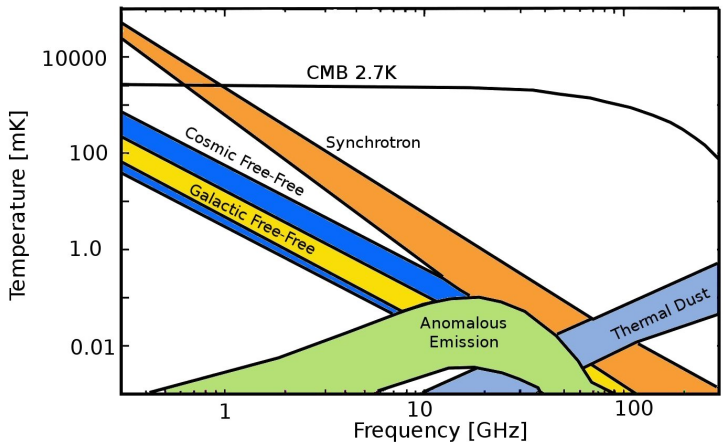

Figure 1.6: The CMB background and the different foreground emission mechanisms. Note the different spectral indices of the the mechanisms as well as their respective contributions at different frequencies.

1.2.2 Anomalous Microwave Emission

Anomalous Microwave Emission, or AME, was first observed at OVRO in the Ring5m experiment (Leitch et al., 1997), a survey that was designed to detect CMB anisotropies on the arcminute angular scale. It consisted of a 14.5 GHz and 30 GHz observations made by the OVRO 40-m dish and 5.5-m telescopes respectively. The survey showed excess emission in certain fields observed for the 14.5 GHz survey that was not predicted by models of thermal dust or synchrotron emission. However, a strong correlation was observed between the distribution of the excess emission, and the Infrared Astronomical Satellite (IRAS) 100μm maps of far-infrared dust emission suggesting that the excess emission was dust-related.

Two mechanisms for the observed radiation were proposed soon after the observation. Draine & Lazarian (1998) explained the excess emission using large numbers of small, optically-excited, dust particles emitting energy by vibration and rotation. Any electric dipole present on the rotating dust particles would also emit electric dipole radiation at a frequency linked to their rotational velocity, estimated to be in the GHz frequency range. This is referred to as the 'spinning dust' model. Draine & Lazarian (1999) proposed another mechanism where a thermally induced magnetic dipole is created by ferromagnetism in dust particles, resulting in magnetic dipole radiation. To create an appreciable dipole, the dust particles need to be much larger than those in the spinning dust model. The two mechanisms could be distinguished from each other by their polarization signature, but it could also be possible that both mechanisms exist in tandem.

The AME spectrum is most significant between 10 and 30 GHz. C-BASS cannot make

direct measurements of the radiation; however, the constraints placed on synchrotron emission will allow better identification of regions of the AME spectrum observed by other instruments.

1.2.3 Thermal Dust Emission

Thermal dust emission is the dominant emission mechanism at frequencies between 100 GHz and 6000 GHz. The emission is caused by heated dust particles ranging between temperatures of 10K and 100K that emit blackbody radiation.

Finkbeiner et al. (1999) used data from the IRAS survey 100μm flux map and the Diffuse Infrared Background Experiment (DIRBE) 100/240μm flux ratio map deployed on COBE to predict full-sky maps of the submillimeter and microwave emission caused by thermal dust.

Thermal dust emission exhibits a degree of polarization caused by asymmetrical dust particles aligned along the Galactic magnetic field.

1.2.4 Free-free Emission

Free-free, or Bremsstralung, emission results when a charged particle is accelerated by another charged particle. This is typical of high energy electrons being deflected by charged atomic nuclei or other electrons present in ionized regions in the interstellar medium.

Synchrotron and thermal dust emission are the dominant Galactic emission at low and high frequencies respectively, making mapping relatively straightforward. Free-free emission is much more difficult to measure, as it is weaker than other mechanisms across all frequencies (Smoot, 1998). The best means of mapping the free-free distribution is by using $H\alpha$ emission as a tracer, since both are related to the same ionized medium.

In the microwave region, Bremmsstralung radiation follows a well-defined power law, $T = \nu^{\alpha}$, with $\alpha = -2.1$. It is expected to be largely unpolarized because the acceleration directions through the ionized medium will be random.

1.3 Science Goals

C-BASS will provide maps of both intensity and polarization at 5 GHz. These maps will aid in constraining the synchrotron emission described in Section 1.2.1, which is essential for any success in CMB polarization studies. Synchrotron emission is also an extremely useful means of probing the Galactic magnetic field and detecting new regions of AME. This section will describe the survey's scientific objectives in more detail.

1.3.1 CMB Polarization Studies

The primary scientific aim of C-BASS is to provide supplementary foreground information to CMB polarization surveys. Detecting the B-mode polarization signature on the CMB depends on extremely sensitive polarization measurements, requiring observations of the CMB across more of the sky than ever before (Carretti, 2010). In the past, areas obscured by foregrounds were typically masked out. However, to detect the 'bump' in the predicted spectrum around $l = 10$, full-sky surveys are required and such masking is not feasible. Detection of the 'bump' at $l = 90$ could be made by smaller angular surveys away from the galactic plane, but even here some level of foreground subtraction will be required to tease out the very low level of the B-mode signal. Higher angular scale B-mode polarization is likely to be caused by gravitational lensing, which maps some of the relatively large E-mode polarization into B-mode polarization. The first detection of this has recently been reported in citeHanson2013.

The importance of accurate foreground models in the Planck data analysis was investigated by Armitage-Caplan et al. (2012). Inflation predicts gravitational waves whose relative amplitude is given by the tensor-to-scalar ratio, *r*. The paper concludes that an estimate of *r* can be biased high by 3σ if incorrect model parameters are used for foreground estimates. Any B-mode polarization surveys will almost certainly be limited by their ability to remove foreground contamination from the surveys (Armitage-Caplan et al., 2011).

1.3.2 Galactic Magnetic Field

The first observations of the Galactic magnetic field were made independently by Hall (1949) and Hiltner (1949), who both described a polarizing effect caused by the interstellar medium on observed starlight. Davis & Greenstein (1951) offered the first plausible explanation for this, suggesting that magnetic fields could align dust particles along the line of sight, polarizing the observed starlight. This provides an indirect mechanism for studying the magnetic field. Other ways were also soon discovered, amongst them the study of synchrotron emission and Faraday rotation of extragalactic sources. A discussion of methods for studying the magnetic field distribution is given in Jansson (2010).

Our understanding of the evolution of Galactic structure relies heavily on evolutionary models of the Galactic magnetic fields (Howes, 2004, Wolleben et al., 2006). These evolutionary models in turn rely on accurate contemporary observations of the Galactic magnetic field.

Low-frequency radio surveys provide a useful probe of the Galactic magnetic field. Studies of the intensity and polarization of the synchrotron emission are a powerful technique of characterising interstellar magnetic conditions. Together with the Rotation Measure (RM)

Survey	Frequency	Sky Coverage	Stokes
Haslam et al. (1981)	408 MHz	All-Sky	I
Testori et al. (2008)	1.4 GHz	South	I,Q,U
Wolleben et al. (2006)	1.4 GHz	North	I,Q,U
Jonas et al. (1998)	2.3 GHz	South	I
S-PASS (Carretti et al., 2013)	2.3 GHz	South	I,Q,U
C-BASS	5.0 GHz	All-Sky	I,Q,U
WMAP (Bennett et al., 2003)	22-90 GHz	All-Sky	I,Q,U
Planck (Planck Collaboration et al., 2013)	30-857 GHz	All-Sky	I,Q,U

Table 1.1: A selection of sky surveys from the low frequency 408 MHz total intensity all-sky survey to the high frequency Planck Surveys. C-BASS (highlighted) provides an important handle on the low frequency synchrotron total intensity emission across the entire sky. The frequency is high enough to avoid significant Faraday rotation of polarized signals.

of pulsars and extragalactic sources, and a suitable electron density model (Cordes & Lazio, 2002), these can provide a deeper understanding of the magnetic conditions present in our Galaxy. Carretti et al. (2009) explain this in greater detail:

1. Total intensity synchrotron emission traces the total magnetic field and is the brightest source of ISM radio emission at low frequencies ($<$10 GHz).

2. Polarized synchrotron emission can be used to study the ordered component of the magnetic field.

3. Rotation Measure of Extragalactic Sources and pulsars provide information of the line of sight to their position. Pulsars located within the Galaxy, are especially useful as (if their distance is known) it is possible to construct a 3D distribution of the magnetic field.

1.3.3 Galactic Total Intensity and Anomolous Emission

A number of other surveys exist at higher and lower frequencies than C-BASS. Properly constraining the low frequency Galactic synchrotron emission requires data between the S-PASS (2.3 GHz) and the WMAP (23 GHz) surveys (see Table 1.1). C-BASS provides exactly these data and will significantly improve our measurements of the Galactic synchrotron distribution.

Free-free emission is well constrained by WMAP and, together with the improved synchrotron intensity model provided by C-BASS, it will be possible to improve the separation of these two components. This will improve our ability to distinguish regions of emission not produced by the synchrotron or free-free mechanisms, such as AME. This is an active field of study, as was discussed in Section 1.2.2.

1.4 Foreground Removal Techniques

Early CMB polarization studies focused on measuring the relatively bright E-mode power spectrum. Small area surveys provided sufficient signal to detect both EE and TE polarization and areas obscured by foregrounds were avoided(Barkats et al., 2005, Barkats et al., 2005, Leitch et al., 2005, Readhead et al., 2004).

With enough data, it should be possible to remove these foregrounds. The CMB produces a nearly perfect blackbody spectrum. However, the foregrounds show strong frequency dependence in their power spectra. We can make use of this information to model the different foreground emission mechanisms and remove their effects from the backgrounds.

A number of different foreground removal techniques exist. An overview of these is provided below to contextualize the important contributions C-BASS data will provide. More detailed information can be found in Dunkley et al. (2009a).

1.4.1 Template Cleaning

Template cleaning is a simple and intuitive method of foreground cleaning. The method assumes that a frequency-dependent map of the microwave sky, $T(p, \nu)$ can be constructed and that we can describe pixel p at frequency ν as a linear sum of fixed spatial templates $X_i(p)$ and noise $n(p, \nu)$ such that

$$T(p, \nu) = \sum_i \alpha_i(\nu) X_i(p) + n(p, \nu). \tag{1.1}$$

The coefficient $\alpha_i(\nu)$ captures the contribution of the foreground template $X_i(p)$ at a given frequency ν. If the noise covariance matrix is known (i.e. the instrument noise and the CMB) then it is possible to solve Equation 1.1 and clean the map by removing the foreground

$$T_c(p, \nu) = T(p, \nu) - \sum_i \alpha_i(\nu) X_i(p). \tag{1.2}$$

The noise of the cleaned map is only slightly affected by the process. While additional templates do increase the resulting noise, the effect is small. Dunkley et al. (2009a) give the example that it would be possible to fit 1000 spatial templates while only increasing the noise amplitude of the cleaned map by 1%.

1.4.2 Blind Component Separation

Blind component separation makes few assumptions about the physical process underlying foreground emission. It instead assumes that the CMB spectrum is thermodynamically uniform across frequencies and that the foreground components are not. The Internal Linear

Combination (ILC) method employed by the WMAP survey (Eriksen et al., 2004) is a well-used example of blind component separation. The ILC technique uses observations of the CMB at multiple frequencies and assumes that what is observed, T_{ν_k}, can be written

$$T_{\nu_k} = T_{CMB} + T_{residual}(\nu_k), \tag{1.3}$$

where the CMB temperature is uniform across frequencies and the foreground component temperatures exhibit frequency dependence.

Breaking the foreground components, $T_{residual}$, into k foreground components with different weightings, w_i, such that

$$T_{residual}(\nu_k) = \sum_{i=0}^{k} w_i T(\nu_k), \tag{1.4}$$

we can rewrite Equation 1.3 as

$$T_{\nu_k} = T_{CMB} + \sum_{i=0}^{k} w_i T(\nu_k). \tag{1.5}$$

The weights, w_i, are then chosen to minimize the final variance of the CMB map, T_{CMB}. This process minimizes the noise contributed by the foreground components. However, it is more difficult to estimate errors in the final CMB map.

1.4.3 Parameter Estimation Methods

If frequency information is available, Bayesian parameter estimation can be used to remove the foregrounds from the CMB. Physical models for the foreground processes are assumed and model-free parameters are fitted using multi-frequency observations. Eriksen et al. (2006) first used this approach, assuming models for the CMB, synchrotron, thermal dust, free-free and AME and applying Markov Chain Monte Carlo (MCMC) modelling to fit the model parameters.

This approach has the advantage that errors in the foreground maps are propagated through to the cleaned CMB map. However, the computational expense associated with Bayesian MCMC is a disadvantage.

Further information on these is also available in Armitage-Caplan et al. (2011), Delabrouille & Cardoso (2009), Dunkley et al. (2009b), Gold et al. (2009), Hansen et al. (2006), Leach & Cardoso (2008).

N_{side}	θ_{pix}	n×C-BASS Solid Angle
1	58.6°	7000
2	29.3°	1750
4	14.7°	440
8	7.44°	110
16	3.66°	27
32	1.83°	7
64	0.91°	1.7
128	0.46°	0.43
256	0.23°	0.1
512	0.11°	0.02

Table 1.2: Angular size, θ_{pix}, of different N_{side} HEALPix pixels (Górski et al., 2005). The C-BASS beamwidth is $\sim$0.7° with a beamsolid angle of 0.15 mSr. We see that the most appropriate HEALPix N_{side} is either N_{side}=128 or N_{side}=256.

1.5 Survey Requirements and Challenges

Having outlined the science case for C-BASS, I now specify the criteria for meeting these objectives: resolution, frequency and sensitivity.

1.5.1 Resolution

In general, the higher the resolution of a survey, the better; however, the scientific goals of C-BASS set a clear minimum resolution requirement. Figure 1.5 shows a distinct B-mode angular power peak at an $l = 90$, corresponding to an angular resolution of $\theta \sim 2°$. This feature needs to be detectable in CMB surveys cleaned using C-BASS data, setting a lower limit to the C-BASS resolution. We require angular sampling at least twice this angular resolution to properly reconstruct the signal, imposing a minimum resolution of 1°. Required dish-size specifications can then be calculated. The resolution, θ_{HPBW}, of a radio telescope is given by

$$\theta_{HPBW} \approx \frac{1.2\lambda}{D}. \tag{1.6}$$

We see from the equation above that achieving the required resolution at the 5 GHz C-BASS frequency requires a dish of at least 4 m diameter.

C-BASS will produce a HEALPix map (Górski et al., 2005). Table 1.2 relates the HEALPix N_{side} with the pixel resolution. For the C-BASS resolution, either N_{side}=128 or N_{side}=256 are most appropriate.

1.5.2 Frequency

C-BASS will constrain the synchrotron component of the low-frequency foreground emission. Synchrotron total-intensity emission increases at lower frequencies, as was seen in Figure 1.6. The highest signal-to-noise measurement of the synchrotron emission would be made at very low frequencies. However, C-BASS aims to measure not only the total-intensity, but also the linear polarization. Polarization measurements are affected by Faraday rotation, which is proportional to the square of the wavelength as shown by

$$\Delta\theta = 420^{\circ} \left\{ \frac{1GHz}{\nu} \right\}^2 \int_0^{L/1kpc} \left\{ \frac{n_e}{0.1cm^{-3}} \right\} \left\{ \frac{B_{||}}{1uG} \right\} \mathrm{d}r. \tag{1.7}$$

Multiple Faraday rotation angles along the line of sight and within the beam lead to depolarization. It is possible to estimate the level of depolarization by comparing surveys made at different frequencies. For example, King (2009) compares the difference in polarization position angle between the 1.4 GHz Dominion Radio Astrophysical Observatory (DRAO) survey and the 23 GHz WMAP survey. The two surveys align well, with an r.m.s. dispersion of approximately 33° attributable to the depolarization in the low-frequency survey. Scaling this angle by λ^2, we estimate a r.m.s. of 2.6° at 5 GHz and can then estimate an r.m.s. error of 9% for Stokes Q and U at 5 GHz.

Given that the total intensity measurements are likely to be confusion-limited anyway (see the next section), a total depolarization of less than 10° seems like a reasonable compromise. This requirement sets the minimum observation frequency at 5 GHz.

1.5.3 Sensitivity

The level of the synchrotron signal at 5 GHz can be estimated by interpolating between lower frequency maps such as the 408 MHz map (Haslam et al., 1981), or the 1.4 GHz DRAO survey (Wolleben et al., 2006), and higher frequency maps like the WMAP 23 GHz survey. The maps and histograms in Figure 1.7 provide an estimate of ∼0.5 mK for the foreground polarization at 5 GHz. For a signal-to-noise requirement of 5:1, the sensitivity should exceed 0.1 mK. The minimum survey sensitivity is therefore set at ∼0.1 mK.

The C-BASS receiver is expected to have a receiver temperature of ∼15 K and will observe across a 1 GHz bandwidth. The sky and atmosphere will contribute approximately 6 K so we can expect a T_{sys} =21K. Using the radiometer equation

$$\sigma_c = \frac{T_{sys}}{\sqrt{\Delta\nu\tau}}, \tag{1.8}$$

we estimate that given T_{sys} =21K and $\Delta\nu$ =1 GHz, we require 20 second integrations to

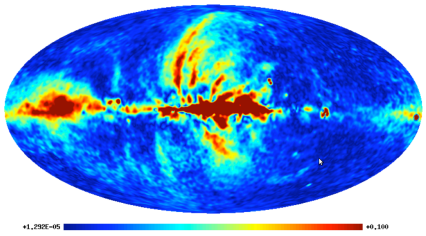

(a) WMAP 23 GHz Polarization Temperature smoothed with a 2° beam.

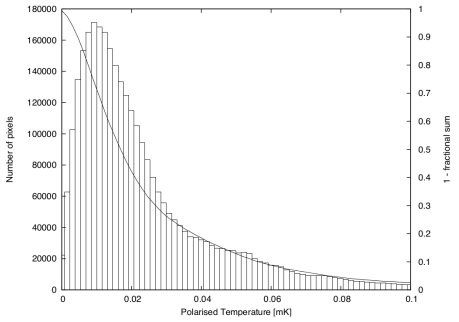

(b) Histogram of WMAP 23 GHz Polarization. 90% of the pixels are greater than 5μK.

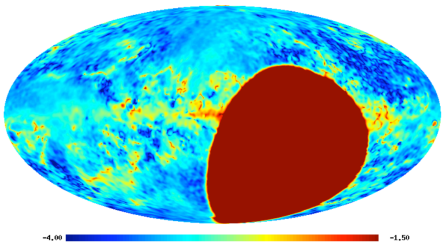

(c) Polarization spectral index across the sky.

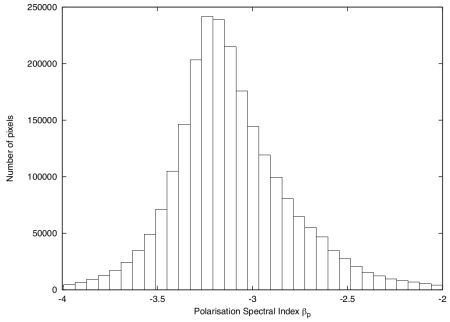

(d) A Histogram of the polarization spectral index.

Figure 1.7: The WMAP polarization map at 23 GHz and the spectral index obtained by comparing this data with the 1.4 GHz DRAO Survey (Wolleben et al., 2006). The comparison suggests a spectral index of $\beta = -3$ between 1.4 GHz and 23 GHz. With a mean value of 5μK/pixel at 23 GHz and $\beta = -3$, we expect an average polarization of 0.5 mK at 5 GHz. To detect this with signal-to-noise of 5:1, we require an r.m.s. noise level of 0.1 mK/pixel in the polarization maps.

achieve 0.1 mK sensitivity in polarization. Using the two independent intensity measurements, we achieve a similar result for intensity; however, the total intensity measurements will be confusion-noise-limited as will be discussed next.

Confusion Limit

Any sources within the antenna beam contribute to the noise level we detect. The larger the beam, the greater the number of sources and the greater their contribution to the background noise. The only way of reducing this background noise is by reducing the field-of-view of the telescope, which would require either increasing the dish diameter or the frequency as per Equation 1.6. A telescope may be more sensitive than the confusion-limit in which case it is confusion limited and the telescope sensitivity is not fully utilized. Sources below the confusion limit will not be observable.

To estimate the confusion noise, σ_c, in the C-BASS total-intensity survey, we can extrapolate from the 1.4 GHz survey in Mitchell & Condon (1985). At the 17.5" resolution of this map, Mitchell and Condon estimate σ_c = 0.010 mJy. At centimeter wavelengths, the flux scales with an average spectral index of β =0.7 (Condon et al., 1998) while the telescope field-of-view scales with the beam solid angle, $\Omega_{beam} \sim \theta^2_{HPBW}$. An estimate of the confusion noise as a function of observation frequency, ν, and beamwidth, θ_{HPBW}, is then given by

$$\sigma_c \propto (\nu)^{-0.7}(\theta_{HPBW})^2 \ . \tag{1.9}$$

Extrapolating from the 1.4 GHz survey, we can estimate the C-BASS total intensity confusion noise as σ_c ∼80 mJy.

The conversion from antenna temperature (in K) to flux (in Jy) is given by

$$S = \frac{2k_B T_{ant}}{\lambda^2}\Omega_{beam} \times 10^{26} Jy. \tag{1.10}$$

where k_B is the Boltzman constant, $\lambda = 6\ cm$ is the observed wavelength, and $\Omega_{beam} \sim 0.01\ sr$ is the C-BASS antenna main beam solid angle in steradians.

Using Equation 1.10, we estimate a σ_c ∼0.9 mK total intensity confusion noise. We cannot reduce this at 5 GHz without altering the primary beam size. Judging from this estimate, we expect to be confusion-limited in our total intensity maps, which would otherwise reach ∼0.1 mK during the survey (see Section 1.5.3). However, the C-BASS polarization measurements are more critical than intensity for CMB polarization surveys. Since most sources are not polarized, we expect a much lower confusion noise for polarization maps. In addition, polarization angles of sources should be randomly distributed and not add coherently within the beam. However, for the sake of certainty let us consider a worst case scenario where

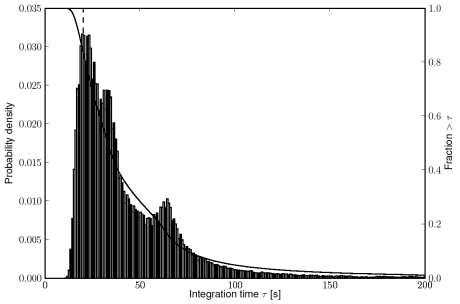

(a) Histogram of pixel integration times, with the minimum integration time of 20 s indicated by a dashed vertical line.

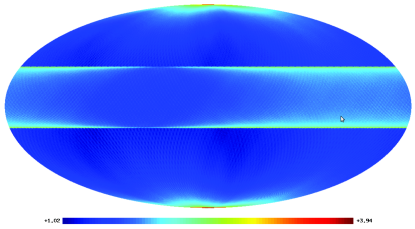

(b) Map of log10(τ) where τ is integration time in an equatorial projection. 85% of the pixels were observed with sufficient sensitivity.

Figure 1.8: The C-BASS double elevation scanning strategy (King, 2009). Over the course of 9 months we scan 85% of the pixels for longer than 20 s. This is the required integration time to achieve the required sensitivity of 0.1 mK per pixel.

$\sim$5% of the sources emit identically polarized radiation at a polarization fraction of $\sim$20%. To estimate the polarization confusion noise, we can scale the total-intensity confusion noise by 0.05×0.2=0.01 or $\sim$1% of the total intensity confusion noise. This is far lower than the 0.1 mK sensitivity requirement estimated in Section 1.5.3. We can heave a collective sigh of relief; the C-BASS polarization maps will not be confusion-limited using the donated antennas.

1.6 Scanning Strategy

Typical scanning strategies in previous ground-based full-sky surveys include: 1) nodding scan and basket-weaving (Haslam et al., 1981), 2) figure-of-eight (Jonas et al., 1998), and 3) drift scans (Wolleben et al., 2006). For C-BASS we used a different approach: scanning the antennas in azimuth at a fixed elevation. A constant reference source is obtained by choosing the elevation to coincide with the North or South Celestial Pole — an approach similar to that used by the S-PASS experiment (Carretti et al., 2010). The ground pickup and atmospheric emission is then (relatively) constant in the data and a horizon map provides a good estimate of what we can expect to see.

The approach was simulated by King (2009) assuming 9 months of observation at an observation efficiency of 20 days a month, 7 hours a night. The simulation results are shown in Figure 1.8 and suggest that 360° azimuth scans at two different elevations (the double-elevation scan strategy) are optimal. This approach provides a minimum of 20 seconds integration time on >85% of the sky pixels — sufficient to reach the 0.1 mK survey sensitivity specified in Section 1.5.3.

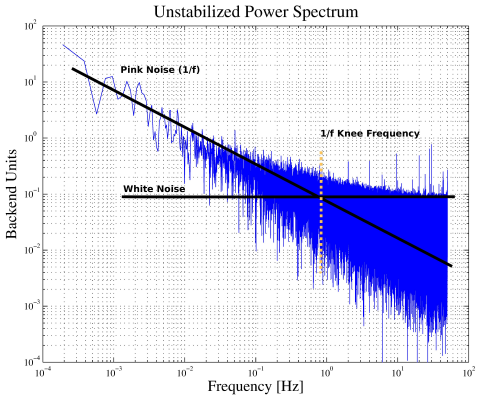

(a) The power spectrum of a typical unstabilized C-BASS LNA.

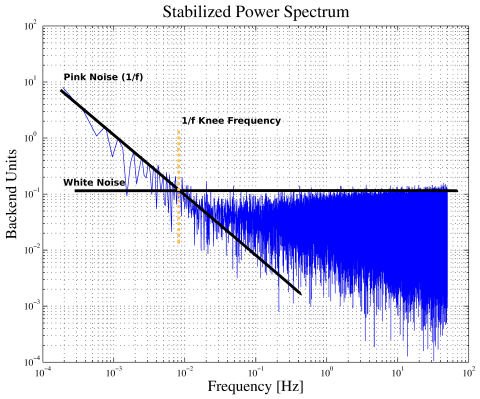

(b) A stabilized intensity measurement using the C-BASS receiver

Figure 1.9: The effect of receiver stability on a traditional dual-polarization total power radiometer design. The gain fluctuations of the LNA increase the noise of long timescale integrations, creating $1/f$ or pink noise. This limits the effectiveness of long integrations. For C-BASS we use the continuous comparison architecture (see Section 1.8.3) to reduce the long term gain fluctuations, and minimize the $1/f$ component.

1.6.1 Receiver Stability

The scanning nature of the survey requires that the instrument gain be stable on timescales longer than the scan period. Stability is important to allow detection of the large angular scale peaks already seen in Figure 1.5. A correlation detector provides wonderful stability because uncorrelated systematic noise from the individual array elements do not affect the signal. Correlation detection is used for the C-BASS polarization measurements, but the intensity measurements are auto-correlation measurements and are therefore susceptible to gain drifts. An interferometer would provide immunity to this; however, because interferometers are not sensitive to angular scales larger than the beam size, they are largely unsuitable for large angular-scale sky surveys.

A traditional single-dish receiver does provide the required large angular-scale sensitivity; however, the noise floor of the instrument is dominated by the systematic noise present on the individual receiver. The long-term gain variations of the low noise amplifiers are particularly troublesome as they produce a pink noise (or $1/f$) power spectrum (Maino et al., 1999, Schmid, 2008, Seiffert et al., 2002), as seen in a C-BASS amplifier power spectrum shown in Figure 1.9a. The $1/f$ noise begins to dominate over the thermal noise on timescales longer that 2 s, limiting the sensitivity improvements obtained by longer integrations. This problem must be resolved to achieve high sensitivity total-intensity measurements.

C-BASS uses a method called continuous-comparison where the sky signal is referenced against a stable source early in the signal chain. Continuous-comparison dramatically im-

proves the data, as shown in Figure 1.9. Further detail is provided in Section 1.8.3, which also shows the advantages of this method over the (possibly) more familiar method of gain stabilization known as Dicke-switching.

1.7 Measuring Total Intensity and Polarization at Radio Frequencies

C-BASS will measure both total intensity (Stokes I) and linear polarization (Stokes Q and Stokes U). In this section, I describe these quantities and explain how an instrument can measure them with maximum sensitivity.

At any instant, a vector describing the electric field of an elecromagnetic wave can be resolved into two components at right angles to each other. For unpolarized radiation, there is no lasting relationship between these two components. For polarized radiation, however, an amplitude and phase relationship does exist and the vector comprised of the two components traces out one of three states: an ellipse, circle or straight line. These states give rise to the terms elliptical, circular and linear polarization respectively.

The Stokes parameters (I, Q, U and V) fully describe the state of polarization. Stokes I characterizes the intensity of the incoming radiation, Stokes Q and Stokes U characterize the state of linear polarization and Stokes V characterizes the state of circular polarization. Measuring these parameters at radio frequencies can be accomplished by correlating orthogonal modes of the incoming radiation.

Tinbergen (1996) provides a review of astronomical polarimetry. The book points out that the Stokes parameters are related to the left (L) and right (R) hand *circular polarization* electric fields by

$$I = LL^* + RR^* \tag{1.11}$$

$$Q = 2\mathbb{R}(LR^*) \tag{1.12}$$

$$U = -2\mathbb{I}(LR^*) \tag{1.13}$$

$$V = LL^* - RR^* \tag{1.14}$$

while for X and Y *linear polarizations*, the parameters are related by

$$I = XX^* + YY^* \tag{1.15}$$
$$Q = XX^* - YY^* \tag{1.16}$$
$$U = 2\mathbb{R}(XY^*) \tag{1.17}$$
$$V = -2\mathbb{I}(XY^*) \tag{1.18}$$

The significance of these equations becomes apparent when we consider the type of polarization we are trying to measure. Since Q, U and V will be small percentages of the total intensity, I, it is preferable to use a form that produces measurements using correlation rather than less sensitive differencing operations. We can see from Equations 1.11 to 1.18 that equipment designed for measuring *linear polarization* (i.e. Stokes Q and Stokes U) should use correlations of circular polarization (Equation 1.12 and 1.13), while equipment designed for measuring *circular polarization* (i.e. Stokes V) should use correlations of linear polarization (Equation 1.18). For C-BASS, we are interested in Galactic polarized emissions that will be dominated by the linearly polarized synchrotron radiation mechanism (Bennett et al., 2003, Reich, 2006). We therefore choose to use correlations between the two circular polarizations, LR^*, to measure Stokes Q and Stokes U.

The auto-correlations, LL^* and RR^*, are just the total power in left and right circular polarizations. These are affected by amplifier gain fluctuations; in the continuous-comparison architecture, we stabilize them by differencing against the two reference temperature loads $E_L E_L^*$ and $E_R E_R^*$, for the left and right polarization measurement respectively. The stabilized measurements are then given by the following relationships:

$$I = (LL^* - E_L E_L^*) + (RR^* - E_R E_R^*) \tag{1.19}$$
$$Q = 2\mathbb{R}(LR^*) \tag{1.20}$$
$$U = -2\mathbb{I}(LR^*)\,. \tag{1.21}$$

These relationships are explored in more detail in Section 1.8.3.

1.8 The C-BASS Instrument

The remainder of this chapter provides a high-level overview of the instrument designed to meet the C-BASS science case. I describe the stabilized C-BASS radiometer and correlation polarimeter and derive the expected system temperature ($T_{\rm sys}$) of the Northern and Southern receivers. I also describe the C-BASS novel optical configuration that features both very low

sidelobes and a radio-transparent foam secondary support. The low sidelobe levels are essential for mitigating the effects of terrestrial pick-up in the far-out sidelobes while the foam support ensures high fidelity polarization measurements by avoiding the scattering of incoming radiation associated with metallic quadropod secondary supports.

1.8.1 The C-BASS Telescopes

In the mid 1990s, two options existed to provide global mobile communications: wireless and satellite. At the time, many believed that the infrastructural costs of establishing wireless infrastructure would restrict their use to densely populated areas. Companies therefore sought to provide mobile communications using satellite technology. However, as the competing wireless infrastructure matured, it became apparent that these networks could provide affordable coverage over wide geographic areas. The market for satellite mobile telephony consequently did not materialize, forcing most newly formed satellite companies into bankruptcy. Much of the ground-based development had already been completed, leaving redundant and identical stations in locations all over the world.

These medium, 8-m-class dishes, capable of efficient operation up to the Ku Band (12 GHz) communications frequencies, meet both the resolution and frequency requirements of C-BASS, described in Section 1.5.1 and Section 1.5.2 respectively. Designing and building new dishes for the survey would cost $\sim$£1.2 M[1]. If an existing antenna can be repurposed for C-BASS, this option should be explored.

If successfully implemented, the C-BASS conversion process has potential to be replicated on a larger scale using redundant, larger dishes located worldwide. Today, almost all of the world's data and voice communications are transmitted via optical fibre. However, in the 1970s, the infrastructural cost of establishing global cable links was daunting. Most voice and data were transmitted using satellite technology. The limited power available on a typical 1970s satellite transmitter required the use of large, 30-m-class ground stations to receive low power satellite transmissions. Today's improved satellite technologies have left these dishes standing largely idle. To the African VLBI network, these offer a possible platform for efficiently deploying a modern, long-baseline radio telescope array (Nordling, 2012). The C-BASS conversion can serve as a useful case-study for this and other similar conversions.

1.8.2 The C-BASS Receiver

The fundamental purpose of a radiometer is to detect radio emission from astronomical sources. Radio emission can usually be approximated by wideband thermal noise characterized by

[1]Personal communication from Willem Esterhuyse *Meer*KAT project manager, 28 April 2009

Gaussian statistics (Wilson et al., 2009), with a voltage probability distribution ($P(v_{source})$) given by

$$P(v_{source}) = \frac{1}{\sigma\sqrt{2\pi}} e^{-v^2/2\sigma_v^2}, \tag{1.22}$$

with:

$$\langle v \rangle = \mu_v = 0 \tag{1.23}$$

$$\langle v^2 \rangle = \sigma_v^2. \tag{1.24}$$

In addition to source emission, the radiometer also detects other sources of thermal noise i.e. the CMB, thermal noise from the atmosphere and ground and thermal noise from the receiver electronics

$$v_{total} = v_{source} + v_{CMB} + v_{atm} + v_{rec}. \tag{1.25}$$

All of these sources can be approximated as Gaussian random variables. We note that the astronomical signal, v_{source}, comprises only a small fraction of the total voltage v_{total} (see Figure 1.11a).

Total Power Radiometer Sensitivity

Most radio receivers do not measure the RF voltage, $v_{total}(t)$, directly. Instead, gain is applied to the signal using an amplifier, the signal is confined to a frequency band $\Delta\nu$ by passing it through a bandpass filter, and this band-limited signal is fed into a square law detector. The radiometer (as the assembly shown in Figure 1.10 is called) produces an output voltage, $P_{total}(t)$, that is proportional to the power of the input voltage

$$P_{total}(t) = a v_{total}(t)^2. \tag{1.26}$$

All of the stages in Figure 1.10 add noise to the measurement of v_{total} by increasing v_{rec}. The amplifier increases the power of the signal of interest, which is necessary to detect very small signals; however, in doing so, it adds electronic noise, adversely affecting the signal-to-noise ratio $S/N = v_{source}/v_{total}$. The signal loss in the filter and detector decreases the S/N ratio by lowering the total RF power while adding thermal noise to the signal.

I now describe techniques used to clarify these concepts. The signal from the radio receiver is band-limited by the bandpass filter allowing us to use the Johnson-Nyquist relationship (Johnson, 1928, Nyquist, 1928). The power in a frequency band, $\Delta\nu$, produced by a resistor at physical temperature, T_{phys}, is given by

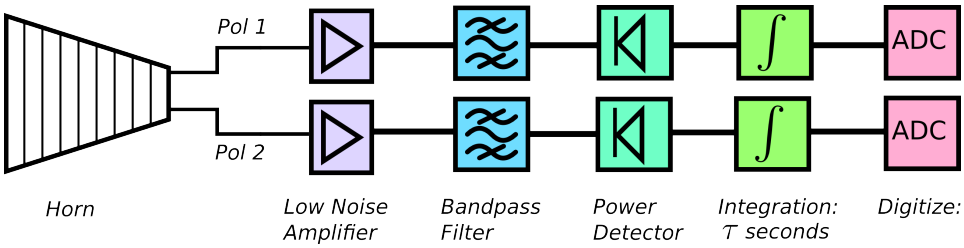

Figure 1.10: A traditional total power radiometer design. This measures the two polarization intensities, and can be used to calculate Stokes I by adding the two arms. Stokes Q can be calculated by differencing the two arms (see Equation 1.16) while measuring Stokes U would require rotating the horn by 45° and differencing the arms. However, this is not an optimal way of measuring polarization for reasons given in Section 1.7

$$P_{noise} = k_B T_{phys} \Delta\nu, \tag{1.27}$$

allowing us to define the concept of a *noise temperature* T_{noise}. This is the equivalent temperature of a resistor at the input that produces a given power spectral density. The noise temperature is then given by

$$T_{noise} = \frac{P_{noise}}{k_B \Delta\nu}. \tag{1.28}$$

Using Equation 1.28, we now rewrite Equation 1.26 to define the *system temperature*, or $T_{\rm sys}$, as the noise temperature of the entire system, such that

$$T_{sys}(t) = a_1 v_{total}(t)^2, \tag{1.29}$$

where only the constant of proportionality has changed. In a radio receiver, we measure the power after the detector. The power we observe (P_{obs}) when observing a source of temperature ΔT_{source} (the Δ indicates that the temperature is small compared to the $T_{\rm sys}$) is the total temperature, $T_{\rm sys}$ +ΔT_{source}, multiplied by the gain and averaged across the frequency band:

$$P_{obs} = P_{sys} + \Delta P = G k_B (T_{sys} + \Delta T_{source}) \Delta\nu. \tag{1.30}$$

The sensitivity of the system is limited by how well we can discriminate ΔT_{source} from $T_{\rm sys}$.

If we assume that the variations in gain are negligible, then our ability to discriminate the small power variations at the detection threshold (where $\Delta T_{source} \ll T_{sys}$) is limited by the noise in our measurement of T_{sys}. For small astronomical signals, lower noise on our measurement of $T_{\rm sys}$ translates to better discrimination of ΔT_{source}. The rest of this section explains how we can calculate the noise on $T_{\rm sys}$ and hence estimate the minimum signal ΔT_{source} we

can detect.

Our measurement of T_{sys} is made by averaging over a period, τ,

$$T_{sys}(\tau) = \frac{1}{\tau}\int_0^\tau T_{sys}dt. \tag{1.31}$$

We estimate the noise using the Central Limit Theorem, relating the measurement noise on a subset, σ_{subset}, to the the number of samples in the subset, N, and the population standard deviation, σ_{pop}:

$$\sigma_{subset} = \frac{\sigma_{pop}}{\sqrt{N}}. \tag{1.32}$$

We then require the population standard deviation of the system temperature (σ_{Tsys}) and the number of samples, N, in each of our measurements of the system temperature itself.

We estimate σ_{tsys} using the relationship between σ_{Tsys} and $T_{\rm sys}$ given by[2]

$$\sigma_{tsys} = \sqrt{2}T_{sys}. \tag{1.33}$$

In order to estimate the number of samples in each of our measurements, we use the Shannon-Nyquist Sampling Theorem (Equation 1.34), which gives us the required sampling rate, f_s, to properly reconstitute an analog signal of bandwidth $\Delta\nu$

$$f_s >= 2\Delta\nu. \tag{1.34}$$

Using Equation 1.34, we can calculate the number of samples over the integration period τ

$$N = \frac{\tau}{1/f_s} \tag{1.35}$$

$$= 2\Delta\nu\tau. \tag{1.36}$$

We can now use the Central Limit Theorem, Equation 1.33 and Equation 1.36 to derive the well known Radiometer Equation. This is given by

$$\sigma_{Tsys} = \frac{\sqrt{2}T_{sys}}{\sqrt{2\Delta\nu\tau}} \tag{1.37}$$

$$= \frac{T_{sys}}{\sqrt{\Delta\nu\tau}}, \tag{1.38}$$

and relates the noise on our measurement of the system temperature to the observed frequency bandwidth, $\Delta\nu$, the integration period τ, and the average system temperature, T_{sys}.

[2]For a full proof see Appendix A

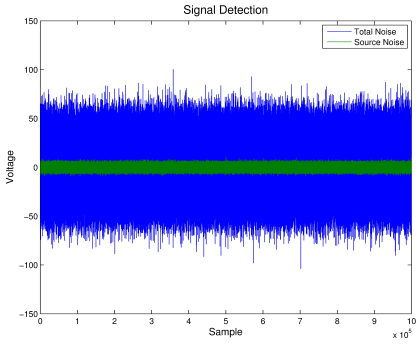

(a) An example of a *voltage* signal (green,σ=2) in the total system noise (blue,σ=20). The *voltage* signals are both Gaussian with zero mean.

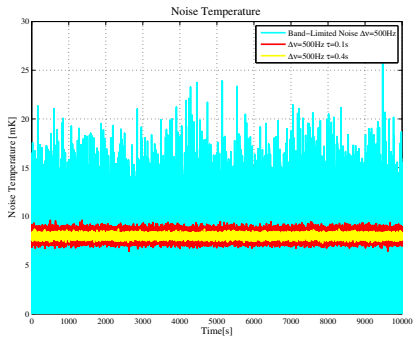

(b) The total detected *power* of the system after the power detection stage. This includes both system noise and astronomical signal. The variance is proportional to the mean of $T_{\rm sys}$.

Figure 1.11: Signal detection in a radiometer. An astronomical signal is buried in the total system noise. The standard deviation of the total system noise, σ_{tsys}, is proportional to the mean (see Equation 1.33), and limits the sensitivity of the detection. To reduce the noise on our measurement of $T_{\rm sys}$, we use the Central Limit Theorem (Equation 1.32) to derive the Radiometer Equation (Equation 1.38). The noise is reduced by either 1)increasing the number of samples per measurement (increasing τ or $\Delta\nu$, or 2)reducing the population standard deviation (decreasing T_{sys}). In this case we show the effect on sensitivity of increasing the integration period τ.

1.8.3 Gain-Stabilized Total-Power Radiometers

The total power detected by a radio receiver of gain, G, bandwidth $\Delta\nu$ and system temperature T_{sys}, is given by $P_{det} = Gk_BT_{sys}\Delta\nu$.

Astronomical signals increase the signal by ΔT_{source}

$$P_{sys} + \Delta P_{source} = Gk_B(T_{sys} + \Delta T_{source})\Delta\nu. \tag{1.39}$$

We estimated the detection threshold of ΔP_{source} in the previous section. However we assumed the signal Gain, G, was constant. In reality, amplifiers exhibit long term gain variations, causing the output power to drift and leading to the relationship

$$P_{sys} + \Delta P_{source} + \Delta P_{gain} = (G + \Delta G)k_B(T_{sys} + \Delta T_{source})\Delta\nu, \tag{1.40}$$

where the variations in detected power ΔP_{source} that we are trying to detect are *indistiguishable* from ΔP_{gain} caused by gain variations. The long time-scale gain variations cause increased low frequency noise following a $1/f$ slope, as seen in Figure 1.9.

Different receiver schemes can be used to mitigate against the effect of these gain drifts. In

the following section we describe the Dicke-switched receiver and the continuous-comparison radiometer employed by C-BASS.

Dicke Switched Receiver

The earliest example of a gain-stabilized receiver architecture was the Dicke-switched radiometer (Dicke, 1946), which is employed on many ground-based telescopes and was also used in the COBE mission (Smoot et al., 1990). The receiver (see Figure 1.12) compares the signal of interest to a load of known power/temperature (either another part of the sky or a stabilized noise diode) by switching between the two signals using an RF switch. Provided the switching period is shorter than the gain fluctuations of the amplifiers, this mechanism significantly improves the long term variations of the system.

For a Dicke-switched system, we switch between the sky and reference so that we only observe the sky for half the time. This has the effect of reducing the integration period τ by a factor of 2 and lowering the sensitivity by a factor of $\frac{1}{\sqrt{2}}$. We also add the noise of the reference load to the measurement, increasing the system temperature to $\sqrt{T_{sys}^2 + T_{ref}^2}$. For the balanced case, where $T_{sys} = T_{ref}$, this increases $T_{\rm sys}$ to $\sqrt{2}T_{sys}$ such that the r.m.s noise is now given by

$$\sigma_{Dicke} = \frac{\sqrt{2}T_{sys}}{\sqrt{\frac{1}{2}\Delta\nu\tau}} \tag{1.41}$$

$$= \frac{2T_{sys}}{\sqrt{\Delta\nu\tau}}. \tag{1.42}$$

The Dicke-switched system is less sensitive than an unstabilized radiometer by a factor of two. If we make two independent observations i.e. measurements in both left and right circular polarization (as we do in C-BASS), then the total power sensitivity is improved by a factor of $\sqrt{2}$.

Continuous-Comparison radiometer

Continuous-comparison architecture uses a stabilized temperature load as a reference. The reference load voltage is combined with the sky voltage using a 180° hybrid, producing two new voltages that differ by 180° in the phase of the load voltage. After the gain stages, the sky voltage and load voltage are separated and the power detected. Since both voltages have been through the same gain chain, we can remove the effect of any gain variations by comparing the detected sky power against the detected load power.

The continuous-comparison radiometer observes the sky continuously, an improvement

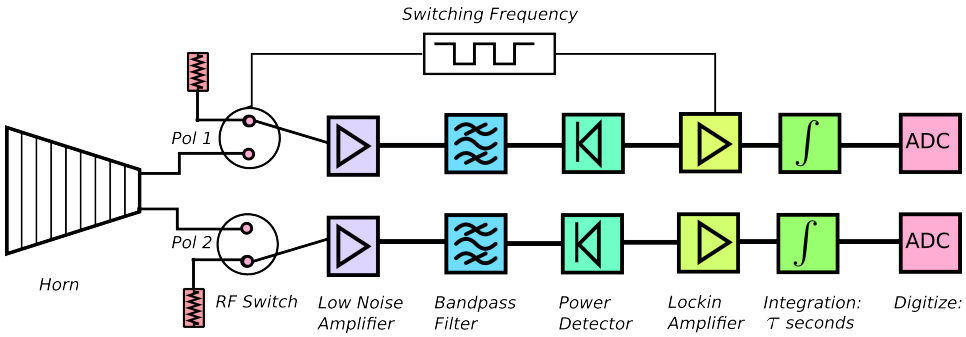

(a) The Dicke Switched radiometer alternates observations between a stabilized reference load and the source of interest

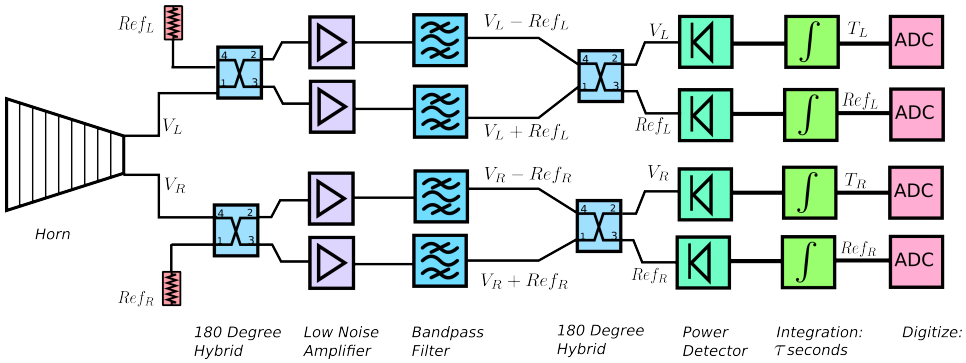

(b) The continuous-comparison architecture where a reference is added both in and out of phase. We use the phase information to separate the reference from the sky after the gain stages. Both reference and sky have been through the same gain stages, and comparing one against the other reduces the effect of any gain fluctuations.

Figure 1.12: Gain stabilized radiometer designs.

over the Dicke switched system. However, it still compares the cold load against the sky, and therefore increases the effective measurement $T_{\rm sys}$ (expressed here as $T_{\rm com}$) in the same way as a Dicke-switched system

$$T_{com} = \sqrt{(T_{sky} + T_{rxSky})^2 + (T_{load} + T_{rxLoad})^2}. \tag{1.43}$$

where the receiver noise added to the sky signal may be different to the receiver noise added to the load signal.

If we assume that sky and cold load are balanced and that the added receiver noise is the same for both the sky and load channels, this reduces to

$$\begin{aligned} T_{com} &= \sqrt{2(T_{sky} + T_{rx})^2}, & (1.44)\\ &= \sqrt{2}T_{sys} & (1.45) \end{aligned}$$

and the intensity r.m.s. noise is $\sqrt{2}$ times better than the Dicke switched system of the previous section described by Equation 1.42.

$$\sigma_{com} = \frac{\sqrt{2}T_{sys}}{\sqrt{\Delta\nu\tau}}. \tag{1.46}$$

For C-BASS, we make independent measurements of total intensity in left and right polarization allowing us to reduce σ_{com} by a factor of $\sqrt{2}$ for a final expected r.m.s. of

$$\sigma_{com} = \frac{T_{sys}}{\sqrt{\Delta\nu\tau}}. \tag{1.47}$$

1.8.4 Correlation Polarimeters

A correlation polarimeter measures polarization by correlating left and right circular polarization, as described in Section 1.7. Since the measurement is made by correlation, it does not require gain stabilization; however, the polarization sensitivity may be affected by the choice of total-intensity gain stabilization architecture.

In a receiver that uses Dicke-switching to stabilize the total-intensity measurement, the sky is only observed half of the time. This will also reduce the polarization sensitivity by $\sqrt{2}$. Continuous-comparison architecture offers a significant advantage since the sky is observed continuously, allowing us to achieve the best sensitivity possible for polarization measurements

$$\sigma_{pol} = \frac{T_{sys}}{\sqrt{\Delta\nu\tau}}. \tag{1.48}$$

The noise of the polarization measurement (σ_{pol}) is thus $\sqrt{2}$ times better than it would be if we used a Dicke-switched receiver.

1.8.5 System Temperature and the Friis Noise Formula

It can be convenient to separate $T_{\rm sys}$ into two components: the sky temperature, T_{sky}, and the receiver temperature, T_{rx},

$$T_{sys} \sim T_{sky} + T_{rx} \tag{1.49}$$

The receiver temperature is the noise added to the signal from the receiver electronics and can be calculated by the Friis Cascaded Noise Equation

$$T_{rx} \approx T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + \frac{T_4}{G_1 G_2 G_3} + ... \frac{T_i}{\prod_{k=1}^{k=i-1} G_k}. \tag{1.50}$$

For C-BASS this becomes

$$T_{rx} \approx T_{horn} + \frac{T_{passives}}{G_{horn}} + \frac{T_{LNA}}{G_{horn} G_{passives}} + ... \tag{1.51}$$

where the noise temperatures of passive devices $T_{passives}$ (such as the OMT, L2C, coupler, isolator etc.) can be approximated by the noise contributed by resistive losses

$$T_{passives} = T_{physical}(1 - G_{passives}) \tag{1.52}$$

while the appropriate temperatures of active devices such as LNAs are given by their noise temperature.

We can use Equation 1.50 to estimate the expected receiver temperature for C-BASS. These will be different for the Northern cryostat (which uses a 4 K Sumitomo coldhead) than for the Southern cryostat (which uses a 10 K Oxford cryosystems coldhead). For the Northern system, we expect the receiver temperature to cascade as shown in Table 1.3, with the final $T_{\rm rx}$ =11.5 K. For the Southern system, we expect $T_{\rm rx}$ =15.9 K with the noise cascade shown in Table 1.4.

To estimate the system temperature for both, we use Equation 1.49 and take into account the sky temperature, T_{sky} (nominally 6 K at 5 GHz) and optical contributions. These are expected to be different for the Northern and Southern antennas, given their different optical configurations (see Figure 1.13). The smaller Northern antenna uses 300 K eccosorb baffles to reduce far out sidelobes while the Southern antenna achieves the same effect by under-illuminating the dish. As a result, we expect the optical contributions on the Northern survey

Component	T_{phys}	Gain[dB]	G[Lin]	$\prod G$	T_{rx}[K]
Horn	300±1 K	-0.010±0.01	1.00±0.01	1.00	0.69±0.01
OMT	5±0.1 K	-0.50±0.05	0.90±0.01	0.89	1.23±0.02
L2C	5±0.1 K	-0.50±0.05	0.90±0.01	0.89	1.84±0.04
Coupler	5±0.1 K	-0.05±0.05	0.99±0.010	0.80	1.91±0.04
Hybrid	5±0.1 K	-0.7±0.1	0.85±0.02	0.79	2.86±0.07
Isolator	5±0.1 K	-0.7±0.1	0.85±0.02	0.67	3.98±0.10
LNA	T_{amp}=3.6 K±1.0 K	40±0.1	10000±230	0.57	13.32±1.10
Cable	40±1 K	-1.2±0.1	0.75±0.02	5735	13.79±1.56

Table 1.3: The Northern receiver (with a ∼5 K cryostat) cascaded noise temperature using the 3 K Low Noise Factory LNAs.

Component	T_{phys}	Gain[dB]	G[Lin]	$\prod G$	T_{rx}[K]
Horn	300±1 K	-0.010±0.01	1.00±0.01	1.00	0.69±0.01
OMT	11±0.1 K	-0.50±0.05	0.90±0.01	0.89	1.88±0.03
L2C	11±0.1 K	-0.50±0.05	0.90±0.01	0.89	3.23±0.06
Coupler	11±0.1 K	-0.05±0.05	0.99±0.010	0.80	3.39±0.06
Hybrid	11±0.1 K	-0.7±0.1	0.85±0.02	0.79	5.48±0.11
Isolator	11±0.1 K	-0.7±0.1	0.85±0.02	0.67	7.93±0.16
LNA	T_{amp}=3.6 K±1.0 K	40±0.1	10000±230	0.57	17.28±1.16
Cable	40±1 K	-1.2±0.1	0.75±0.02	5735	17.74±1.62

Table 1.4: The Southern receiver (with a ∼11 K cryostat) cascaded noise temperature using the 3 K Low Noise Factory LNAs.

to be larger than for the Southern survey. Simulations suggest that the Northern survey optical contributions will be 3 K higher than the Southern: North(T_{sys}) =11.5 K+6 K+3 K=20.5 K and South(T_{sys}) =15.9 K+6 K=22 K. For simplicity we assume that both will have $T_{\rm sys}$ =22K

In Section 1.8.3 and Section 1.8.4, I described how to estimate the r.m.s. of the noise in both total intensity and polarization respectively. In Section 1.6, I also described the double elevation scanning strategy that provides 20 s integrations on 85% of the sky pixels over the course of 9 months.

Using $T_{sys} = 22$ K, $\Delta\nu = 1$ GHz, we expect a noise level of 6.95 mK for a typical $\tau = 0.01$ s integration. This is reduced to 0.15 mK on each pixel if we are able to observe for a total $\tau = 20$ s. This is applicable for the polarization measurements. In the case of the intensity measurements we need to modify $T_{\rm sys}$ to take into account the stabilization signal provided by the cold load (see Equation 1.46). In this case we use $T_{com} = \sqrt{2}T_{sys}$=30.8 K and we expect a noise level of 9.84 mK per integration in each polarization. In the final map we will combine both left and right polarization intensity measurements and gain a factor of $\sqrt(2)$ for a final intensity map sensitivity of 0.15 mK per pixel.

1.8.6 The C-BASS Beam Shape

The final C-BASS survey will be a combination of both Northern and Southern surveys. Combining the two surveys will be easier if the two telescope beams are similar.

Since the two primary reflectors used in the C-BASS were different (having been donated by JPL for the Northern survey and by Telkom for the Southern survey), no control over their design was possible. Designing different horns for the two antennas would have been prohibitively expensive. In order to keep the Northern and Southern beam shapes similar, we designed different secondary mirrors and also used different antenna optics for the two telescopes (Gregorian and Cassegrain for the Northern and Southern telescopes respectively). Photographs and diagrams of the differences between the two telescopes are shown in Figure 1.13. Simulations of the two antenna beam shapes are given in Holler et al. (2013) and later discussed in Section 3.2.3.

Antenna Directivity, Gain and Effective Area

The directivity of an antenna, D, is a measure of the ratio of the beam solid angle, Ω_A, to the solid angle of a sphere; a perfectly isotropic antenna would have a directivy of 1. Given the beam pattern, $U(\theta, \phi)$, the beam solid angle is given by (Wilson et al., 2009):

$$\Omega_A = \int_{\phi=0}^{\phi=2\pi} \int_{\phi=0}^{\phi=\pi} U(\theta, \phi) sin(\theta) d\theta d\phi \tag{1.53}$$

(a) The 6.1 m Northern Survey Dish

(b) The 7.6 m Southern Survey Dish without secondary optics

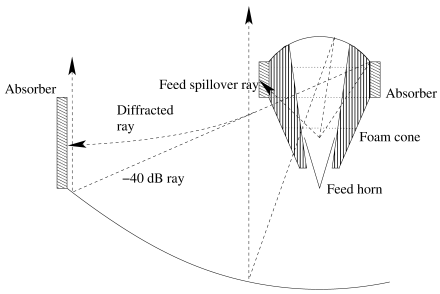

(c) Optics of the Northern antenna

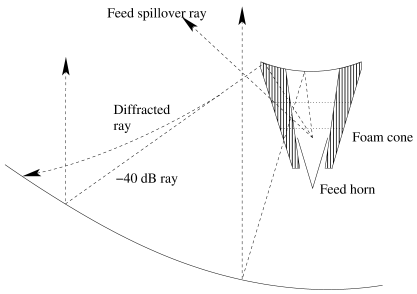

(d) Optics of the Southern antenna

Figure 1.13: The two dishes used for the Northern and Southern C-BASS surveys and their associated optical configurations. A Gregorian optical configuration was used for the Northern survey while a Cassegrain optical configuration was used for the Southern. By designing different secondary reflectors, we were able to use the same low-sidelobe horn for both surveys. Absorbing baffles were also installed on the Northern antenna. These minimize the sidelobes at 50° and 110° (Holler et al., 2013). The 7.6 m Southern antenna is under-illuminated and any spillover is reflected back onto the sky. Absorbing baffles are therefore not required for the Southern antenna.

	Horn	Antenna
G[dBi]	17.0$\pm$0.1	47.4$\pm$0.1
A_e (m^2)	0.0146$\pm$0.0003	15.7$\pm$0.4

Table 1.5: The gain and effective area of the horn and full C-BASS antenna

The $sin(\theta)$ term in the integration corrects for the projection effect one observes, namely that in a spherical coordinate system, the length of an arc subtended by the angle ϕ, depends on the elevation, θ of the arc.

Since the solid angle of a sphere is 4π, it follows that the directivity, D is given by:

$$D = \frac{4\pi}{\Omega_A} \tag{1.54}$$

The directivity of an antenna is very similar to the gain, G. In fact, if we neglect antenna losses, the gain of the antenna and directivity are the same and $G = D$.

The effective area of an antenna, A_e is the area over which the antenna collects power and is used to predict the power that will be detected by an antenna from a source of known flux.

The effective area depends on the gain of the antenna, G, and the wavelength of observation, λ, so the relationship is given by,

$$A_e = \frac{\lambda^2}{4\pi} G. \tag{1.55}$$

We remind ourselves that the the gain, G, and directivity, D, are functionally equivalent.

For C-BASS we can use the simulated beam patterns for the horn and for the entire antenna to measure G for both of these cases. Using a wavelength of 6 cm for C-BASS we arrive at the values in Table 1.5,

The effective area of the horn or antenna (A_e) can be used to estimate the antenna temperature due to a source of flux (S_ν) using (Wilson et al., 2009)

$$T_{ant} = S_\nu A_e / 2k_b \tag{1.56}$$

Antenna Sidelobes

Bright non-astronomical sources can contaminate the signal we detect. The ground with a temperature of 300 K, is challenging. Since the temperature is significantly brighter than the $\sim$6 K sky, it is easily picked up in the antenna sidelobes. This is even more problematic in polarization because the strongly polarized ground reflections can completely drown the small polarized signals we wish to detect. To prevent this, C-BASS is designed with $>$60 dB rejec-

tion in the far out sidelobes (Holler et al., 2013). These low sidelobes reduce the contamination from the ground significantly.

The ground spillover can be calculated by integrating the ground temperature over the solid angle subtended by the sidelobes. This is:

$$
\begin{aligned}
T_{groundspill} &= \frac{\int_{\phi=0}^{\phi=2\pi}\int_{\phi=0}^{\phi=\pi} U(\theta,\phi)\,T(\theta,\phi) sin\theta\; d\theta d\phi}{\Omega_{mainbeam}} && (1.57)\\
&\approx \frac{T_{ground}\,\Omega_{sidelobe}\,U(\theta_{sidelobe})}{\Omega_{mainbeam}} && (1.58)\\
&\approx \frac{300\ K \times 300\ \mathrm{sq.deg} \times 10^{-6}}{0.5\ \mathrm{sq.deg}} && (1.59)\\
&\approx 200\ mK. && (1.60)
\end{aligned}
$$

With the 60 dB rejection, the contribution from the ground in the sidelobes is about 200 mK. This varies due to the topography which is especially important at the OVRO site which is surrounded by the Sierra Nevada mountains. We are currently developing a set of ground templates to remove this contribution from the intensity and polarization measurements by modelling the ground spillover as a function of azimuth.

1.8.7 Secondary support

The design of the secondary support is critical for C-BASS polarization measurements. Any mechanical supports in the path of the beam will scatter the incoming radiation, negatively impacting the quality of the polarization measurements. This was indentified as a key design challenge early in the C-BASS project. Our solution for both the Northern and Southern antennas was to use radio-transparent foam to support the secondary mirrors directly off the horn. The foam is fashioned into a circularly symmetric structures and the mirror and mounting structure are epoxied onto the foam. The entire secondary structure is then bolted onto the horn directly and adjusted into focus using height adjusting screws. This is described in Holler et al. (2013).

1.9 Conclusion

Understanding the low frequency galactic emission mechanisms — especially the dominant synchrotron component — can only be improved by all-sky total-intensity and polarization surveys at appropriate frequencies.

In this chapter I have introduced C-BASS, a highly sensitive, gain-stabilized galactic foreground survey that will produce 5 GHz HEALPix maps in both total intensity and linear po-

larization with an r.m.s. noise of 0.1 mK/pixel. The frequency was chosen as a compromise between the competing requirements of high signal-to-noise measurements of synchrotron emission, which increases at lower frequencies, and high quality polarization measurements that are adversely effect by depolarization due to Faraday rotation, which increases with wavelength.

C-BASS and the resulting total-intensity and polarization maps will be essential for measuring the low-level B-mode polarization signature of the CMB. These B-mode surveys will almost certainly be limited by the difficulties of effective polarized galactic foreground removal, something a low-frequency polarization survey like C-BASS is absolutely essential for. The predicted first peak in the B-mode angular power spectra at l =100 places a minimum 1° angular resolution requirement on C-BASS. Fortunate secondary scientific results of the C-BASS survey will be significant improvements in our understanding of galactic magnetic structure and a useful low frequency constraint to aid in identifying new regions of AME.

I concluded the chapter with a high-level description of the C-BASS instrument and introduced the novel design features used to meet the scientific objectives.

In the following chapter, I discuss the first iteration of the receiver design for C-BASS. This receiver was a broadband analog system which makes stabilized total-intensity measurements using a dual polarization continuous-comparison design and linear polarization measurements using a correlation polarimeter. The receiver is currently being used for the Northern survey and deployed on the telescope at the OVRO C-BASS site. The experience gained during the analog receiver's design, manufacture, integration and commissioning proved instrumental during the design and the manufacture of the Southern receiver, which will form the primary focus of this thesis.

Chapter 2

The Northern C-BASS System

> Science, my lad, is made up of mistakes, but they are mistakes which it is useful to make, because they lead little by little to the truth.
>
> Jules Verne

The C-Band All Sky Survey began with the development and deployment of an analog radiometer/polarimeter at the Owens Valley Radio Observatory in November 2009. The design and deployment are described in detail in the excellent D.Phil thesis King (2009). The Northern survey is nearing completion and is expected to meet the stated survey goals, but the intervening years have taught us a number of lessons.

This chapter will describe the analog receiver in detail and will elaborate on a few of the technical challenges we faced during the commissioning and survey operations. The digital C-BASS receiver, which is the primary subject of this thesis, is an improvement to the analog receiver and is best described with reference to this, our first attempt at building an instrument capable of delivering on the C-BASS scientific objectives.

2.1 The Analog Radiometer/Polarimeter

Unlike the digital Southern C-BASS receiver, the Northern receiver used an analog system because suitable digital alternatives were not readily available at the time of design.

In practical terms, deploying a broadband analog receiver is very difficult, especially in ground-based all-sky surveys. The instrument's sensitivity across broad frequency ranges is a significant weakness in the face of artificial radio frequencies, both terrestrial and satellite borne.

The analog receiver system diagram is shown in Figure 2.1. The system uses the continuous-comparison architecture described in Section 1.8.3 to stabilize and measure Stokes I.

The linear-to-circular convertor, or L2C, converts the linear polarized signals from the orthomode transducer, or OMT, to circular polarization. Each signal goes through a 180° hybrid, coupling in the reference load voltage used for the continuous-comparison stabilization. The hybrid outputs are then amplified by the LNAs. The signal is split into the radiometer and polarimeter sections.

The radiometer signals are phase-switched and passed through a second 180° hybrid, producing a signal that alternates between sky and load according to the state of the phase switch. The phase-switching gives immunity to noise introduced in the post-detection diode electronics, and differencing the detected sky and load powers provides immunity to gain drifts.

The polarimeter signals are also passed through a 180° hybrid to separate the sky voltage from the load voltage. For polarization measurementss, which require no stabilization, the load voltage is simply dumped into a 50 ohm termination while the sky left and right circular polarizations are cross-correlated to detect Stokes Q and Stokes U (as described in Section 1.7). Phase switching again provides immunity from post-diode detection noise.

The RF signal is converted to a power measurement using detector diodes, and the voltage is then sampled using a FPGA-based digital-acquisition card with a 2 MHz bandwidth. Phase switching is also implemented to provide immunity to gain fluctuations and noise present in the signal path after the detector diodes. More detail can be found in King (2009).

2.2 Receiver Temperature and Low Noise Amplifiers

We originally assembled the Northern receiver using 4–8 GHz LNAs supplied by Jodrell Bank. These amplifiers had a noise temperature of 8 K and 27 dB gain.

As the Northern survey progressed, it became apparent that the noise temperature of these amplifiers was too high to meet the required C-BASS survey sensitivity. Since the continuous-comparison architecture of the C-BASS receiver requires multiple lossy components prior to the first stage of amplification, the amplifier noise temperature is critical, as evidenced by the Friis Equation introduced in Equation 1.50. As a result we decided to switch to the 3 K $LNF-LNC4\_8A$ LNAs from the Low Noise Factory.

I estimated the expected $T_{\rm rx}$ =13.79±1.6 K using the 3 K $LNF-LNC4\_8A$ amplifiers in Table 1.3. The same process is applied in Table 2.1 using the 8 K Jodrell Amplifiers. The estimated receiver temperature of $T_{\rm rx}$ =26.15±1.56K is 13 K higher than the receiver using the $LNF-LNC4\_8A$.

In Section 1.8.5 I estimated an r.m.s. noise of $\sigma = 0.15$ mK/pixel using the double elevation strategy for 9 months of observations with the Low Noise Factory amplifiers. Given the values of $T_{\rm rx}$ estimated in this section and a sky temperature of ∼6 K, we would have required

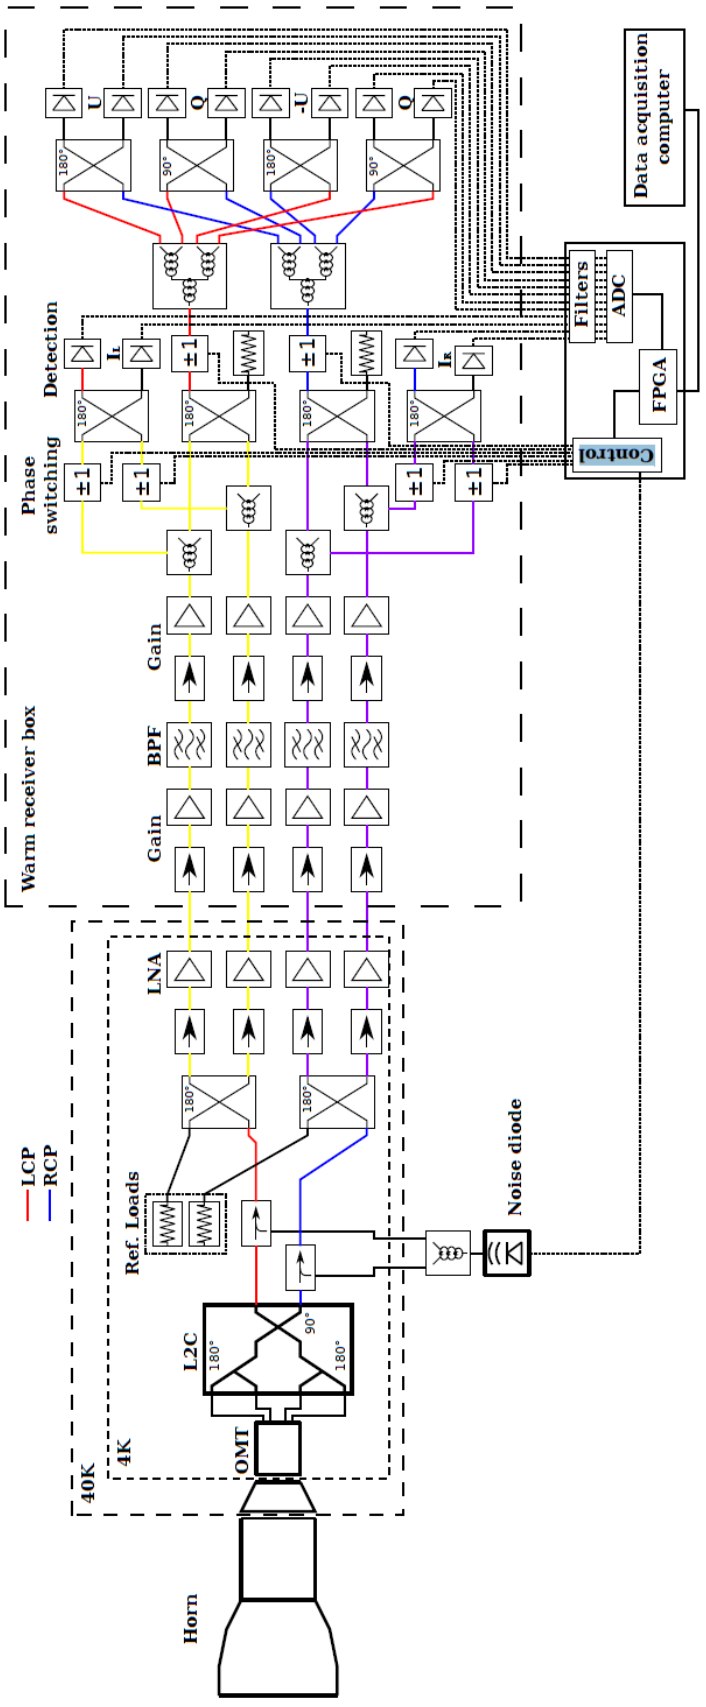

Figure 2.1: The analog C-BASS radiometer/polarimeter designed by Dr. Oliver King and currently installed on the Owens Valley Antenna (King, 2009).

Component	T_{phys}	Gain[dB]	G[Lin]	$\prod G$	T_{rx}[K]
Horn	300±1 K	-0.010±0.01	1.00±0.01	1.00	0.69±0.01
OMT	11±0.1 K	-0.50±0.05	0.90±0.01	0.89	1.23±0.02
L2C	11±0.1 K	-0.50±0.05	0.90±0.01	0.89	1.84±0.04
Coupler	11±0.1 K	-0.05±0.05	0.99±0.010	0.80	1.91±0.04
Hybrid	11±0.1 K	-0.7±0.1	0.85±0.02	0.79	2.86±0.07
Isolator	11±0.1 K	-0.7±0.1	0.85±0.02	0.67	3.98±0.10
LNA	T_{amp}=8.0 K±1.0 K	27±0.1	500±11	0.57	26.08±1.10
Cable	40±1 K	-1.2±0.1	0.75±0.02	397	26.15±1.56

Table 2.1: The Northern receiver cascaded noise temperature using the 8 K Manchester LNAs.

18 months of survey time to reach a similar level of sensitivity using the Jodrell amplifiers.

To reduce the unacceptably long survey time, we switched from the Jodrell LNAs to the Low Noise Factory LNAs over a 3 week period in September 2012. To evaluate the improvement from this change, we measured the $T_{\rm rx}$ using both the Jodrell and Low Noise Factory LNAs using the Y-Factor method. The Y-Factor method measures $T_{\rm rx}$ by observing the output power with a hot optical load and a cold optical load. We typically used a 300 K radio frequency absorber for T_{hot} and the sky, which has a temperature ∼6 K at the C-BASS frequency, for the cold load. The C-BASS receiver combines the optical load temperature ($T_{eccosorb}$ or T_{sky}) with the 12 K reference load temperature (T_{load}) using the 180° hybrid, producing observed temperatures T_{hot}=($T_{eccosorb} + T_{load}$)/2 and T_{cold}=($T_{sky} + T_{load}$)/2.

The power we observe in the two cases is then given by

$$P_{hot} = GK_B\Delta\nu(T_{hot} + T_{rx}) \tag{2.1}$$

$$P_{cold} = GK_B\Delta\nu(T_{cold} + T_{rx}). \tag{2.2}$$

We calculate the ratio $Y = P_{hot}/P_{cold}$ and, assuming both G and $\Delta\nu$ are constant, to solve for $T_{\rm rx}$

$$T_{rx} = (T_{hot}) - Y(T_{cold})/(Y - 1) \tag{2.3}$$

The results are summarized in Table 2.2. $T_{\rm rx}$ was found to have improved from 22.6±5.2 K to 13.6±4.6 K. This is largely consistent with the estimates obtained in Table 1.3 and Table 2.1 that predict an improvement from 26.2±1.6K to 13.8±1.6.

2.3 Coldhead Signature and OMT alignment

The Northern data were contaminated by a signal correlated strongly with the cold head refrigeration cycle that has confounded us for some time. It could arise because of a number of

LNA	P_{hot}[dBm]	P_{cold}[dBm]	T_{load}	Y[dB]	T_{rx}	Date
Manchester	-41.0$\pm$0.1	-48.4$\pm$0.1	12$\pm$0.1K	7.4$\pm$0.2	22.6$\pm$5.2 K	18/9/2012
Low Noise Factory	-14.5$\pm$0.1	-23.9$\pm$0.1	3.6$\pm$0.1K	9.4$\pm$0.2	13.6$\pm$4.6 K	10/10/2012

Table 2.2: Measured improvement of the Northern receiver $T_{\rm rx}$ by changing LNAs. Manchester LNA have a $T_{\rm amp}$=8 K while Low Noise Factory have $T_{\rm amp}$=3 K. $T_{\rm rx}$ is estimated using the modified Y-factor measurement where T_{hot}=($T_{eccosorb}$+T_{load})/2 and T_{cold}=(T_{sky}+T_{load})/2.

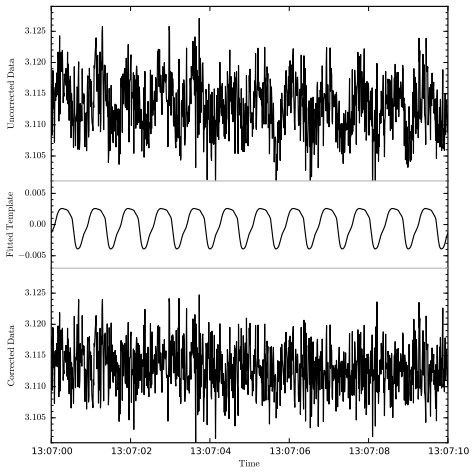

(a) Time Series. Top: raw data. Middle: fitted template. Bottom: corrected data.

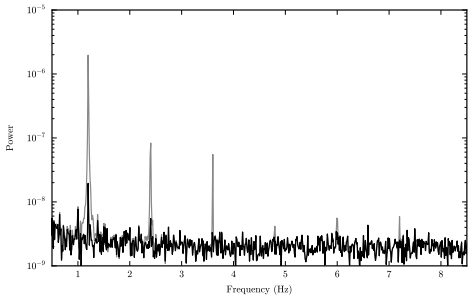

(b) Power Spectrum. Dotted line is before template removal and black is after template removal.

Figure 2.2: Removal of the cold head refigeration cycle effect from the data using template fitting. Figures from Matthew Stevenson.

possible mechanisms,

1. Loose connectors
2. Vibration sensitive LNAs
3. Vibrations on connectors in LNA bias lines
4. Lateral movement of the OMT relative to the horn

With these in mind, we replaced a number of internal cryostat RF cables that had developed cracks, tightened connectors, and succeeded in reducing the signal level significantly. In tandem with these mechanical efforts, Matthew Stevenson (Caltech) developed a template fitting algorithm to remove the signal from the recorded data. The template fitting reduced the level of signal to the level seen in Figure 2.2.

We made extensive modifications during the manufacture of the Southern cryostat as a result of these experiences, specifically improving and stabilizing the alignment of the OMT with the horn. More detail is provided in Section 3.4.

2.4 Low Frequency Contamination

The C-BASS antenna scans at a relatively high speed of 4°/s in order to minimize the effect of residual gain drifts not removed by the continuous-comparison architecture. These drifts would otherwise increase the noise on large angular scales. With the $\sim$0.7° beam size, this scan rate translates to $\sim$6 beams/s and consequently modulates the small angular scales onto the RF signal at $\sim$6 Hz.

This makes low-frequency artefacts problematic. From the beginning of the deployment the Northern system suffered from significant low-frequency interference caused by mains frequency pickup. The harmonics of the 60 Hz U.S.A. mains frequency are aliased by the 100 Hz integration period. The 5th harmonic is at 300 Hz which aliases to 0 Hz, affecting the sky signal.

To remove this, a series of filters were implemented in the FPGA backend and the phase switching was properly utilized. Modulating the sky signal at the phase switch frequency, f_{sw}, modulates the $\sim$6 Hz sky signal up in frequency to a band centred around f_{sw} =500 Hz. Any low-frequency contamination after the phase switching, is then separated in frequency from the sky signal by f_{sw}. When we demodulate the signal, we switch the two around in frequency, moving the sky signal back down to baseband, and the contaminating low-frequencies up to f_{sw}. Filtering the demodulated data with a suitable low-pass digital filter removes the contamination, and leaves only the sky signal.

2.5 Radio Frequency Interference

After deploying the Northern receiver, it became evident that the survey would be limited by the Radio Frequency Interference (RFI) environment at OVRO, despite it being a relatively good site. Direct-broadcast TV C-Band communications at $\sim$4.2 GHz were still visible despite being below the cut-off frequency of our band pass filter, and 5 GHz point-to-point terrestrial links are exceedingly popular in the Owens Valley (see Figure 2.3). Aircraft communications, both from commercial aircraft and the military, were also visible, although their sporadic nature makes them less problematic. The effect of these artificial contaminants is clear from Figure 2.8. Visible contamination is evident up to 50° elevation, making post-measurement removal very difficult.

To remedy this on the Northern receiver, we were forced to employ an expensive (both in time and in instrument performance) mitigation strategy. To reduce the effect of in-band RFI, we designed and a built a set of narrowband notch filters, and to improve the out-of-band rejection we added a second stage of band-defining filters.

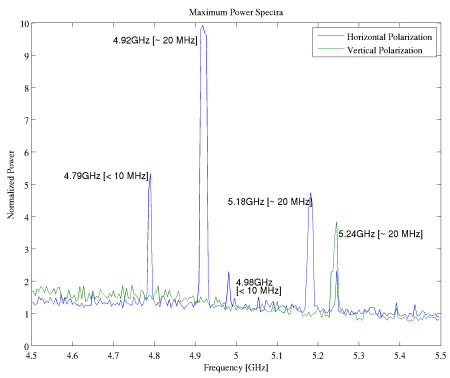

(a) Radio Frequency Interference as captured by a spectrum analyser attached to the output of the C-BASS cryostat. We think the 4.79 GHz spike is sporadic and does not require filtering

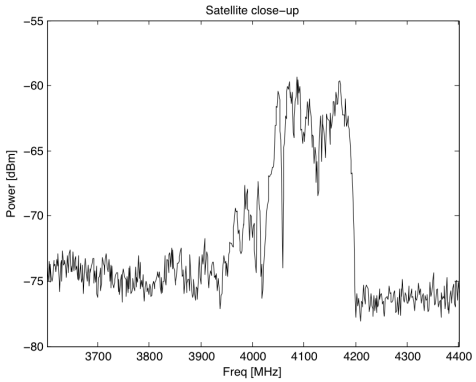

(b) Spectra of geostationary satellites. While they are outside the C-BASS band, these are still powerful enough to be detected with only one band-pass filter installed.

Figure 2.3: Radio Frequency Interference at Owens Valley Radio Observatory. There were four sources of narrow in-band RFI, and a number of sources of out-of-band RFI strong enough to be detected despite the bandpass filtering stage.

2.6 Notch Filters

Narrow band notch filters are a possible solution to the problem of in-band RFI. Examining Figure 2.3 suggests that each RFI band is approximately 25 MHz wide, which at 5 GHz requires a $Q = \frac{f_c}{\Delta f}$ =200. This is a challenging task.

However upon more careful examination of Figure 2.3, it is clear that some relaxation of this is possible. Each of the major RFI bands (i.e. 4.92 GHz and 5.18 GHz) has a second proximate RFI peak that could be included into a single notch filter. To attenuate both of these a stop-band of $\sim$80 MHz is required. This is an achievable goal.

2.6.1 Design

The design I implemented uses resonance structures placed in parallel with a transmission line. These structures resonate at a specific frequency, selectively dissipating this energy through ohmic and radiative losses in the resonating structure. The resonant frequency can be changed by adjusting the dimensions of the resonating structure.

For C-BASS we used the square resonance structure from Garcia-Garcia et al. (2005) shown in Figure 2.4. The split-ring resonator structure (SRR) acts as a self-resonant circuit where the two loops/squares act as inductors and the gaps act as capacitors. Varying the loop dimensions tunes the resonant frequency. The square resonance structures edge-couple to the micostrip transmission line and couple energy out at the resonant frequency. This is easy to

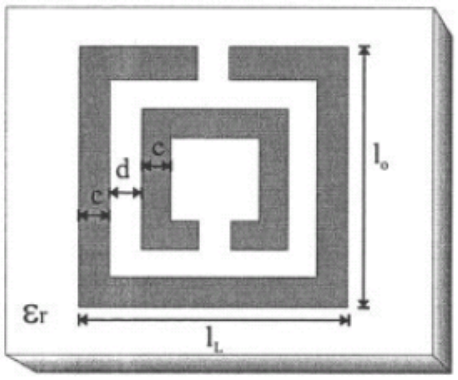

(a) Square Resonator

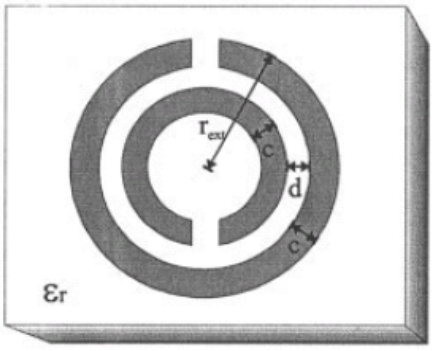

(b) Circular Resonator

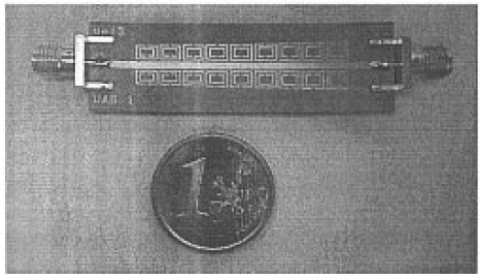

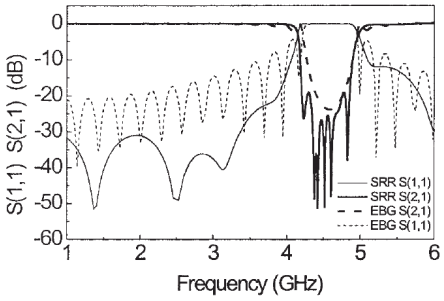

(c) Manufactured Bandstop filter with a series of differently tuned square resonator structures

(d) S21 Bandstop filter Result- See SRR plot

Figure 2.4: Resonance structures for the notch filter (Garcia-Garcia et al., 2005). These diagrams show the dimensions that can be adjusted. We used the square type resonance structure in (a).

tune and multiple structures can be stacked along a transmission line to increase the attenuation in the stop-band. The width of the resonance is controlled by ohmic and radiative losses in the loop and can be made more selective by cooling the filters to cryogenic temperatures thereby reducing the ohmic losses in the circuit and improving the Q.

We simulated and optimized the structures shown in Figure 2.4 for the RFI frequencies in Figure 2.3 (4.79 GHz, 4.92 GHz, 5.18 GHz and 5.24 GHz) using *Ansoft Designer*. The designs were etched onto 10 mil Rogers RO4350 substrate with a dielectric constant, ϵ_r=3.66. This produced satisfactory results as shown in Figure 2.5.

2.6.2 Filter Installation at Owens Valley Antenna

Figure 2.6 shows the filters after being installed in the Northern cryostat. The positive effect on data quality is clearly demonstrated in Figure 2.8 and, with the 0.25s/pixel integration period used to produce these scans, RFI is no longer visible. Unfortunately, this improvement comes

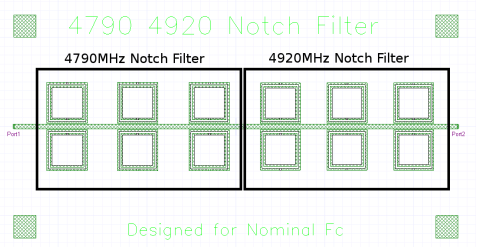

(a) The 4.79 GHz and 4.92 GHz notch filter

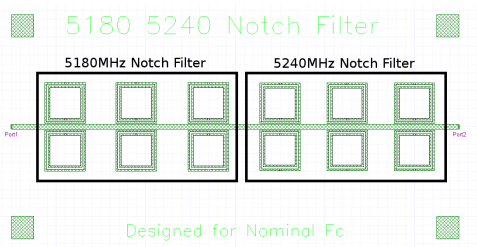

(b) The 5.18 GHz and 5.24 GHz notch filter

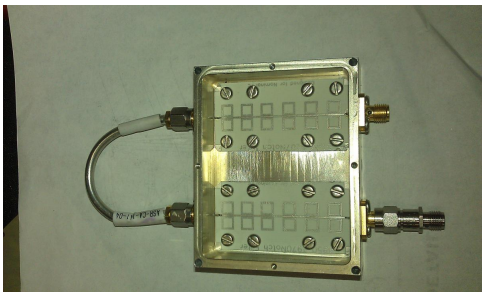

(c) Manufactured notch filter (with 6 dB attenuator to improve input match)

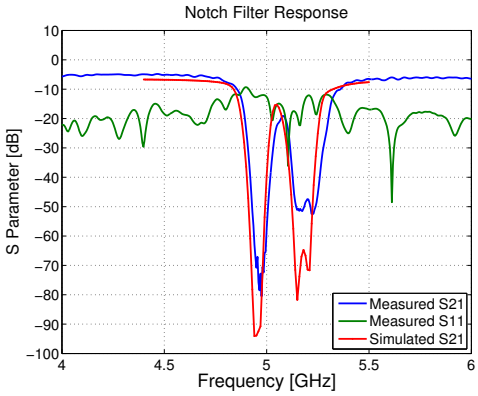

(d) Measured vs Simulated Responses

Figure 2.5: The design and manufacture of the notch filters for the Northern C-BASS receiver. We used the square resonance structures introduced in Figure 2.4 to simulate and build filters for 4.79 GHz, 4.92 GHz, 5.18 GHz and 5.24 GHz RFI present at the Owen's Valley. The substrate used was 10mil RO4350 (ϵ_r=3.66) The filters were packaged in milled, RF tight aluminium boxes and cooled to 40 K to reduce the ohmic losses and improve the Q factor. High reflection in the filter stop-band was originally mitigated using 6 dB attenuators which were present in the measurements and simulations above. The attenuators were replaced with low loss (0.3dB) cryogenic isolators in the cryostat to improve the total T_{sys}.

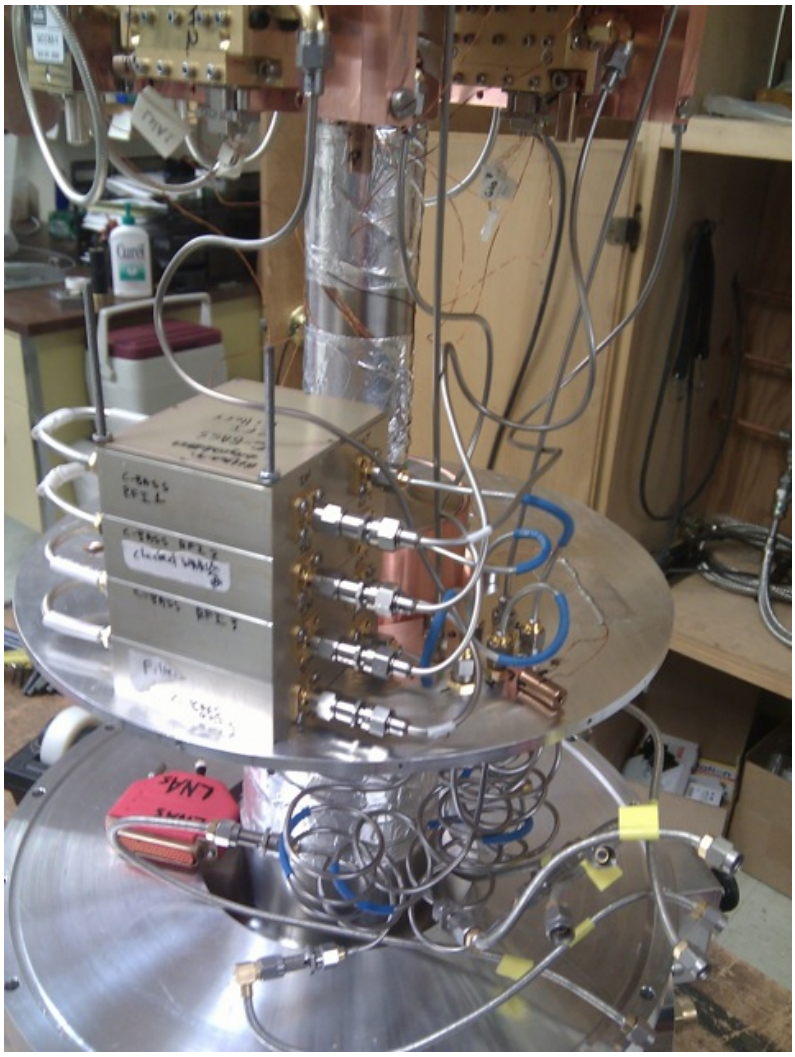

Figure 2.6: The notch filters installed in the cryostat at Owens Valley.

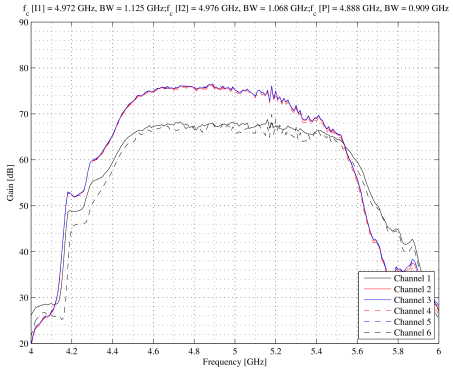

(a) The C-BASS pass band before installing any additional filtering

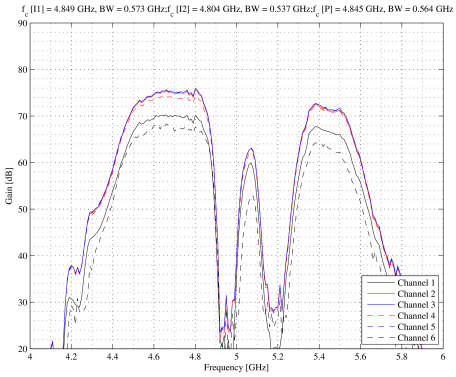

(b) The C-BASS pass band after installing the Notch filters and additional Band pass filters

Figure 2.7: The effect of the notch filters on the C-BASS pass band. Removal of the in-band RFI reduces the effective bandwidth by 340 MHz, reduces our sensitivity by 21%.

at a price, namely a reduction in sensitivity by 21%, because of the 340 MHz decrease in effective bandpass shown in Figure 2.7.

2.6.3 Possible Improvements

High temperature superconducting filters (Futatsumori & Hikage, 2008), could achieve higher Q values, potentially allowing us to increase the remaining bandwidth after notch filtering. However the manufacture process is complicated, expensive and, in the case of C-BASS, the frequency distribution of the four RFI signals would require extremely high-Q filters for HTS filters to make any significant difference.

The most appealing possibility is to replace the analog receiver with a digital backend and do the digital signal processing in the frequency domain. Provided there was sufficient cross-channel isolation we would be able to remove RFI with a precision given by the frequency resolution by simply zeroing the appropriate frequency channels. This is the route taken for the Southern receiver. which has an 8 MHz spectral resolution and achieves >60dB out-of-channel rejection using a polyphase filter bank implementation of a Hamming window. For the RFI seen in Figure 2.3 we would only need to discard 10×8 MHz channels, which would improve the sensitivity by 18% over the current notch filtering solution.

2.7 Northern Survey Status

I have not been intimitately involved in reducing the Northern data, however it would be remiss not to include any preliminary results. The Northern survey is progressing well, and we have

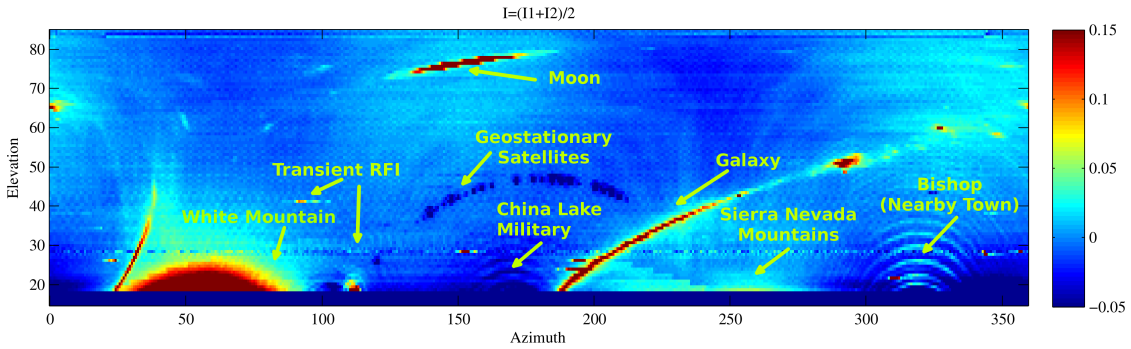

(a) Sky before the installation of notch filters. Note the diffraction patterns produced by strong RFI sources detected in the antenna sidelobes

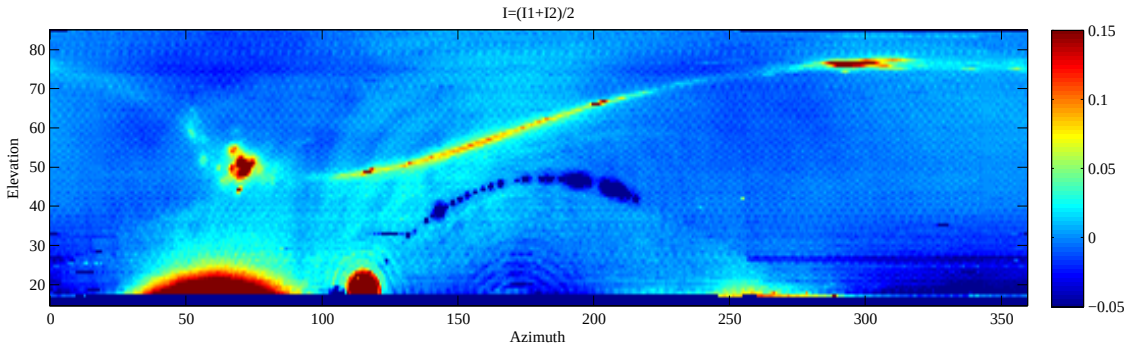

(b) Sky after the installation of notch filters. All terrestrial RFI aside from a strong source to the South (180° Azimuth) has been removed

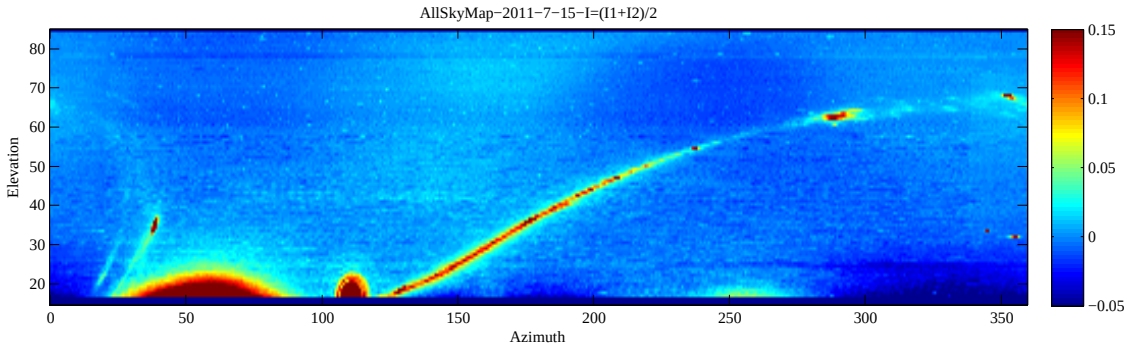

(c) Sky after the installation of notch filters and additional band pass filters. All terrestrial radiation removed.

Figure 2.8: The effect of notch filters on the data quality. These images show the dramatic improvement in data quality brought about by the addition of notch filters and additional band defining filters in the C-BASS RF path

developed a data reduction pipeline that will be used for reducing the Southern data.

This data reduction pipeline provides automated RFI flagging, and produces a signal calibrated against the instrumental noise diode. The calibrated signal is passed to the DESCART map-maker which uses a de-striping algorithm to produce a final map in a HEALPix format. The map produced from the daily reduction pipeline are shown in Figure 2.10a. Data taken over two months are given in Figure 2.10b. The polarization survey maps are currently limited by ground pickup. An accurate ground template is being made for Stokes I, Q and U which will be used to remove this.

Figure 2.9 shows an early small scale area of total-intensity mapped around Cygnus A and compares it against the lower frequency 408 MHz Haslam map and the 23 GHz WMAP map.

I have also included preliminary all-sky total-intensity map in Figure 2.10. This was produced using the DESCART de-striping map maker.

2.8 Conclusion

I have introduced the first iteration of the C-BASS instrument, a broadband analog continuous-comparison radiometer combined with an analog correlation polarimeter. After the original deployment we upgraded the sensitivity of the receiver by changing from 8 K Manchester LNAs to 3 K Low Noise Factory LNAs. We mitigated against a signal associated with the coldhead cooling frequency using both mechanical improvements and a template fitting approach. RFI is a real weakness of broadband analog receivers and I have described the design and manufacture of high-Q cryogenic notch filters that were used to remove in-band interference from the C-BASS data. I have tackled this weakness in the Southern receiver by using a digital backend and doing the digital signal processing in the frequency domain. The chapter concluded by presenting images constructed from the Northern instrument that show preliminary results of C-BASS. These are promising and show the advanced nature of the C-BASS data reduction pipeline and map making algorithms.

The next chapter focuses on the Southern C-BASS instrument. This uses an improved digital implementation of the analog receiver detailed in this chapter that provides an appreciable level of immunity to in-band RFI - a similar implementation on the Northern telescope would improve the sensitivity by an estimated 18% over the cryogenic notch filters. I will describe the telescope at a system level, and provide detail on a number of wide-ranging aspects that will include the optical design, telescope servo controller, the cryogenic receiver front-end, the analog signal processing, the digital receiver backend and the high level integrated control system.

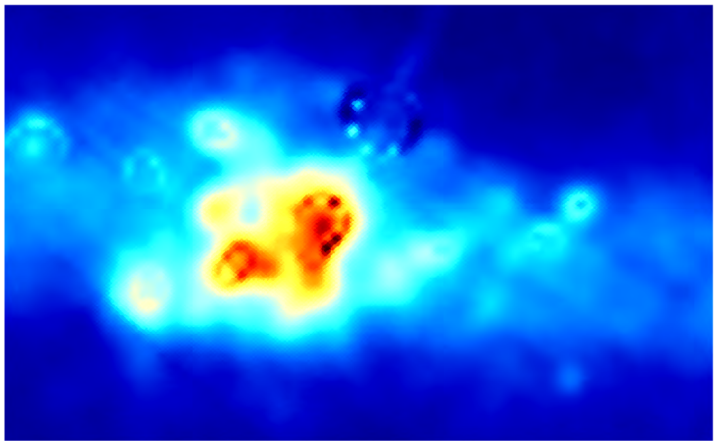

(a) Haslam 408 MHz

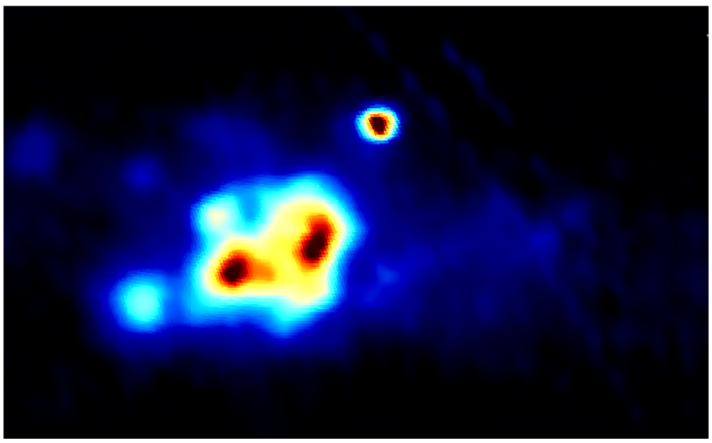

(b) C-BASS 5 GHz

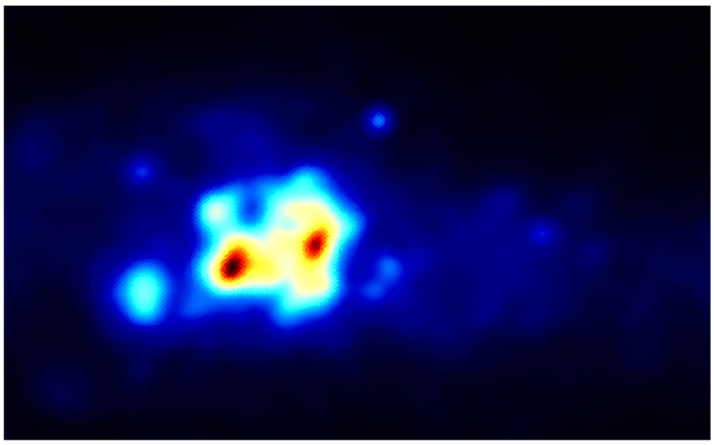

(c) WMAP 23 GHz

Figure 2.9: Image of Cygnus X region at three different frequencies- produced by Matthew Stevenson- C-BASS Collaboration.

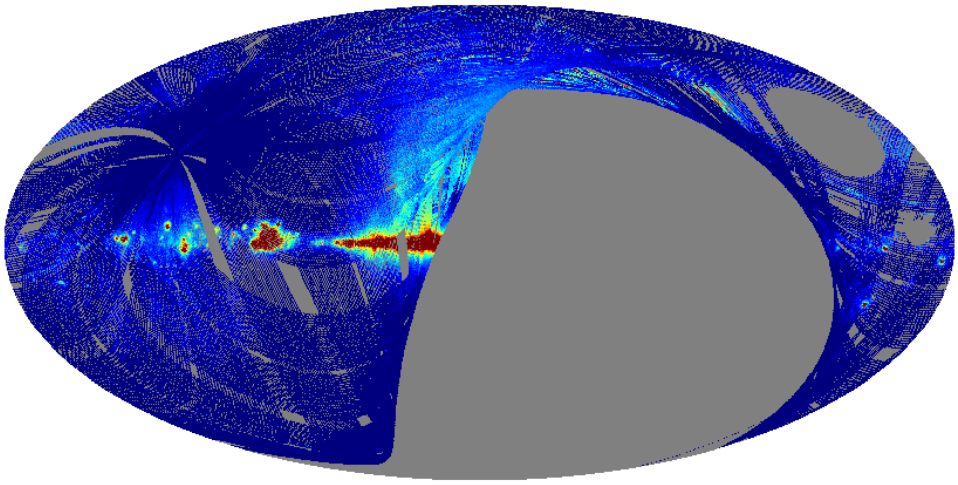

(a) One day.

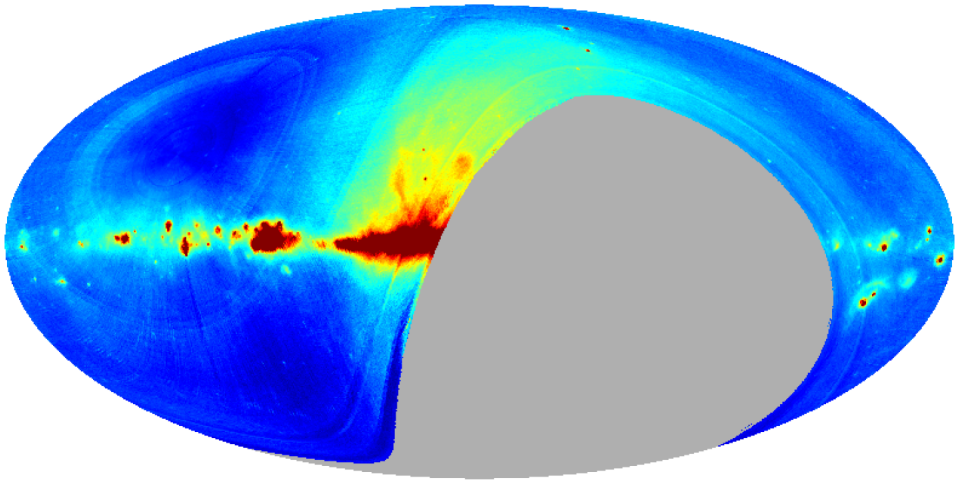

(b) Two months.

Figure 2.10: All-sky total intensity maps- the daily map is produced by an automated daily reduction pipeline written by members of the C-BASS collaboration, and the two month map was produced by Melis Irfan with data between November and December 2012.

Chapter 3

The Southern C-BASS Telescope

> This digital world is completely fascinating to me.
>
> Graham Nash

This chapter will describe the development of a new C-BASS receiver which boasts a number of improvements and modifications over the analog receiver described in the previous chapter. The most significant of these improvements is a fully-digital spectral backend that provides a level of immunity to in-band RFI as well as opening the possibility of studying the foreground total-intensity and polarization across the 1 GHz C-BASS operating band. The polarization measurements across the band will make it possible to measure the Faraday polarization rotation between 4.5–5.5 GHz. The digital receiver I will describe is scalable in frequency, allowing much wider operating bandwidths to be covered, and could also easily be duplicated to provide a tested solution for deployment on future similarly renovated antennas with minimal risk or cost. This is made even more applicable by describing the renovation of a donated antenna for use in the C-BASS survey. Together, these provide a roadmap for other similar antenna upgrades.

The chapter is divided as follows:

1. Optical Measurement and Design (Section 3.1 and Section 3.2)
2. Servo Control (Section 3.3)
3. Cryogenic Receiver (Section 3.4)
4. Signal conditioning (Section 3.5)
5. ROACH Digital Backend (Section 3.7)
6. Integrated Control System (Section 3.8)

7. Site Infrastructure (Section 3.9)

3.1 Measuring and Verifying Antenna Optics

The dish that was donated from Telkom did not arrive with any detailed shaping information. We therefore required measurements of both the primary and secondary reflectors in order to evaluate the antenna optics. This is likely to be the case in any potential telescope conversion. This section describes the techniques used to measure the shapes and also presents a cross-scan of a geostationary satellite.

3.1.1 Primary Dish Measurements

We required an accurate model of the primary surface shape, together with an estimate of the surface accuracy. The size of the dish and parabolic shape make it difficult to use traditional measurements techniques such as theodolite and laser ranging. Instead, we chose to use close-range photogrammetry.

Photogrammetry

Close-range photogrammetry is a recent innovation that allows highly accurate measurements of a distribution of optical targets to be made in a fast, automated fashion (Luhmann, 2010). We used close-range photogrammetry via the commercially available *Photomodeller* software package[1] to measure the shape of the primary reflector.

In close-range photogrammetry, optical targets are distributed across a surface and photographs are taken from different angles. Using these photographs, it is possible to determine the positions of the individual targets using a least-square fitting procedure. By fitting centroids to circular target positions, photogrammetry routinely achieves sub-pixel resolution in the point's position. A few of the targets are usually geometrically coded as seen in Figure 3.1a, which allows image-processing software to identify these points in the photographs. This makes it possible to semi-automate the post-processing. The measurement accuracy varies as a function of camera resolution, target contrast, and separation angle of the photographs; accuracies of up to 1:250 000 are routinely achieved when using carefully calibrated equipment. For the C-BASS measurements, we used a suitable high-end consumer camera to take the photographs; this typically has an accuracy of 1:10000 (Deng & Faig, 2001). On the 7.6 m C-BASS dish this is better than 0.8 mm.

[1] http://www.photomodeler.com/index.html

Target Choice

The most common targets used in photogrammetry are retroflective targets. These are usually black adhesive 'stickers' with reflective sections hand-pasted into place and can be expensive to manufacture and place.

For C-BASS, we experimented with the slightly unconventional approach of projecting targets onto the dish surface using a data projector. Figure 3.1a shows the general idea. We hoped that this would provide both a higher density distribution of points across the surface and a less expensive implementation. Unfortunately, we found that the resolution and contrast of the projector limits the effectiveness of this approach and we were forced to abandon this technique in favour of the retroreflective targets.

Primary Shape Measurement results

The results of the photogrammetry measurements are presented in Figure 3.1. We see that the primary surface is not a 'strict' paraboloid, but rather falls into the category of 'shaped optics'. This type of design features an offset from the traditional parabolic reflector described in detail in Galindo (1964), and is a technique used to optimise the forward gain of antennas.

We modelled this shaped surface using a 5th order polynomial to express the depth of the primary reflector, z, as a function of the radius, r. The best fit is given by

$$z(r) = -0.000524r^5 + 0.004680r^4 - 0.016472r^3 + 0.120840r^2 - 0.015305r. \quad (3.1)$$

3.1.2 Secondary reflector measurements

A secondary reflector was provided with the antenna. This mirror could have been measured in the same way as the primary dish using photogrammetry, however it was small enough to place on a milling machine. This allowed us to measure the depth to a high precision. The profile is shown in Figure 3.2.

3.1.3 Optical Simulation

Using the measurements of the primary and secondary profiles (Figure 3.1 and Figure 3.2 respectively), we were able to simulate the optics of the entire system. The *GRASP* ray-tracing diagram is shown in Figure 3.3a. This diagram shows the predicted focal point, $A_{\rm fp}$, of the optical arrangement and it is clear from this figure that the optics are out of focus. This was not easy to correct, but while science would not have been possible with the system in this state, it does provide an indirect way of checking the accuracy of the photogrammetric measurements.

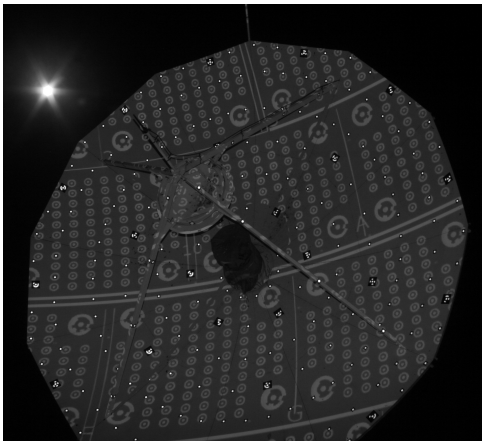

(a) The high-contrast, bright points are the retroflective targets, while the other (more densely distributed) points were projected onto the dish from a data projector. The targets consist of both standard circular targets and geometrically coded targets to help in the post-processing photograph orientation.

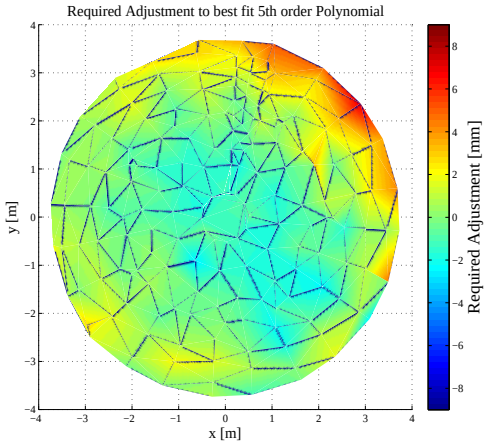

(b) Colour shows the deviation from model used in optical design (Equation 3.1) before transport. The r.m.s. error was 1.61 mm

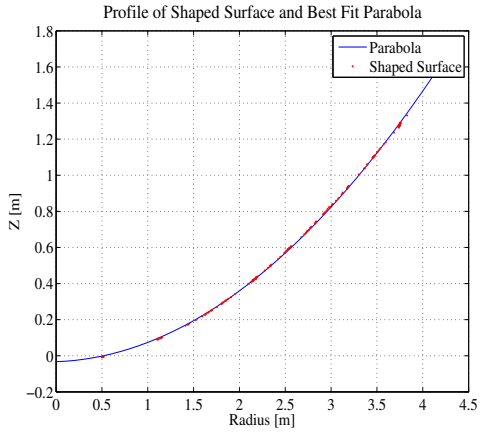

(c) Section through the dish surface relative to the best-fit parabola

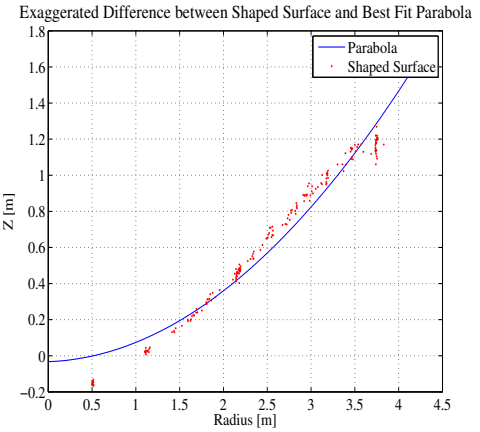

(d) Exaggerated (x20) section difference from the best-fit parabola to show the shaping of the primary

Figure 3.1: Photogrammetry uses a series of photographs such as those shown in (a). The primary shape used in the optical design is derived by fitting a 5th order polynomial to the cross-section in (c) and (d). The offsets across the face of the dish shown in (b) are with respect to this 5th order polynomial.

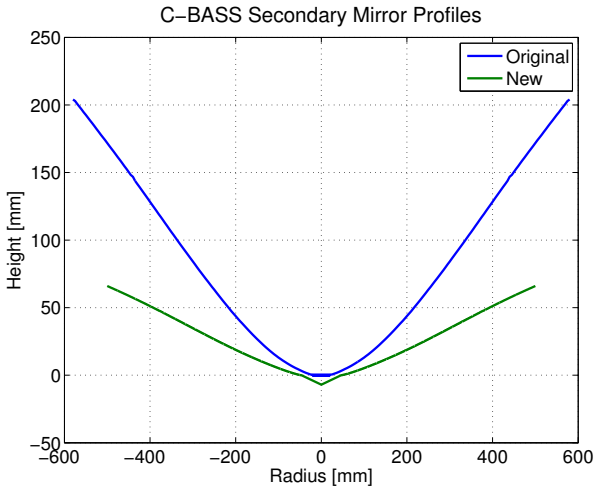

Figure 3.2: Old and new secondary reflector profiles. The old secondary was measured by setting it up on a milling machine. The new mirror was manufactured in the Oxford Physics workshop. The small bump in the new secondary design is to direct any reflections from the horn back onto the dish, preventing the formation of standing waves.

Satellite cross-scans

If we move a horn through the antenna focal point, $A_{\rm fp}$, while simultaneously measuring the detected power, we would expect a maximum when the horn phase centre, $H_{\rm fc}$, is coincident with the $A_{\rm fp}$. This provides a way of checking the photogrammetry measurements used in the *GRASP* simulation. If $A_{\rm fp} = H_{\rm fc}$, we can deduce that the photogrammetry results must be reasonably accurate.

To make the horn phase centre measurement, I modified a 12 GHz receiver and built a custom rectangular feed horn (Meeks, 1976) to illuminate the secondary mirror. The phase centre of these horns is well understood and is very close to the horn aperture. This test setup

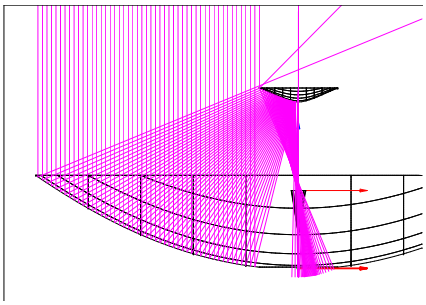

(a) Ray tracing diagram of the system as it was given to us by Telkom. The optics are not in focus, and it was not possible to refocus them.

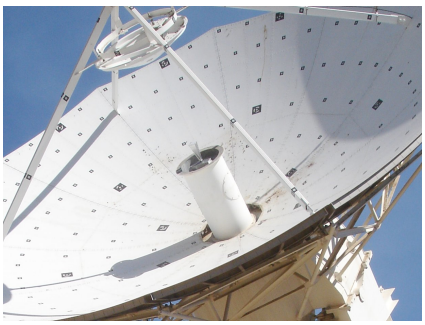

(b) 12 GHz Receiver installed at the focus. It is possible to move the receiver along the z-axis of the antenna

Figure 3.3: Simulations of the optics before the C-BASS upgrade.

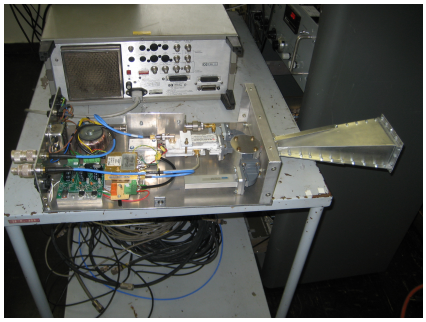

(a) 12 GHz Receiver and newly designed feed

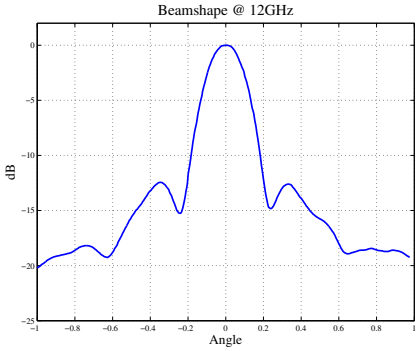

(b) Beamshape of 12 GHz optical setup confirms first order optical alignment

Figure 3.4: The 12 GHz satellite observations.

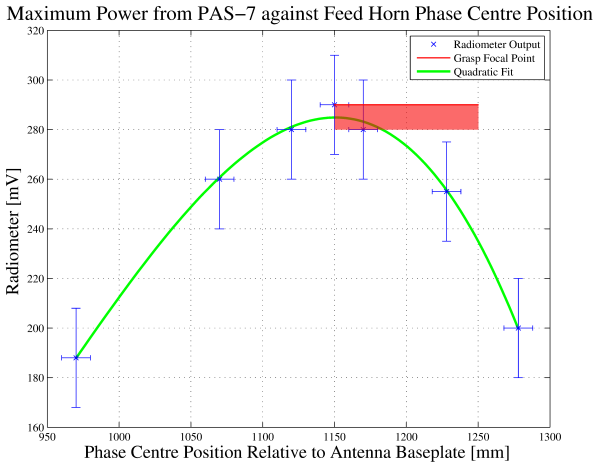

Figure 3.5: Measuring the antenna phase centre. Maximum radiometer output voltage when scanning through the PAS-7 12 GHz satellite, is plotted against feed phase centre position during the scan. The *GRASP* predicted phase centre is 1200±50mm

allowed us to observe ∼12 GHz transmitters from geostationary satellites. By moving the horn vertically to measure the maximum detected power, I was able to determine where $A_{\mathrm{fp}} = H_{\mathrm{fc}}$. A photograph of the assembled receiver is included in Figure 3.4a. An azimuth cross scan through the 12 GHz satellite in Figure 3.4b shows that there is no gross optical misalignment.

The measured power against position is plotted in Figure 3.5. Maximum power is expected when the horn phase centre and antenna focal point are coincident. The red marker shows the *GRASP* prediction of the antenna focal point in Figure 3.3a, $A_{\mathrm{fp}} = 1200 \pm 50$ mm. The measurements show a maximum power with the horn phase centre position $H_{\mathrm{fc}} \approx 1150$ mm. The consistency between A_{fp} and H_{fc} is reassuring and gives confidence to the photogrammetric measurements.

3.2 New Optical Design

Using the photogrammetry measurements of the Southern antenna primary reflector shape in Equation 3.1 and similar measurements obtained for the Northern antenna, Dr. Christian Holler re-designed the C-BASS optics to produce similar beams with heavily suppressed side-lobes for both Northern and Southern telescopes. Simulated results are shown in Figure 3.9, taken from Holler et al. (2013). The similarity of the two telescope beams will help when combining the Northern and Southern maps to produce the final C-BASS data product. This section describes the process of designing these optics, from measurement to implementation.

3.2.1 Horn

In order to meet the C-BASS sensitivity goals, we require a system with very low side-lobes. This reduces ground pick up, which is both difficult to characterize and and adds unwanted noise to the sky measurements. A suitable horn was designed by Richard Wylde of Thomas Keating Ltd. using a circular profiled corrugated design. These designs are often used when such low sidelobe performance is required. The simulations predicted extremely low sidelobes and beam pattern measurements conducted verified these simulation, as seen in Figure 3.6. On the C-BASS receiver the transition between the OMT is implemented by a tapered section machined as part of the cryostat structure. For the beam pattern tests we manufactured a separate transition piece between the OMT and the corrugated horn sections that allowed the horn to be tested independently of the cryostat.

When we measured the bandpass power spectrum using the complete receiver (see Figure 3.7), we discovered a suite of extremely narrow bandwidth features in the lower 500 MHz of the receiver band. These were also evident in T_{sys} measurements, confirming that these narrow features were lossy. They were not detectable without the horn in place. Confirmation that the horn was responsible for these features was obtained by making measurements of the S11 of the horn (Figure 3.8b) that showed these features clearly. The original *CORRUG* horn simulations did not provide sufficient frequency resolution to resolve this structure in Figure 3.8a. To try and understand the cause of the features, we simulated the horn response at a higher frequency resolution using the *CHAMP* simulation software.

The *CHAMP* higher-frequency resolution simulations predict similar structure to what was observed in our bandpass measurements. Similar effects were described in Bhattacharyya (2001), who attributes them to two simultaneously propagating modes coupling at the horn aperture or at transitions between corrugated and smooth horn sections. The effect was removed by remachining the corrugations in the throat of the horn, breaking the regularity of the geometry inside the horn. Simulations and measurements of the S11 of the horn in both states

(a) The OMT was clamped onto the horn transition section and rotated to align it with the transmit horn polarization

(b) The transmit horn was rotated to change the polarization of the test signal

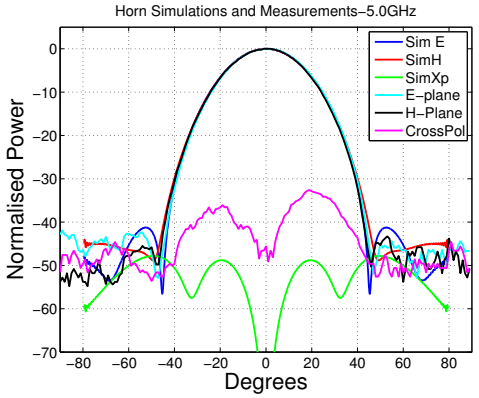

(c) The simulated and measured horn beam patterns

Figure 3.6: Horn copolar and crosspolar beam patterns. The horn beam pattern was measured by rotating the C-BASS horn through the beam of a 5 GHz rectangular transmit horn. The results at 5.0 GHz are shown, together with 5.0 GHz *GRASP* simulations.

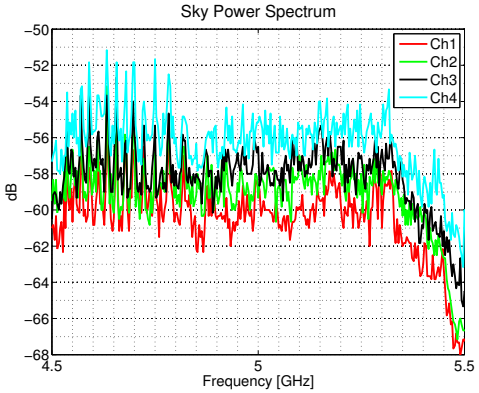

Figure 3.7: Measured bandpass spectrum of the system between 4.5–5.5 GHz. The narrow frequency structure was present on all four channels.

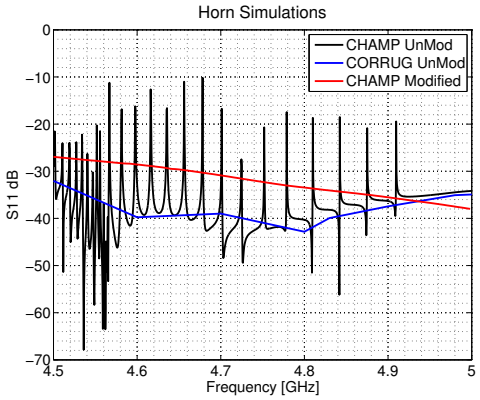

(a) *CHAMP* and *CORRUG* simulations of the unmodified and modified horn between 4.5–5.0 GHz

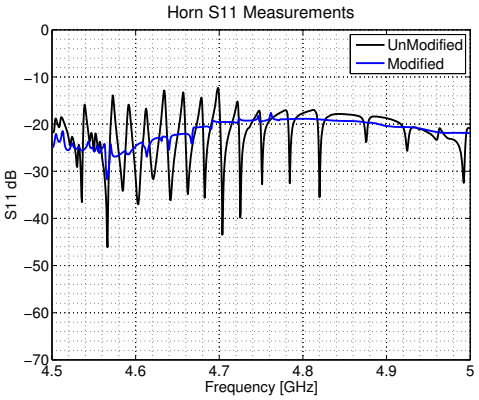

(b) S11 measurements of the unmodified and modified horn between 4.5–5.0 GHz

Figure 3.8: C-BASS horn narrowband structure in lowest half of the band. The resolution of the original *CORRUG* simulation was not sufficient to resolve this structure, however it is clearly visible in the higher resolution *CHAMP* simulation. This was removed by remachining the horn. Simulations by Dr. Jamie Leech.

are shown in Figure 3.8.

3.2.2 Secondary Mirrors and Foam support

The original secondary mirror was not suitable for C-BASS as it does not produce a beam matched to the Northern survey and has high sidelobes due to scattering. As such we designed and built a new secondary reflector. To avoid the quadrupod support polarization problem alluded to previously, the C-BASS secondary is mounted on a radio-transparent foam support. This has been tested to be rigid enough to support the secondary mirror on the Northern survey and has undergone extensive FEA modelling. The foam avoids the polarization problems associated with the quadrupod feed supports, which is critical for maintaining the polarization purity of the C-BASS astronomical signal.

3.2.3 Simulated Telescope Beam Patterns

The beam pattern of the Northern telescope was simulated in Holler et al. (2013). The results for the Northern and Southern telescope are shown in Figure 3.9. The beam pattern was also measured astronomically in this paper using data from the Northern telescope. However, the Northern telescope is not able to point lower than 15° elevation, preventing us from making high dynamic range measurements of the main lobe using a terrestrial signal source. These measurements will be made on the Southern telescope, which is able to drive to the horizon.

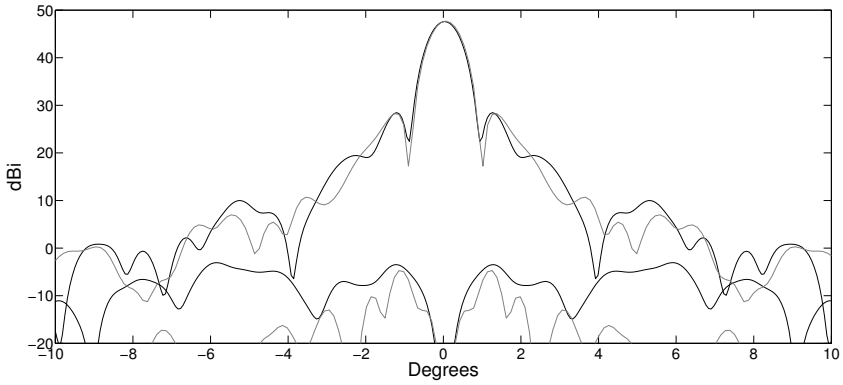

Figure 3.9: Simulation of the co- and cross-polar forward beams of the Northern telescope (black) and the Southern telescope (grey). (Holler et al., 2013)

3.3 Servo Control

The donated antenna was originally designed to track satellites, which emit a strong beacon signal specifically for this purpose. The design is not suitable for astronomical observations, where we must be able to drive the telescope to any specified absolute position on the sky. Rather, we required a servo controller able to point to a suitable accuracy for our 0.7° beam. In addition, our azimuthal scan strategy (see Section 1.6) and stability requirements need high speed scans across the sky to minimize the effect of residual gain drifts on the data. We needed to implement a control system capable of both these modes of operation.

3.3.1 Dish Servo System

The existing servo controller consisted of four DC motors driven by DC servo amplifiers. The output current of these amplifiers is controlled by a ± 10 V DC control voltage. Tachometers were present on the motors, providing velocity feedback in the form of a voltage, and 16 bit encoders using an SSI (serial synchronous interface) communications protocol were used for position feedback.

In order to drive the antenna, we required hardware to generate a control voltage for the servo amplifiers, sample the tachometer voltages and read the encoder feedback. We also required control logic for the various relays and electrical contactors in the system.

The servo controller itself needs a control loop bandwidth significantly higher than the lowest natural frequency of the antenna, f_0. This can be estimated by driving the antenna up and down at set frequencies and monitoring the response; however, as a rule-of-thumb Gawronski (2008) studied the fundamental frequency of a number of telescope structures and modelled f_0 as a function of the diameter, d. This is given by:

$$f_0 = 20d^{-0.7}, \tag{3.2}$$

providing an estimate of ∼5 Hz for the 7.6 m Southern antenna. I wanted to oversample this frequency, so I settled on a control loop bandwidth requirement of 100 Hz.

The pointing accuracy needed for an antenna depends on the beam size and is typically less than HPBW/20. For C-BASS this is 35 mdeg. In order to accurately position on this angular scale, the angle encoders need to measure ∼17.5 mdeg, or twice the required pointing accuracy. This requries an encoder resolution of at least $log2(360/0.0175) = 14.3$ bits.

3.3.2 Control Loop

A feedback control loop accepts a command and uses feedback from various sensors to execute the command. In the case of a radio antenna, the command is a time dependent position and tracking velocity (Gawronski, 2008).

Several different types of feedback control loop exist. Common loop examples used to control antennas are the PID (proportional, integral, derivative), LQG (linear quadratic gaussian) and the H∞ (Gawronski, 2001).

The C-BASS servo control loop uses a PID control loop to accurately set the antenna position, while the tracking velocity is controlled using a PI loop, which is identical but without derivative feedback. The complete form is shown in Figure 3.10, with two copies running independently on the azimuth and elevation axes.

For the positioning, the angle encoder position, $P(encoder)$, is compared against the desired position, $P(desired)$, to calculate a position error term, $e_p(t)$=$P(desired)$-$P(encoder)$. The error term is then used to calculate three new products: the proportional, integral and derivative products. The proportional product produces a response to the instantaneous error, $e_p(t)$, the integral responds to the accumulated error, $\int e_p(t)dT$, while the derivative predicts future error using the derivative $de_p(t)/dT$. The three products are then scaled by adjustable multiplicative constants, K_p, K_i, and K_d respectively, and the results are added together to produce a command voltage. This command voltage is presented to the servo amplifiers and can be thought of as a torque command to drive the antenna to a desired position.

Astronomy requires not only positional control, but also controllable tracking velocity control in order to accurately follow astronomical sources. To provide this, we implement another control loop using the velocity feedback provided by the motor tachometers. The desired drive velocity is compared against the tachometer feedback, producing a velocity error term, $e_v(t)$, akin to the position error, $e_p(t)$. We calculate proportional and integral products using the velocity error term, $e_v(t)$ in the same way as for the position error, and scale using the coef-

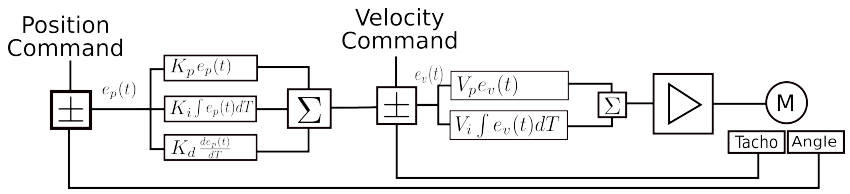

Figure 3.10: The C-BASS servo control loop. The coefficients K_p, K_i, K_d, V_p and V_i, were determined heuristically.

ficients V_p and V_i for proportional and integral products respectively. We omit the derivative term since the tachometer feedback is generally quite noisy, and including the derivative term can result in instability.

The result is a new adjusted torque command incorporating any additional torque required to track moving objects accurately. The torque command is given to the servo amplifiers, which in turn provide current to drive the motors.

The scaling coefficients, K_p, K_i, K_d, V_p and V_i were determined heuristically using a standard PID tuning technique described in Tim Westcott (2000). After tuning we were able to drive the antenna to within 1-bit of the encoder resolution. A response to a step command is given in Figure 3.13.

3.3.3 Servo Controller Implementation

The C-BASS servo controller block diagram is given in Figure 3.11. This section provides a description of the individual hardware components.

Hardware

I wanted the servo controller for the antenna to be inexpensive and capable of running an operating system. Linux provides a the number of open-source tools such as SSH for remote access, NTP for timing, and a thoroughly tested IP stack for communications. This makes makes a number of design aspects much easier. These requirements were met by the ARM9260 development board made by Olimex. The board runs the Debian flavour of Linux, and has hardware access to a number of GPIO banks for hardware communications to peripherals, ethernet access and USB ports. High level control can be built directly into low-level servo controller, improving scalability. This is desirable as such servo controllers could be used in future large-scale telescope arrays.

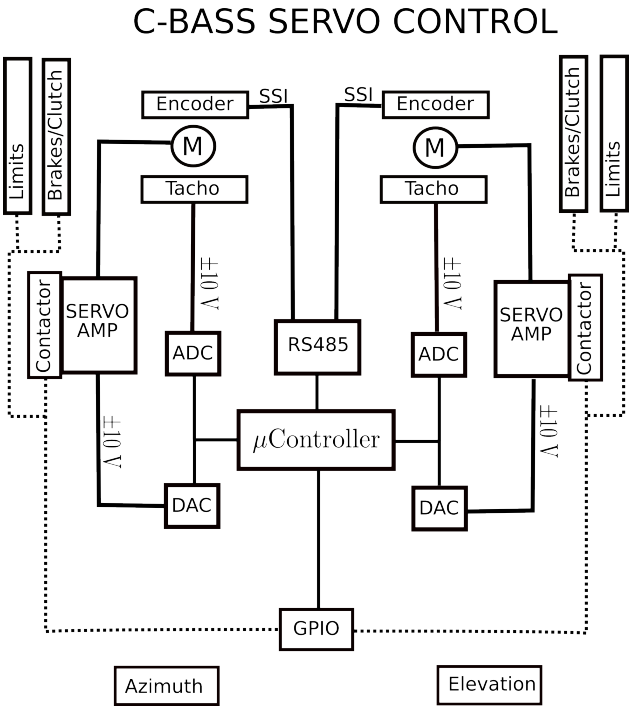

Figure 3.11: The C-BASS servo control implementation.

Realtime within an Operating System

Despite the numerous advantages outlined above, a significant problem does exist: Operating systems do not usually function in a deterministic real-time fashion. Operations are scheduled and guaranteed to occur, but the latency between scheduling and execution is variable. As a result operating systems are often avoided for high-bandwidth servo applications that require well-defined and deterministic timing.

These problems are surmountable. Real-time support for Linux does exist[2] and preemptive kernels are available in in almost all major operating systems. These allow a level of real-time control, provided one is willing to run this at the kernel level. We can take advantage of this to build a servo controller, by sampling the feedback sensors at regular well-defined intervals generated at the kernel level, and doing the less time-critical PID calculations in between the intervals. This works well, provided that the calculations do not usually take longer than the sample interval. This puts a limit on the achievable bandwidth. We also need to be aware that because of the nature of operating systems, we cannot guarantee the time required for these calculations. Fortunately we only need the vast majority to be suitably short.

I investigated running the sampling at 100 Hz and found that this was an achievable bandwidth as seen in Figure 3.12. Approximately 85% of the loops had latencies of 10500 μs, a majority that appears to be suitable for the servo control loop.

[2] `https://rt.wiki.kernel.org/index.php/Frequently_Asked_Questions`

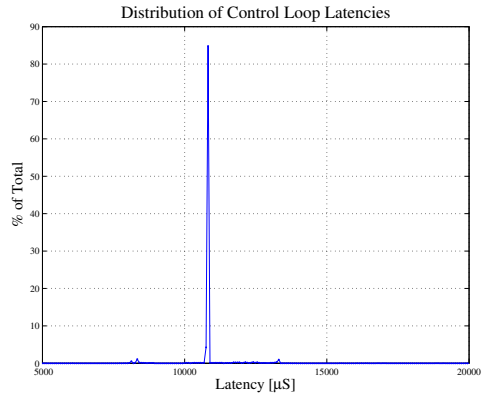

Figure 3.12: Histogram of the control loop latency. The loop is fairly stable at 10500 μS per loop (85% of the distribution), however the variation is still high. Later versions of the control system will use FPGAs to provide hard-timing of the PID loop.

In the next iteration of the control system (to be deployed on the African VLBI Network antennas) the low level servo control will be performed on inexpensive FPGA boards. These will be interfaced to inexpensive deployable ARM based CPU solutions capable of running operating systems, such as the nearly ubiqitous Raspberry PI. The FPGAs will provide the required deterministic sampling while the CPU will provide the same computational advantages as those outlined earlier. This solution makes the best of both worlds,

3.3.4 Sensor Hardware Design

Sensors were chosen to standardize the communications interface, and limit the peripheral circuitry required. Most sensors used the SPI protocol, a serial protocol with minimal hardware interconnect requirements. We designed a series of PCBs to allow for interfacing between the sensors and control PC, and to allow control voltages to be passed to the servo amplifiers.

ADC PCB

The ADC was implemented using an AD7367 IC. This provides dual 14-bit $\pm$10V sampling for the tachometer feedback. Interfacing is done through the SPI protocol directly to the ARM9260 development board, providing high bandwidth sampling of the tachometer directly to the control loop.

DAC PCB

We chose the AD5362 to supply the analog control voltage to the servo amplifier. This IC can provide quad $\pm$10V signals, allowing us to control all four servo drives from a single board.

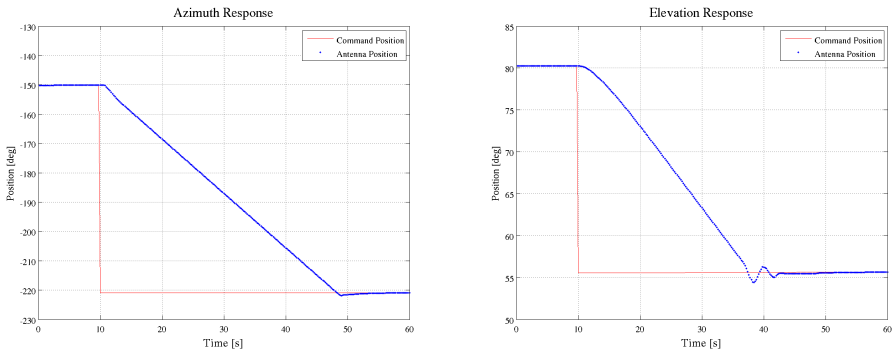

(a) (b)

Figure 3.13: Servo response to a step command. The elevation response is underdamped as we had not yet balanced the elevation axis.

The IC is interfaced to the ARM9260 development board using the standard SPI interface

Tachometer

These were built into the DC motors, and provide a DC voltage as a velocity reference. This is sampled by the ADC PCB and fed into the ARM9260 servo controller.

GPIO

The GPIO relay control and feedback is provided through a series of MAX7301 ICs. These control a number of contactors, relays etc. and sample the state of a number of mechanical switches

Angle Encoders and Interface

The encoders used were FRABA 16 bit shaft encoders. They are coupled onto the azimuth and elevation axis using a set of flexible couplings. This mitigates against first-order axial misalignments. The encoders themselves are interfaced to the ARM9260 using SSI protocol. This is similar to RS485, allowing us to make use of a number of inexpensive optically isolated RS485 drivers to prevent over-voltage or over-current conditions.

3.3.5 Servo Results

The antenna has been successfully integrated with the servo controller allowing the antenna to be controlled. This is fundamental to the results presented in Chapter 6. Figure 3.13 shows the antenna response to a step function in both azimuth and elevation.

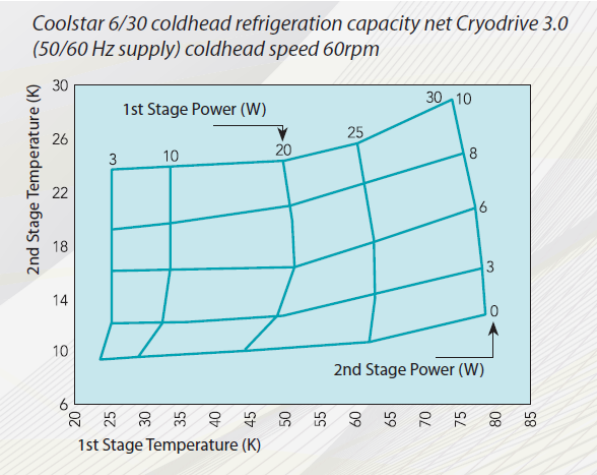

Figure 3.14: The cooling capacity of the Oxford Cryosystems 6/30 coldhead (Oxford Cryosystems, 2013)

3.4 Cryogenic Front End

The cylindrical cryostat (see Figure 3.15 for a photograph of the interior) is cooled by a two-stage 6/30 Gifford-McMahon, or GM, mechanical cooler from Oxford Cryosystems capable of cooling to 10 K on the second stage. The heat extraction capacity is given in Figure 3.14. This was a change from the 4 K Sumitomo coldhead used for the Northern system. We had originally anticipated requiring cooling of the reference load to $\sim$6 K to balance the reference power against the expected sky temperature, however our experience with the Northern system suggested that cooling the load to 10 K would be sufficient.

The front-end components are cooled to 10 K to minimize additional thermal noise. The components housed on the 10 K stage are identified in Figure 3.17. These were the orthomode transducer (OMT), the linear-to-circular convertor (L2C), the noise diode RF couplers, the 180° hybrids, the reference cold load, the cryogenic isolators, and the low noise amplifiers. Since the noise diode is critical to the system calibration, we built this unit into the cryostat on a temperature-controlled plate on the 300 K stage. The location inside the cryostat has the advantage that it maintains the integrity of the interconnects between the noise diode and the coupler.

We learned a great deal in the Northern system regarding coldhead synchronous pickup (see Section 2.3). These issues were taken into account while assembling the Southern cryostat. We spent a great deal of time in the lab commissioning, trying to identify and deal with possible causes of these microphonics, minimizing the number of connectors on the LNA bias lines, and ensuring all RF connectors were tightened. To test the effectiveness I terminated the RF inputs with 50 ohm terminations at 10K and monitored the detected output power. This

had the power spectrum shown in Figure 3.16b; in this configuration there was no evidence of a coldhead signal.

Once we added the horn optics however, we saw evidence of a raised $T_{\rm rx}$ (Figure 3.16a) and a strong microphonic signal (clearly visible in Figure 3.16c). Examining the interior of the cryostat we found that there was a 5 mm axial misalignment between the OMT and the horn, shown in the diagram of Figure 3.15. The OMT was mechanically supported and thus aligned using the cold head, and after examination it was determined that the misalignment was caused by a deformation in the coldhead cylinder, probably caused during transportation from Oxford to South Africa.

To remedy this we redesigned the mechanical support of the OMT. We mechanically decoupled it from the cold head and positioned it using four stainless steel travelling rods extended through the heatshield. The rods were thermally insulated from the heat shield using phenolic grommets preventing any uneccessary thermal leakage. These rods allowed for horizontal positioning of the OMT, while vertical positioning was provided by screws that were extended down from the cryostat lid. The OMT was kept in thermal contact with the 2nd stage using a sprung copper sheet. This provided sufficient thermal contact to cool the OMT to $\sim$14 K. This alignment modification has the added benefit of removing any possibility of lateral movement of the OMT relative to the heat shield horn section. Four stainless steel pins in the cryostat lid locating into phenolic grommets in the heatshield, provide further rigidity between the horn segments. We also added G10 rods between the 300 K and 10 K stages to provide further mechanical stiffness and prevent the possibility of further deformation of the cold head.

Diagrams and photos of cryostat both before and after are included in Figure 3.15. The modification improved the $T_{\rm rx}$ significantly as seen in Figure 3.16a and also reduced the cold head signature in the data. The improvement in $T_{\rm sys}$ is because of the improved return loss of the properly aligned OMT. Future work will include investigations of the increased $T_{\rm rx}$ in the upper half of the band.

3.5 Signal Conditioning

The C-BASS signal conditioning is straightforward. The RF 4.5–5.5 GHz band is defined by a series of gain stages and bandpass filtering. We then mix the RF band down to a baseband of DC–1000 MHz. The ready availability of 1 Gsps ADCs with 500 MHz sampling bandwidth allow us to easily sample this full 1 GHz bandwidth by splitting the 1 GHz band into two sub-bands, a DC–500 MHz sub-band corresponding to the lower half of the RF band, and a 500–1000 MHz sub-band corresponding to the upper half of the RF band. The signal conditioning

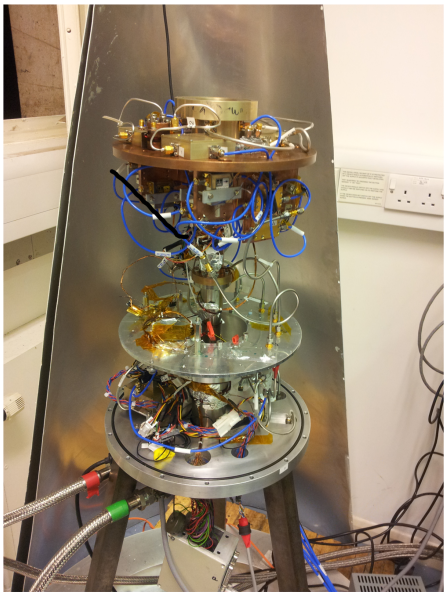

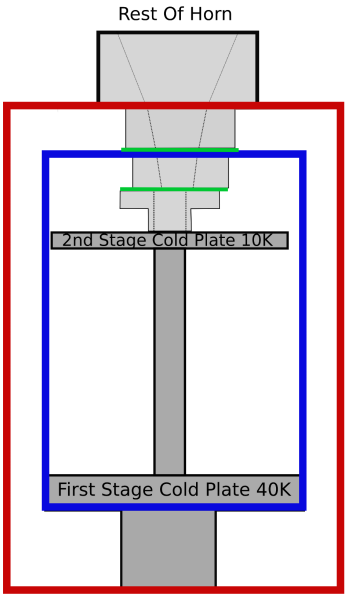

(a) Cryostat before modifications.

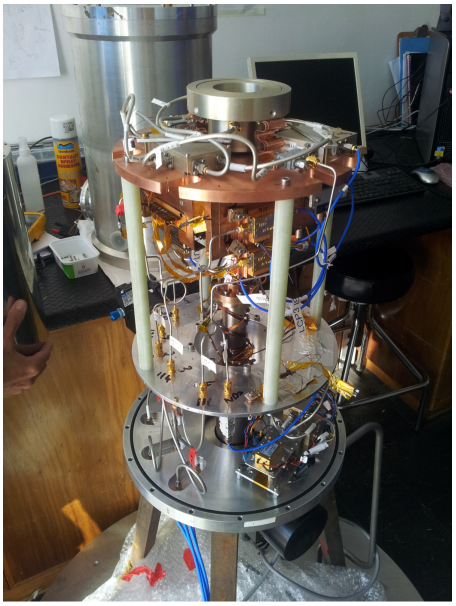

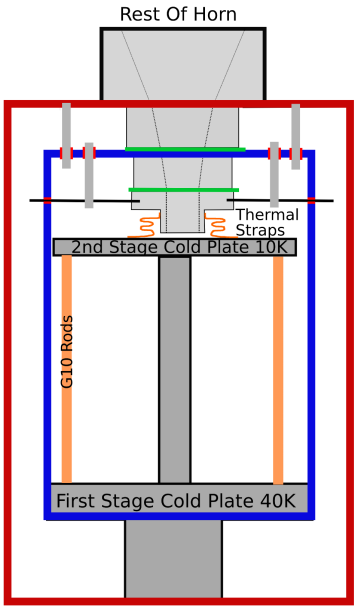

(b) Cryostat after modifications.

Figure 3.15: The cryostat was modified to allow more accurate alignment of the OMT to the external optics, and to rigidify the interior with G10 rods.

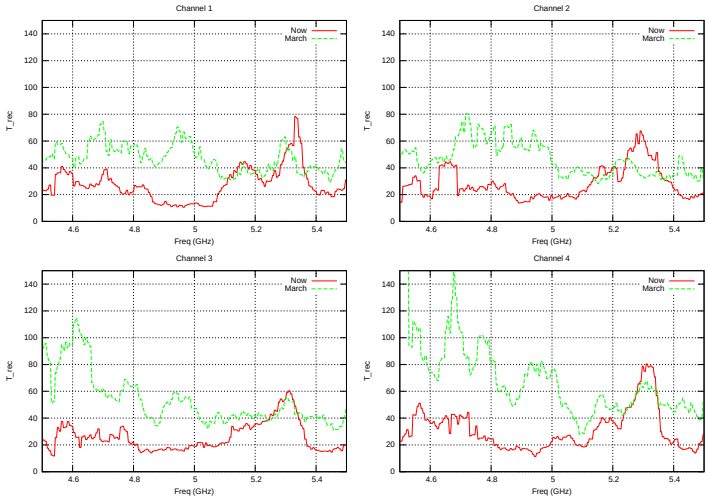

(a) The T_{rx} calculated using spectrum analyser measurements.

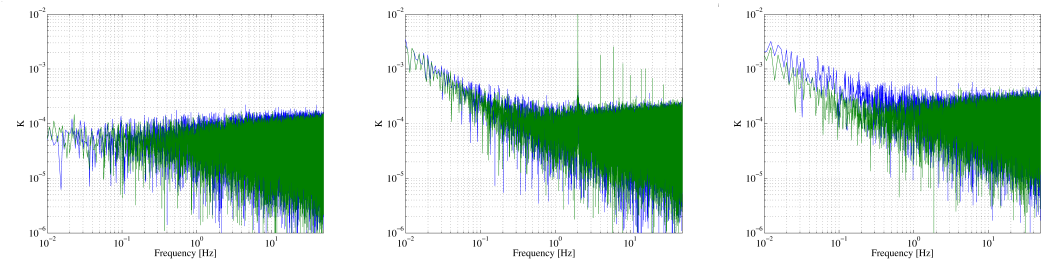

(b) RF chain terminated with reference load at 12.000 K

(c) OMT Misaligned

(d) OMT Aligned

Figure 3.16: Effect of correcting the OMT alignment on detected signal. In (a) The green line shows the T_{rx} before the OMT alignment was corrected, while the red line shows the improved T_{rx} after the correction. In (b), (c), and (d) I have plotted the signal power spectrum in different states. In (b) we see that the coldhead does not affect the signal, while (c) shows the strong coldhead signal present when the optics are added but the OMT is misaligned. Aligning the OMT improves - but does not completely remove- this coldhead signal. Future work will include diagnosing the increased T_{rx} in the upper half of the band. The $1/f$ component in (c) and (d) is due to power fluctuations from outside the cryostat coupled to the system through the feed horn.

is illustrated in Figure 3.17, while the digital signal processing after digitization is described further in Section 3.7.

A great deal of effort was put into maintaining signal fidelity. The amplifiers were individually regulated to minimize cross-coupling between power supplies. All components were firmly bolted onto temperature-controlled aluminium sheets and these aluminium sheets were also used as the ground return for all the active components in an effort to prevent the formation of ground loops in the DC power distribution.

3.6 Noise Diode Signal Calibration

The signal is calibrated using a temperature stabilized noise diode. This is described in more detail in the following chapter, Section 4.2. The noise diode calibration signal is split into identical signals for the right and left circular RF branches, and both are combined with the respective arms using 30 dB directional couplers. A novel feature of the calibration signal is the addition of 6-bit digital phase shifters to both of the calibration signal paths. By shifting the phase it is possible to change the noise diode signal from a purely Stokes Q signal when the noise diode calibration signal has identical phase in both left and right polarizations, to a purely Stokes U signal when there is a 90° phase difference between the noise diode signal in left and right circular polarizations. This is an improvement over the analog instrument which only allows Stokes Q to be calibrated with the noise diode.

3.7 Digital Backend

A number of recently deployed instruments (e.g. Allen Telescope Array (Welch et al., 2009), KAT-7/MeerKAT (Jonas, 2009), Low-Frequency Array for radio astronomy, or LOFAR (de Vos et al., 2009), Precision Array for Probing the Epoch of Reionization (PAPER) (Parsons et al., 2010)) have demonstrated the significant advantage of the flexibility provided by a generic and scalable signal processing architecture based around FPGAs or GPUs (de Vos et al., 2009, Parsons et al., 2008, Wawrzynek & Brodersen, 2005). This approach allows instruments to be developed on rapid timescales, reducing the time-to-science (Parsons et al., 2010), and allowing designers to take better advantage of Moore's Law (Moore, 1965). We were able to leverage these developments for the Southern C-BASS system.

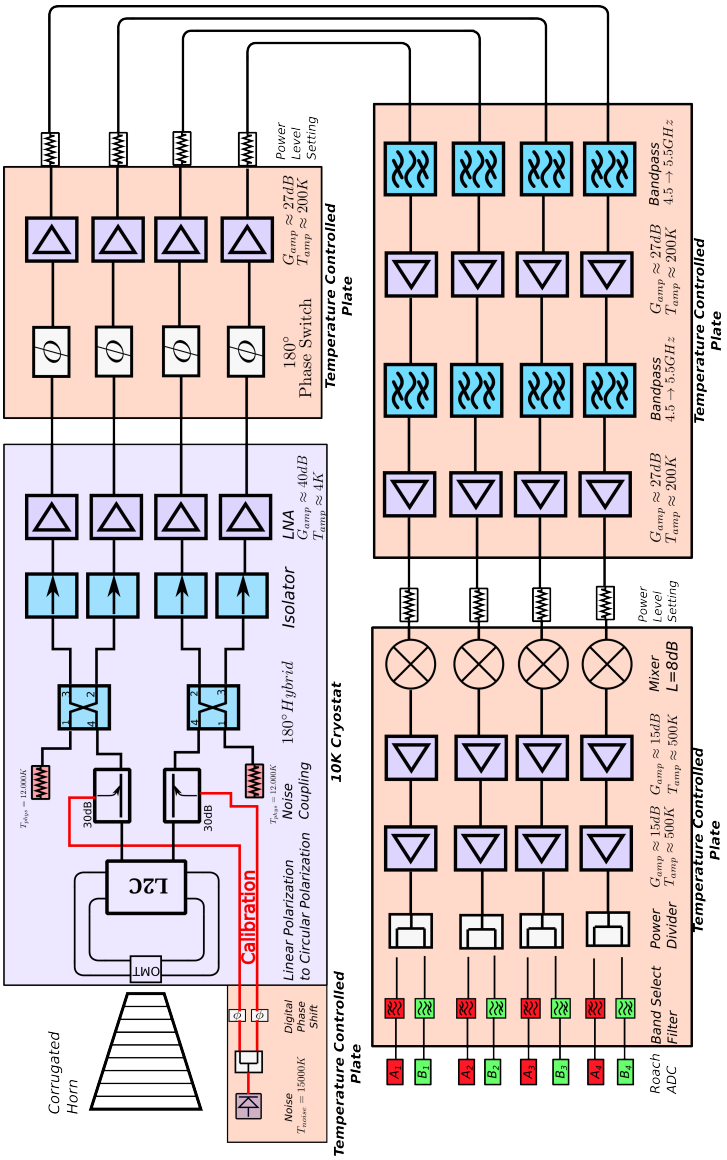

Figure 3.17: Diagram of the analog front end of the digital receiver. There are significantly fewer analog components than the Northern Hemisphere polarimeter of Figure 2.1. The diagram shows signal conditioning, and the down-conversion of the 1 GHz band (into two 500 MHz bands) followed by the ADC capture. For C-BASS we use a mixing frequency of 4.5 GHz but this could easily be changed for projects interested in other bands. The digital processing is shown on Figure 3.19.

3.7.1 Collaboration for Astronomy Signal Processing and Electronics Research

The Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) is a worldwide collaboration that provides both hardware and software for the radio astronomy instrumentation community. The aim is to reduce the time required to develop astronomical instrumentation. Hardware platforms include a number of high speed ADCs, 8 bit DACs and FPGA development boards. The software libraries include highly optimized fast Fourier transforms (FFTs), Polyphase filter banks, PFBs, and other useful digital signal processing tools. These are all packaged into the *MATLAB Simulink* environment. Using these libraries together with the FPGA hardware platforms, it is often possible to reduce instrumentation development timescales from multi-year to multi-month.

3.7.2 ROACH

The ROACH board in Figure 3.18a, forms a part of of the CASPER hardware suite. The board is based around a Xilinx Virtex 5 FPGA and provides a versatile platform for digital signal processing. The FPGA is accessed through a Power PC, operating a light-weight linux-based operating system, the Berkeley Operating System for Re-programmable Hardware, or BORPH. The FPGA programming is done through BORPH, and the FPGA registers and memory are mapped to devices in the operating system. This allows for remote programming of the FPGA hardware, and also allows users to access the FPGA hardware through a familiar environment.

3.7.3 iADC

The iADC, again in Figure 3.18a, is a multi-purpose CASPER ADC convertor that uses the Atmel/e2V AT84AD001B IC. In non-interleaved mode a single ADC is capable of sampling two signals at 1 Gsps providing 500 MHz of RF bandwidth for the two sampling channels. As each ROACH can host two iADC boards, it is possible to sample four 500 MHz RF bandwidth on each ROACH. For the 1 GHz C-BASS bandwidth, we used two ROACH boards, hosting four iADCs. The band was split into a lower 500 MHz (processed by the first ROACH) and an upper 500 MHz (processed by the second ROACH).

3.7.4 Digital Signal Processing

The digital signal processing is done on the ROACH board, and consists of the steps outlined in Figure 3.19. The logical operations are programmed into the FPGA using the CASPER

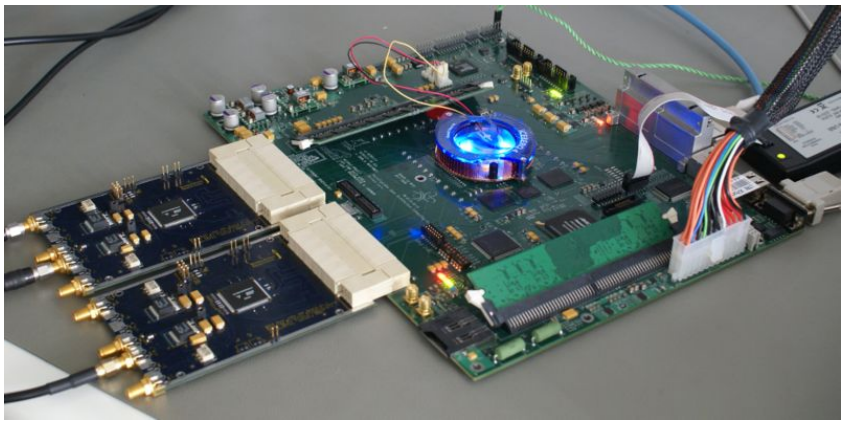

(a) ROACH board. The two smaller boards on the left are the iADC boards. From https://casper.berkeley.edu/wiki/ROACH

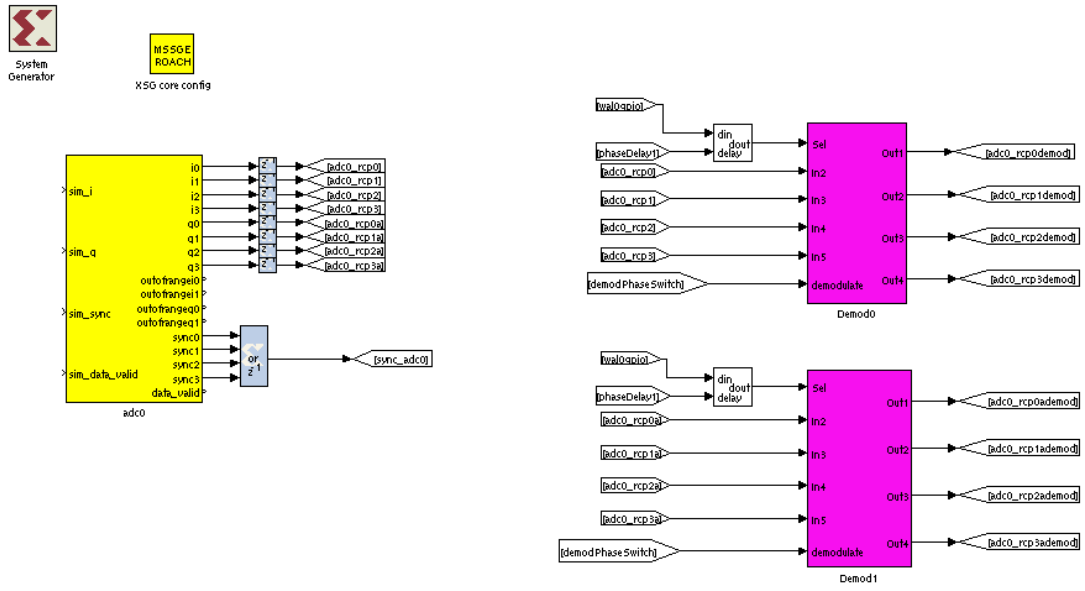

(b) CASPER toolflow in *Simulink*.

Figure 3.18: Typical CASPER hardware and design. The CASPER collaboration provides both hardware and an associated toolflow. The open-source model allows standard software to be reused, allowing for much faster instrument development.

toolflow, a suite of generic FPGA digital signal processing tools programmed in Simulink. A screenshot of a typical design is shown in Figure 3.18b.

Sampling at 1 Gsps on the iADCs requires running the ROACH FPGA at 250 MHz. This is unusually fast within the CASPER community where a fast clock is typically considered to be 200 MHz, and was especially challenging because of the large size of the design, which complicated the signal routing across the FPGA fabric. This was dealt with by identifying the routes that were failing to meet timing using the *Xilinx Timing Analyser* (a part of the Xilinx FPGA toolflow), and adding delays in appropriate places to relieve pressure on the route optimizer. An alternative approach would have been to manually route the FPGA using timing tools such as *Planahead*. I found these to be unneccessary.

The digital backend is comprised of six major stages:

1. Baseband analog sampling
2. Polyphase filter bank/ fast Fourier transform
3. Gain/Phase Adjust and Sky/Load Separation
4. Stokes I, Q, U detection

Accurate timestamping of the data is provided using a 1 pulse-per-second generated using an inexpensive Thunderbolt GPS receiver from Trimble GPS. We also phase-lock the ADC sampling clock to GPS using the 10 MHz GPS reference signal to minimize long-term clock drift.

The rest of this section describes these different stages, and how they were implemented on the ROACH board.

3.7.5 Baseband analog sampling

We introduced the iADC in Section 3.7.3. For the C-BASS continuous-comparison architecture, we require 4 channel sampling, with a total bandwidth of 1 GHz. Each ROACH can host two iADCs, capable of 8-bit sampling four analog signals at 1 Gsps; so each ROACH can sample 500 MHz of C-BASS bandwidth.

In the Southern C-BASS receiver we mix the 4.5–5.5 GHz RF band to a DC–1000 MHz baseband, and split this into 2x500 MHz sub-bands for digitization. These sub-bands are digitized by parallel iADC/ROACH systems providing the full 1 GHz sampling.

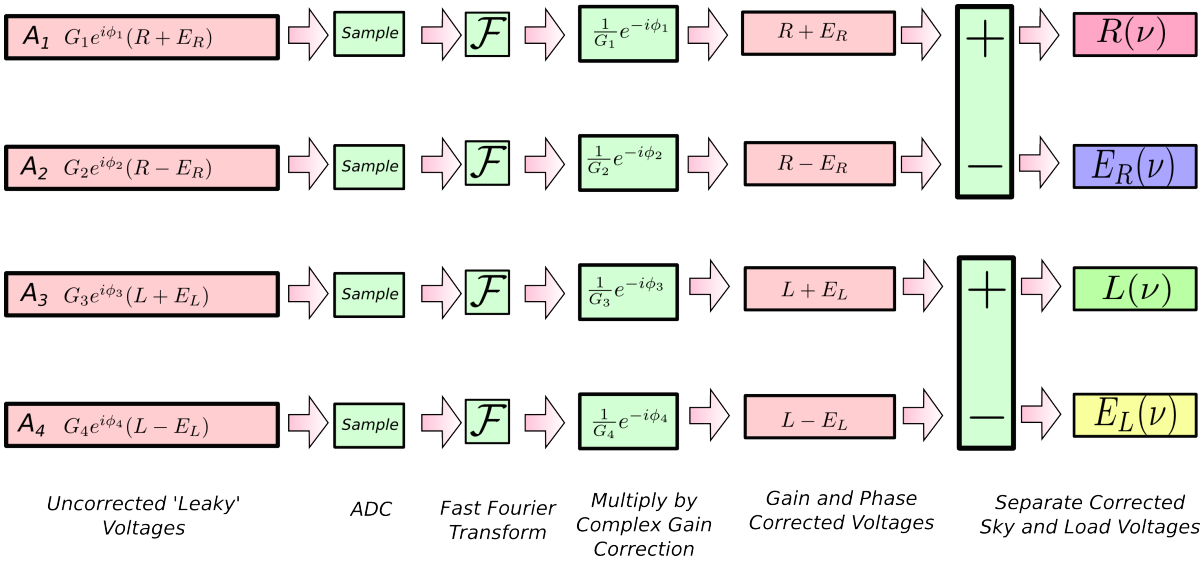

(a) The uncorrected signals from the analog receiver (see Figure 3.17), are sampled, polyphase filtered and fast Fourier transformed, and then gain and phase adjusted. The sky and load voltages are then separated from each other.

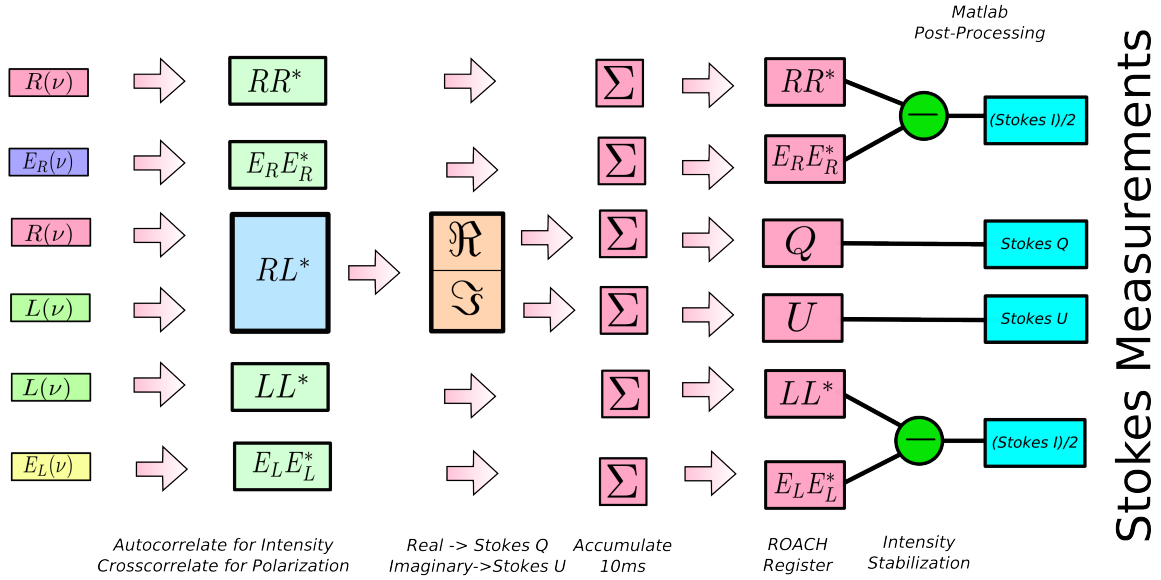

(b) Next we perform the auto-correlations (for Stokes I) and cross-correlations (for Stokes Q and U). These are accumulated to produce each 10 ms time integrated sample.

Figure 3.19: The digital signal processing. See Equations 1.11 and 1.14 for the derivation of the Stokes Parameters calculations.

3.7.6 Polyphase Fast Fourier transform

The novel aspect of the Southern C-BASS receiver is the digital implementation of a spectral backend to support the continuous-comparison architecture and correlation polarimeter introduced in Section 1.8.3 and Section 1.8.4. The main reason for doing this is to provide immunity to RFI, avoiding the costly process of analog filtering required for the C-BASS North receiver.

This begs the question: what level of immunity to RFI can we expect? In signal processing we commonly use the Fast Fourier transform (FFT) to efficiently convert from the time sample domain, $x(n)$, to the frequency domain, $X(k)$. The transform is given by

$$X(k) = \sum_{n=0}^{N-1} x(n) e^{-2\pi nk/N}. \tag{3.3}$$

In order to estimate our immunity to RFI, we need to calculate the response, $X(k)$ to a typical RFI event. In the case of RFI, the time domain signal, $x(n)$, is typically a sinusoid. In Equation 3.3 we sample this across a square window of size N. Except for special instances where the period of the sinusoid matches the window size, there will be a discontinuity at the edge of each sample window, and this will result in *spectral leakage* (Harris, 1978)

Another way to think of this is in terms of the Fourier transform pairs. In the time domain we have two signals, the sinusoid and the square window. The Fourier transforms of these are a delta function and a sinc function respectively. Multiplication of the sinusoid and the square window in the time domain is the same as convolving a delta function and a sinc function in the frequency domain. The transform, $X(k)$, will be then be a sinc function centred at the location of the delta function; most of the power will be concentrated in the bin associated with the frequency of $x(n)$, however there will be leakage into adjacent frequency bins caused by the frequency response to the square window i.e. the sinc function.

In order to minimize this leakage in $X(k)$, we can apply a filter $h(n)$ to the sampled data, $x(n)$, before taking the FFT so that

$$X(k) = \sum_{n=0}^{N-1} x(n) h(n) e^{-2\pi nk/N} \tag{3.4}$$

These window functions are often called taper functions as they reduce the leakage across frequency channels by minimize the discontinuity at the edge of the sample window. An infinite window length would allow us to completely suppress the discontinuity. Given the limited FPGA resource this is not practical and windows are truncated at the edges; however the larger the window, the better the reduction in spectral leakage.

Polyphase Filter Banks

The FFT size is defined by the required frequency resolution, and a typical window would be the same size as the FFT. As I have mentioned though, a larger window will lower the spectral leakage. In non-resource limited world we would make the FFT larger, apply the larger window function and discard the unnecessary frequency bins. This is resource-intensive though, and it turns out that it is possible to avoid doing this by using the technique of time-aliasing. This allows us to to both enlarge the window size and benefit from the improved spectral leakage, while simultaneously maintain the FFT size (Dahl, 2003, Lyons, 2004).

To take advantage of this, we start with a large window of data, $y[q]$, of size $M \times N$. We subdivide this into M smaller windows, and call these $x_i[n]$. These are each of size N.

we introduce a quantity, $x^{(\alpha)}$, which is given by the sum of the M smaller windows,

$$x^{(\alpha)}[n] = \sum_{i=0}^{M-1} x_i[n] \tag{3.5}$$

This operation is known as time-aliasing since we can also express the quantity $x^{(\alpha)}$ using indexing in the large data window, $y[p]$, using

$$x^{(\alpha)}[n] = \sum_{m=0}^{M-1} y[n + mN] \tag{3.6}$$

The FFT of this series, $X^{(\alpha)}[k]$, is given by

$$X^{(\alpha)}[k] = \sum_{n=0}^{N-1} x^{(\alpha)}[n] e^{\frac{2\pi kn}{N}}, \tag{3.7}$$

while the FFT, $Y[p]$ of the larger sample window, $y[q]$ is given by

$$Y[p] = \sum_{MN=0}^{MN-1} y[q] e^{\frac{2\pi pq}{MN}}, \tag{3.8}$$

It can be shown (Dahl, 2003) that the two FFTs are related by,

$$X^{(\alpha)}[k] = Y[Mp], \tag{3.9}$$

We see that time-domain aliasing is equivalent to decimation in the frequency domain. Time-domain aliasing allows us to take advantage of a large window function, while not requiring larger FFTs. Dahl (2003) shows the various savings that can be achieved in Chapter 2.8. Polyphase filterbanks are an efficient way of taking advantage of this property. A

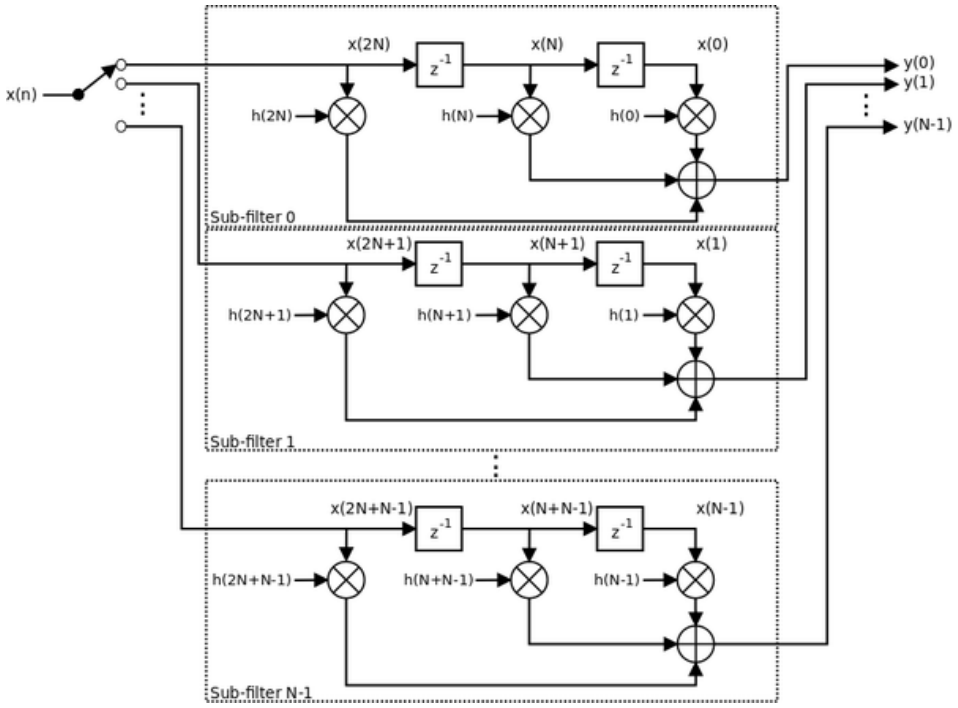

Figure 3.20: A 3-tap polyphase bank with N channels. The 3-taps allow a window function of three times the FFT size to be efficiently implemented using the time aliasing technique. From `https://casper.berkeley.edu/wiki/The_Polyphase_Filter_Bank_Technique`.

digital implementation is given in Figure 3.20.

Expected Sidelobe Suppression

The C-BASS backend uses a $N = 128$ point FFT with a $N = 512$ point Hamming window function given by

$$\omega(n) = \alpha - \beta cos\left(\frac{2\pi n}{N-1}\right) \tag{3.10}$$

with $\alpha = 0.54$,, $\beta = 1 - \alpha = 0.46$.

The Hamming window is applied using a polyphase filter bank architecture.

We can simulate the expected response of both an un-windowed FFT and the Hamming windowed FFT. Both of these are illustrated in Figure 3.21. We expect a minimum cross-channel isolation of 50 dB using this configuration. The flat shape of the response across the band minimizes the 'scalloping loss' between the individual frequency bins and is far superior to the scalloping loss of a standard FFT-based channelizer.

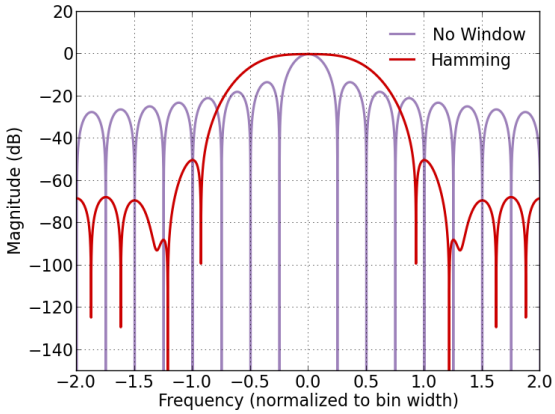

Figure 3.21: The un-windowed and Hamming windowed frequency response. C-BASS uses an $N = 128$ size FFT with 4-tap, Hamming Windowed polyphase filter bank. We expect 50 dB channel-to-channel isolation using this windowing function. 'Scalloping loss' is minimized by flattening the response across the channel.

3.7.7 Gain/Phase Adjust and Sky/Load Separation

The four RF channels were originally produced by coupling each of the sky polarization voltages - R and L for right and left polarization respectively - with the reference temperature voltages - (E_R and E_L for the right and left polarization references respectively - the using the 180° hybrid inside the cryostat. If the combined signal, C, of the sky and load were sine waves of the form $C(x,t) = sin(\omega t - kx + \phi)$ we would expect the voltage of the four RF channels to be described by

$$C_1(x_1,t) = G_1(R(\omega t - kx_1) + E_R(\omega t - kx_1 + 0))) \quad (3.11)$$
$$C_2(x_2,t) = G_2(R(\omega t - kx_2) + E_R(\omega t - kx_2 + 180°)) \quad (3.12)$$
$$C_3(x_3,t) = G_3(L(\omega t - kx_3) + E_L(\omega t - kx_3 + 0)) \quad (3.13)$$
$$C_4(x_4,t) = G_4(L(\omega t - kx_4) + E_L(\omega t - kx_4 + 180°)) \quad (3.14)$$

where $k = 2\pi/\lambda = 2\pi f/c$ and $G_i, i = 1...4$ are the RF channel gains.
Making use of the fact that $sin(0) = -sin(180°)$ we can rewrite this as follows:

$$C_1(x_1,t) = G_1(R(\omega t - kx_1) + E_R(\omega t - kx_1)) \quad (3.15)$$
$$C_2(x_2,t) = G_2(R(\omega t - kx_2) - E_R(\omega t - kx_2)) \quad (3.16)$$
$$C_3(x_3,t) = G_3(L(\omega t - kx_3) + E_L(\omega t - kx_3)) \quad (3.17)$$
$$C_4(x_4,t) = G_4(L(\omega t - kx_4) - E_L(\omega t - kx_4)) \quad (3.18)$$

If the electrical lengths $x_1 = x_2$ and $x_3 = x_4$, and we can trivially disentangle the sky voltage and load voltage by addition and subtraction. This is done using another 180° hybrid in the Northern C-BASS receiver, and is done digitally in the digital C-BASS receiver:

$$A = (C_1 + C_2)/2 = (G_1 + G_2)R + (G_1 - G_2)E_R \quad (3.19)$$

$$B = (C_1 - C_2)/2 = (G_1 + G2)E_R + (G_1 - G_2)R \quad (3.20)$$

$$C = (C_3 + C_4)/2 = (G_3 + G_4)L + (G_3 - G_4)E_L \quad (3.21)$$

$$D = (C_3 - C_4)/2 = (G_3 + G4)E_L + (G_3 - G_4)L \quad (3.22)$$

Since the gains G_i are similar, $(G_1+G_2) \gg (G_1-G_2)$, and the second terms are negligible.

The world is not perfect however. The physical lengths of the RF signal paths $x_1 \neq x_2$ and $x_3 \neq x_4$ since small path differences exist in the cascaded RF components. These path length differences create a frequency dependent phase mismatch in the four RF signals, resulting in 'leakage' of the sky voltage into the load voltage in Equation 3.22.

This 'leakage' is corrected in the Northern system by adjusting the physical lengths of the coaxial interconnects, such that $x_1 = x_2$ and $x_3 = x_4$. This is a time-consuming process, and on the Southern receiver we can avoid this by correcting the signals digitally. This can either be done by adding variable digital delays after the ADC sampling, or by using the Fourier relationship given in Equation 3.23 and applying the equivalent frequency dependent phase shift in the digital domain.

$$F(S)e^{-2i\pi\phi S} \Leftrightarrow f(x - \phi) \quad (3.23)$$

I chose to do the latter, and after channelizing the data using the PFB/FFT, we simply phase shift each channel appropriately using Equation 3.23.

The phase corrections are calculated by assuming that the optimal set of phase corrections maximizes the noise diode contribution to the sky power.

In addition to the phase correction, gain corrections are also applied at this stage, allowing the passband to be flattened.

3.7.8 Stokes I, Q, and U detection

After the previous step, we are left with FFTs of the voltages of the two sky polarizations and the FFTs of the two reference load voltages. We can square the sky voltages to calculate the two sky temperatures LL^* and RR^*, and square the load voltages to get the two load temperatures, $T_L T_L^*$ and $T_R T_R^*$ for the left and right polarizations respectively. Together these

can provide the stabilized Stokes I. The linear polarizations, Stokes Q and Stokes U can be calculated by multiplying the two sky polarizations.

These relationships were originally defined in Section 1.7. They are repeated here for ease of readability:

$$I = (LL^* - E_L E_L^*) + (RR^* - E_R E_R^*) \tag{3.24}$$

$$Q = 2\mathbb{R}(LR^*) \tag{3.25}$$

$$U = -2\mathbb{I}(LR^*) \tag{3.26}$$

The relationships above show us what we need to record for the C-BASS South receiver. These are the following six measurements,

1. LL^* and RR^*
2. $E_L E_L^*$ and $E_R E_R^*$
3. Q
4. U

These are accumulated in 10 ms integrations and recorded by the integrated control system described in Section 3.8. The resulting data are shown in Figure 3.22.

3.7.9 Mains Frequency and other Low Frequency Contaminants

A significant problem in scanning single-dish observations is low-frequency pickup. The signal of interest from the sky is a convolution of the antenna beam with the sky. For C-BASS scanning at 4°/s, the 0.7° beam means the sky appears at frequencies of order 6 Hz and lower. Any artificial signals near this frequency are exceedingly difficult to distinguish from the sky signal. A good example is the well known 'mains hum'- in the U.S. mains is at 60 Hz. When sampling at 100 Hz, the first harmonic at 120 Hz is aliased down to 20 Hz, uncomfortably close to our sky signal. This proved to be a significant and challenging problem to address on the Northern survey. The digital C-BASS architecture completely avoids this problem. Since both channels are sampled directly by ADC's, and low frequency components will be common to both systems and will cancel in the same way as $1/f$ noise is reduced. We see this when looking at Figure 3.22a where low frequency contaminants in the unstabilized power spectrum (red), are removed in the the stabilized power spectrum (blue).

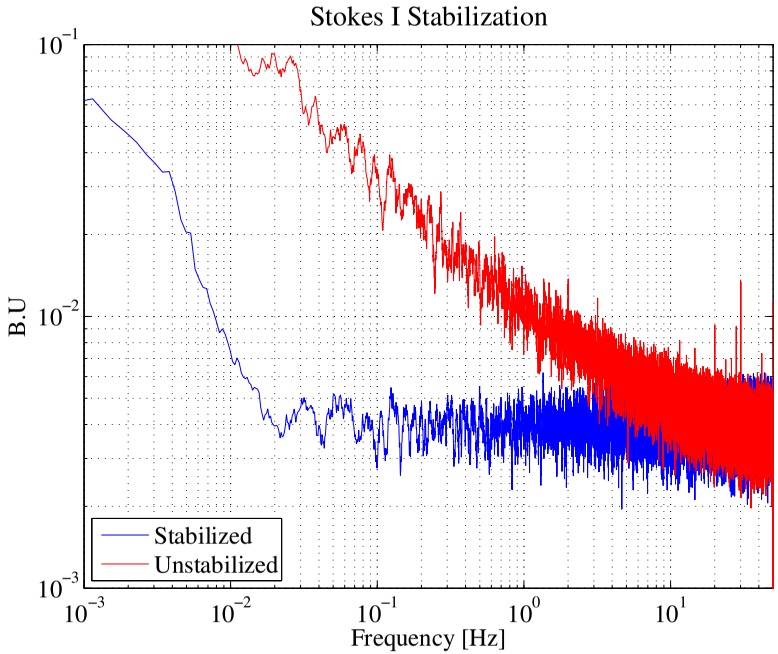

(a) The unstabilized amplifier (red), suffers from low frequency noise as a result of amplifier gain drifts. In the continuous-comparison receiver (blue), the receiver chain is stabilized by comparing against a stabilized reference load. This reduces $1/f$ noise at the expense of increasing the white noise power. The (digitally implemented) continuous-comparison method has the added benefit of removing any low frequency signals picked up in the RF chain. This is clearly seen as the spikes in the unstabilized (red) signal, are absent in the stabilized (blue) signal.

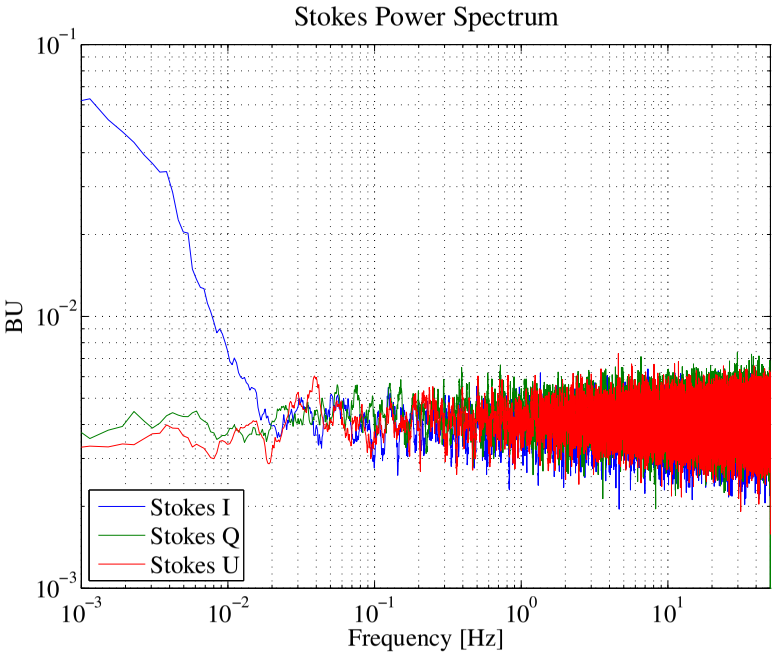

(b) The power spectrums of Stokes I, Stokes Q and Stokes U as detected by C-BASS during lab testing. The polarization measurements are made by correlating the two circular polarizations, a measurement insensitive to the uncorrelated amplifier gain drifts. They do not require any stabilization mechanism to reduce the $1/f$ noise.

Figure 3.22: Power spectra of different C-BASS Stokes measurements.

3.8 Integrated Control System

The antenna subsystems are integrated using a control suite originally designed for the Cosmic Background Imager (CBI) by Martin Shepherd (Caltech). This was significantly altered for use by C-BASS North by Stephen Muchovej. Porting the code over for to the Southern antenna servo system was done by Stephen and myself over a three week period in October 2010

Further additions were added over the course of the receiver development. The control suite now provides data acquisition and control support for all of the antenna subsystems, as well as providing a number of extremely useful general features, such as

1. ROACH Backend Data Acquisition
2. Environment temperature monitoring
3. Cryogenic temperature monitoring
4. Cold load temperature control
5. Calibration noise diode control
6. Frame grabber integeration, which allows automated optical pointing
7. Servo control
8. Pointing model
9. Weather data
10. Automated schedule run
11. Realtime plotting
12. LNA Bias Monitoring

3.9 Klerefontein Site Establishment

The Southern C-BASS survey is currently being commissioned at a test facility at HartRAO at a location S$25^\circ 53.458'$ E$27^\circ 41.160'$, altitude $\sim$1400 m. The deployment will be to the final site near the town of Klerefontein in the Northern Cape Province S$30^\circ 58.290'$ E$21^\circ 58.817'$, altitude $\sim$1300 m. The exact locations are in Figure 3.23 and can be located in *Google Earth* using the KMZ file located at `http://goo.gl/gmB9mx`. We were able to carefully select this site and optimized the location for science, choosing a elevated site with a clear horizon to prevent

Figure 3.23: The two South African C-BASS locations. C-BASS will be initially commissioned at the HartRAO site near Pretoria, and will then move to the Klerefontein site in the Northern Cape Province. Image from *Google Earth*. KMZ file available at `http://goo.gl/gmB9mx`

any azimuth dependent ground pickup. The disadvantage is that the instrument is exposed to a much higher risk of RFI, however the digital backend and the radio-quiet legislation enacted to protect the Klerefontein vicinity from artificial terrestrial RFI are likely to mitigate against the impact RFI can have on the Southern survey.

3.9.1 Infrastructure

The antenna had a number of infrastructural requirements including power, data connectivity, weather information, road access, and foundations. This sections describes these issues.

Power Provision

The helium compressor requires 3kW and the four drive motors require 12 kW total. We budgeted for a total power requirement of 20 kW. The C-BASS antenna is 500 m from the nearest high voltage transformer. The high power requirement combined with the long cable run can cause a significant voltage drop. This can either be mitigated by reducing the cable resistance by increasing the copper diameter, or by running high voltage lines and stepping-down the voltage at the antenna using a transformer. We chose the prior to avoid the need for a transformer near the antenna, eliminating any potential broad-band RFI in close proximity to the antenna. The cable required to do this was expensive.

Optical Fibre

Single-mode optical fibre connectivity provides communication access to the C-BASS antenna. This was run directly from the Klerefontein support base to the C-BASS container and provides network connectivity to the site.

Access

Transportation access to the site is critical. For the C-BASS installation we built a concrete road to allow us to transport the antenna to the final location.

Control Shelter Considerations

The C-BASS control shelter is located under the primary reflector. It was specifically designed to be out of the antenna sidelobes to prevent any reflection. As such the container is fairly compact.

Weather Conditions and Monitoring

The weather conditions at the site can be characterized as dry and stable throughout the year as seen in Figure 3.24. We will be installing a weather station that will provide realtime access to humidity, temperature and wind speed/direction.

3.9.2 Current Status

The Northern Cape site is complete. The antenna has been successfully installed (see Figure 3.25) with both power and data connectivity. The antenna drive system has also been tested using a copy of the servo controller described earlier. After completing the commissioning period at HartRAO we will move the survey to the Northern Cape to take advantage of the optimal survey conditions there.

3.10 Conclusion

This chapter has introduced the Southern C-BASS instrument and I have outlined the instrument at a high level. I have described the photogrammetry measurement technique and also briefly evaluated the accuracy of the measurements using optical observations. I have described all the subsystems including the optics, servo controller, and the novel C-BASS digital receiver. The C-BASS antenna system has been successfully integrated with the same high level control system as that used for the Northern survey, which will facilitate re-use of both

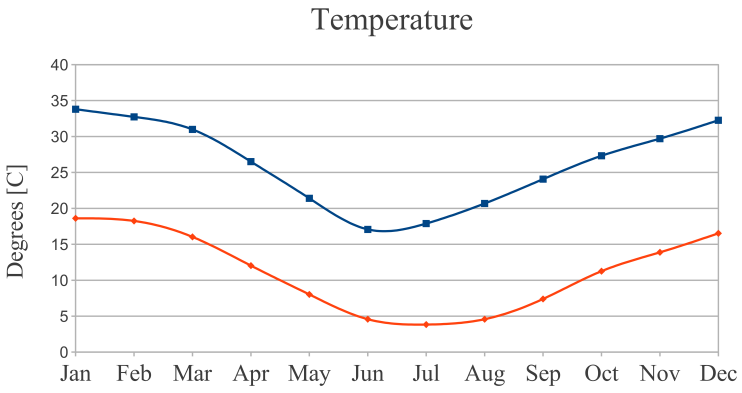

(a)

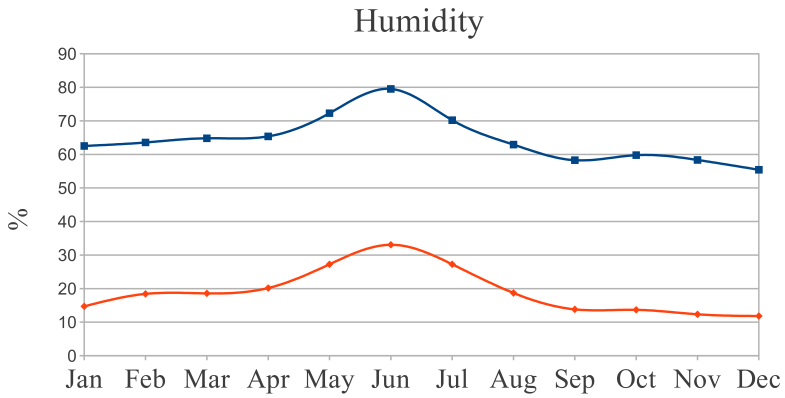

(b)

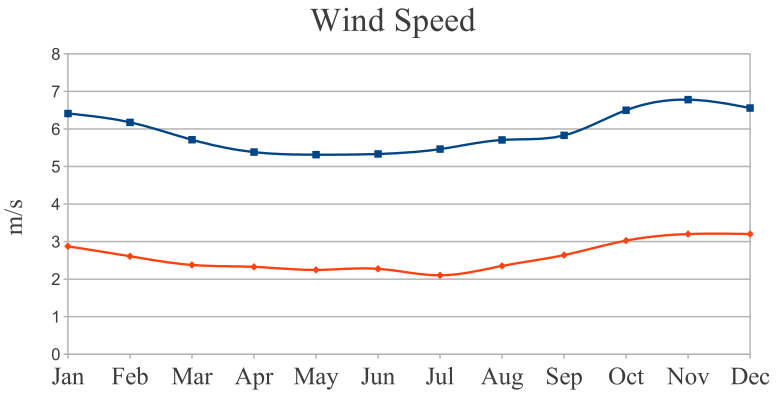

(c)

Figure 3.24: Average yearly weather conditions at the Klerefontein site, average maximum in blue and average minimum in orange. From private communication with Dr. Adrian Tiplady, Square Kilometer Array South Africa Site Bid Manager.

Figure 3.25: The final destination for C-BASS South at Klerefontein together with a number of the HartRAO team involved in establishing the site and installing the antenna.

survey schedules as well as data reduction pipelines. The status of the final site in the AGAA, Northern Cape Province has been described; it is ready for deployment and is in a perfect location for C-BASS with minimal ground horizons and a stable weather pattern. A key aspect of this chapter is the applicability of a number of the techniques and systems to other, future, antenna upgrades.

In the following chapter, I examine and characterize the individual components used to build the instrument. Most of these components are wide bandwidth allowing the bandwidth of the receiver to be easily increased in the future, either for C-BASS or for other similar projects.

Chapter 4

The Southern C-BASS Receiver Component Design

> I built it one piece at a time and it didn't cost me a dime,
> You'll know it's me when I come through your town.
>
> Johnny Cash

The C-BASS science case requires a cutting-edge instrument, and no single subsystem is more important in achieving these goals than the receiver. The C-BASS receiver needs to provide immunity to $1/f$ gain variations implicit to amplifier technology, as well as meet stringent sensitivity requirements. In the previous chapter, I discussed the high-level receiver design; in this chapter I focus on the development and characterization of a number of components designed and built specifically for C-BASS. Despite their bespoke nature, a number of these components would be useful for general purposes in other radio receivers.

4.1 Temperature controlled resistor

The reference load needs to provide both a well defined and stable temperature for the continuous-comparison measurement of Stokes I. The Stokes I measurements are made relative to this temperature reference, and any variation in the temperature impacts directly on their quality. This section describes the development of a suitable reference load for the C-BASS receiver.

4.1.1 Temperature controlled bobbin

Our first attempts were made during the development and deployment of the Northern receiver. Initially we tried to stabilize the temperature of a small copper block, or 'bobbin' seen in

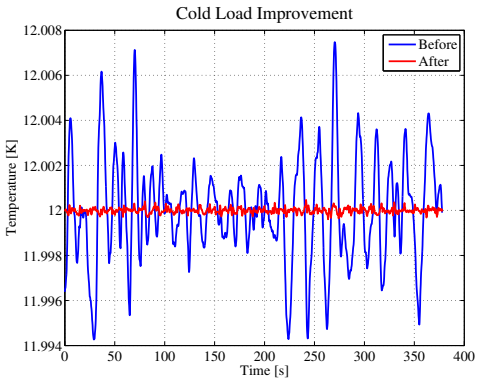

(a) Bobbin temperature with time before and after the modifications.

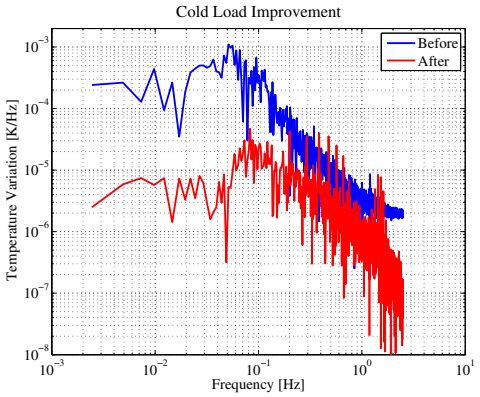

(b) Temperature spectrum before and after the modifications.

Figure 4.1: Temperature stability of the new cold load.

Figure 4.2a, which we had firmly bolted onto the second stage cold plate. The temperature was controlled by heating the bobbin, by passing a current through high resistance wire wrapped tightly around the block. We controlled the current using the Cryocon 32B PID temperature controller with feedback from a Cernox temperature sensor[1].

The choice of temperature sensor is important as it needs to be suitable for use at the cryogenic temperatures, provide a highly accurate temperature, and have low thermal response time to provide the required measurement bandwidth to the controller. The Cernox temperature sensor has an accuracy of $\sim$5 mK and a thermal latency of $\sim$50 ms, providing an ideal temperature feedback to the Cryocon 32B controller. The temperature sensor was mounted inside a hole drilled in the bobbin (see Figure 4.2b) using the General Electric *IMI 7031* varnish to provide a good thermal contact. This, together with an alternative *Epibond 121* epoxy, was suggested by personal correspondence with Lakeshore representatives.

The Northern analog receiver was originally deployed in this configuration. However, during the Southern receiver commissioning, further investigation found that this was suboptimal. The bobbin is exposed to the short-timescale, $\sim$1 Hz, temperature variations associated with the coldhead cooling cycle, a fact that was originally used to explain the coldhead signature in the data. In preference, we decided to thermally decouple the bobbin from the cold stage using nylon washers, decreasing the effect of short term coldhead temperature variations. This dramatically improves the temperature stability as is clearly demonstrated in Figure 4.1. The reduced cooling capacity on the bobbin, has only marginally affected the minimum achievable temperature of the reference load.

[1]See `http://www.lakeshore.com/Documents/LSTC_Cernox_l.pdf`

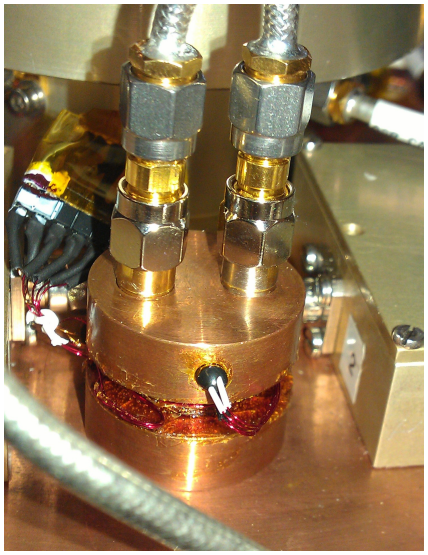

(a) Bobbin 1- Commercial 50 ohm load friction fitted into bobbin.

(b) Bobbin 2- Commercial 50 ohm load soldered into bobbin.

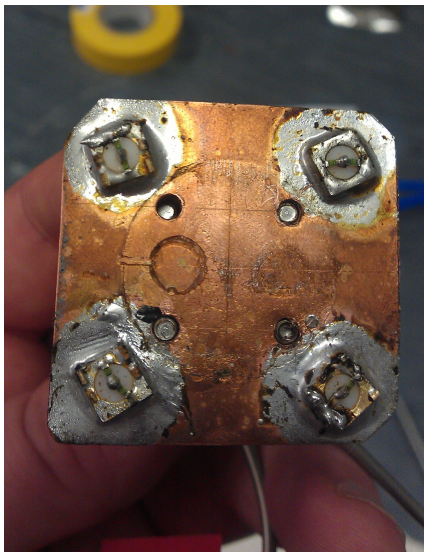

(c) Final Design- Two 100 ohm resistors soldered directly to an SMA connector following a suggestion from Eddie Blackhurst, Manchester. Four SMA connectors are soldered into a 200μm copper sheet.

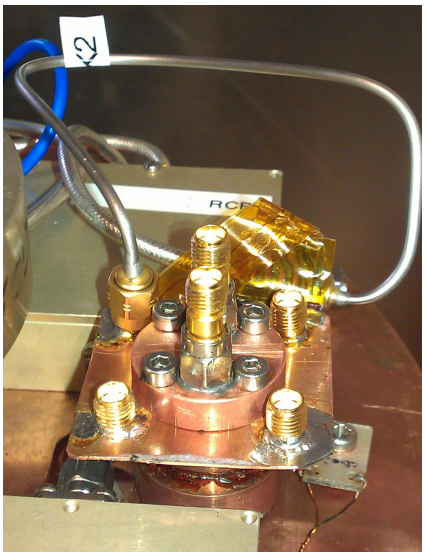

(d) Final Design- The copper sheet is thermally bonded to bobbin, and installed on the cold plate/ Only two of the four loads are used, the others being available as spares.

Figure 4.2: The evolution of the C-BASS reference temperatures.

4.1.2 Controlling the Resistor Temperature

After successfully controlling the temperature of the copper bobbin, the next challenge was ensuring that the reference loads tracked this temperature accurately. A number of different schemes were attempted.

We initially inserted commercial 50 ohm loads into holes machined into the bobbin. However, this was difficult as machining the holes to the correct diameter was complicated by the lack of standardization in commercial 50 ohm load outer diameters. This is clearly seen in Figure 4.2a where the loads do not fit properly into the cutouts.

Our initial intuition told us that the problem was in the thermal contact between the bobbin and the outer casing of the 50 ohm load. After discussing the situation further with Lakeshore representatives, we decided to improve the thermal contact by soldering the load in place. This was non-trivial, as heating the Cernox sensor past 400 K would have damaged it, and it couldn't be removed as it had been affixed in place using varnish. To avoid damaging the sensor, the loads were soldered into a separate piece of copper which was then firmly bolted onto the bobbin. *Apiezon* cryogenic thermal grease (again suggested by Lakeshore) was used to improve the thermal contact between the two interfaces. The bolts provided a positive pressure between the two pieces of copper not present when the loads were simply fitted into the machined holes, and providing a significantly better thermal bond. Figure 4.2b shows this configuration.

Unfortunately, this did not improve the issue. A number of tests conducted after this suggested that the resistor was still not tracking the bobbin temperature accurately, and what was worse, the resistor power seemed to exhibit occasional steps suggesting an intermittent fault in the thermal contact. The outer shell was now in very good thermal contact with the bobbin through the solder, however the 50 ohm resistor inside the load might not maintain thermal contact with the outer casing of the load during the cooldown.

A surgical procedure on the 50 ohm load revealed that the thermal contact between the resistor and outer can was made through a friction-fit. It was entirely possible that this thermal contact would be lost or would change during the thermal cycle.

To avoid this uncertainty we manufactured our own 50 ohm load by soldering two 100 ohm 0603 surface mount resistors[2] between the central conductor and the outer conductor of an SMA connector, a sugestion from Eddie Blakchurst, Manchester. The connector itself was then soldered into a copper sheet, clamped firmly onto the temperature controlled bobbin. The modification is shown in Figure 4.3. *Apiezon* grease was used to improve the thermal bond between the two copper sheet and the bobbin. This modification ensures an excellent

[2]Resistors used were part Number RN731JTTD1000F50:Thin Film Resistors - SMD 1/16W 100ohm 1% 50ppm

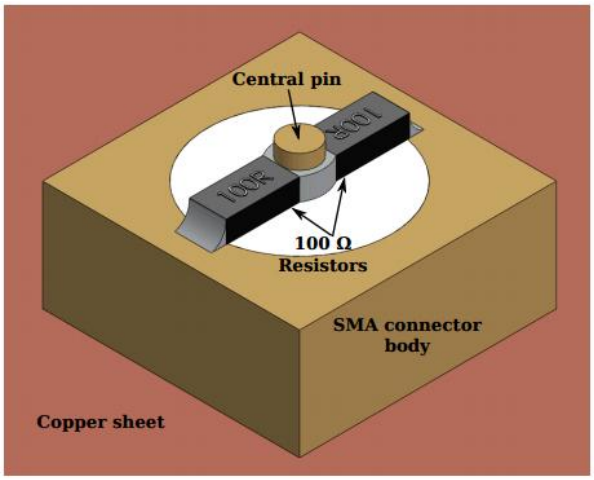

Figure 4.3: CAD drawing of the cold load. Two 100 ohm resistors were soldered between the ground and central conductor of the SMA pin.

thermal contact between the resistor and the bobbin, through the intermediate copper sheet. The reference load now provides a highly stable output power as measured by the digital backend. Stepping the temperature by 1 K produces a similar power increase to turning on the noise diode, which is reassuring as we expect $\sim$1 K noise temperature contribution from the noise diode as we shall see in the following section.

4.2 Noise Diode

The noise diode provides a means of calibrating astronomical source during the survey. In addition, by using the fact that the noise diode power should only be present in the sky signal, the noise diode provides an effective means of accurately measuring the separation between the sky and cold load signals in the digital backend, Figure 3.19. Noise is generated by an Atlantic Microwave AS5254 noise diode module mounted of a temperature stabilized plate, and is combined with the sky signal using a 30 dB directional coupler directly before the 180° hybrid combines the sky and temperature controlled load voltages.

The noise diode has an average excess noise ration (ENR) of 17.01 ± 0.01 dB at 5 GHz (see Table 4.1). This corresponds 15300 ± 50 K. Before combining the signal with the sky signal, the noise diode power is attenuated by cables, a two way power divider, the digital phase shifter (Section 4.3) and the 30 dB coupler. This is a total attenuation of $\sim$42 dB, reducing the original 15300 K signal to $\sim$1 K noise temperature. I have plotted the relationship between ENR, total attenuation and expected noise temperature in Figure 4.4.

Frequency [MHz]	ENR [dB]
10	15.99
100	15.76
1000	15.26
2000	16.05
3000	16.38
4000	16.87
5000	17.01
6000	17.94

Table 4.1: Noise diode calibration provided by Atlantic Microwave

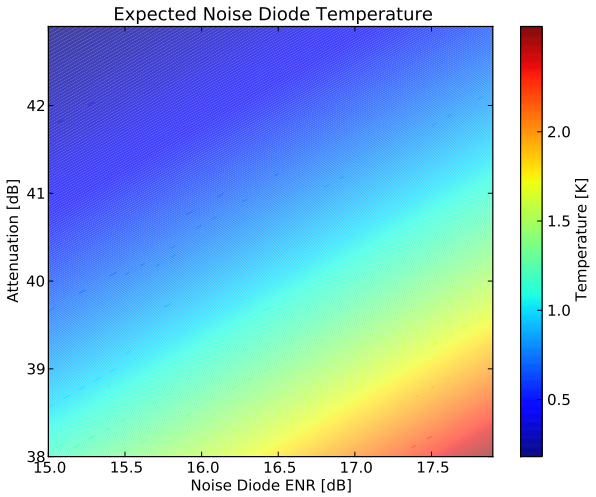

Figure 4.4: Effect of attenuation and noise diode ENR on the expected noise temperature of a noise diode event. We expect to observe $\sim$1 K from the noise diode.

4.2.1 Noise Diode Physical Temperature Stabilization and Location

The noise diode ENR is insensitive to physical temperature variations, with the manufacturers datasheet quoting the dependence at 0.01dB/°C. Despite this we felt that keeping the noise diode physical temperature stable was very important, as nois source stability is an important aspect of the instrument calibration.

We placed the noise diode inside the cryostat, and added a PID temperature controller to the mounting plate. The temperature control uses an Arduino Microcontroller to control a heater current using feedback from a DS18B20 temperature sensor. This allowed us to reduce the noise diode temperature variations to $\sim$0.1°C, as seen if Figure 4.5. Having the noise diode inside the crystat has the added benefit of preventing the interconnecting cables from being disturbed, minimizing the chance of unwanted changes in the noise diode calibration signal.

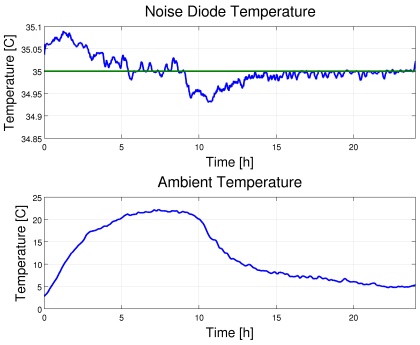

(a) Temperature of the stabilized plate and ambient temperature over a 24 hour period.

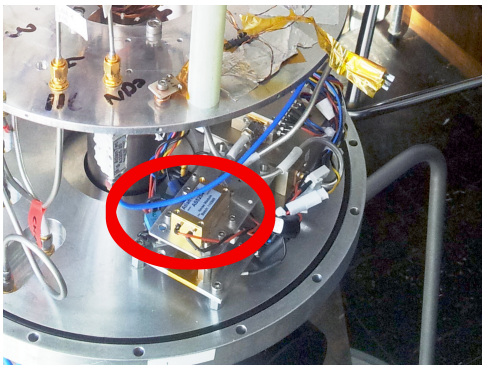

(b) The noise diode and temperature controlled plate installed at the base of the cryostat. The noise diode is circled. The plate also temperature stabilizes the phase shifters in Section 4.3.

Figure 4.5: Ambient, and stabilized noise diode temperature over 24 hours. The noise diode temperature (blue line in Fig(a)) is stabilized at 35.00±0.04°C, and shows not evidence of the the 22°C diurnal ambient temperature variation in Fig(b). This suggests that the noise diode physical temperature will be stable over long time scales.

4.3 Digital Phase Shifter

The noise diode provides an important calibration of the total-intensity measurements, but taken on its own it is not possible to use it to calibrate the polarization measurements. If both cable lengths from the power splitter are identical we expect to inject a 100% Stokes Q signal into the signal path, which does not help in calibrating Stokes U. However, if it were possible to change the phase of one of the noise diode arms, we would be able to change the noise diode from a purely Stokes Q signal to a purely Stokes U signal. This is useful for calibration purposes.

Suitable phase shifting packages are not readily commercially available, so instead we packaged the HMC649 digital phase shifter chip from Hittite Microwave. The HMC649 nominally provides 6 bit digitally controlled phase shifting between 0° and 360° with ∼6.5 dB loss in all states. I measured both the insertion loss and phase shift applied in all 64 states, and characterized the effect of temperature on both phase and insertion loss. The results are included plots in Figure 4.6. Phase appears to be stable with temperature, however the insertion loss does vary. We require long term stability on our calibration source, requiring us to stabilize the temperature of these units. Temperature stabilization is provided by the temperature controlled plate (already described in Section 4.2.1) on the 300 K stage inside the cryostat, which is maintained 35 ± 0.1°C.

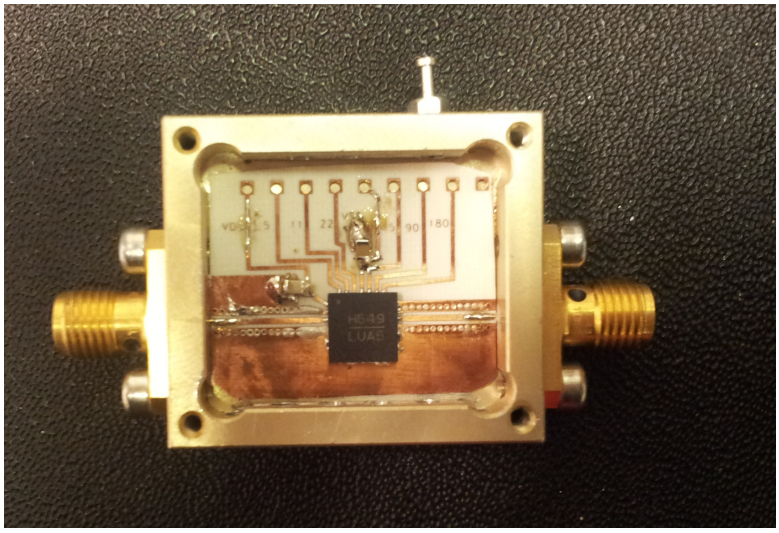

(a) HMC649 digital phase shifter. Capacitively filtered feed throughs (not shown in this photo) provide the power and control lines

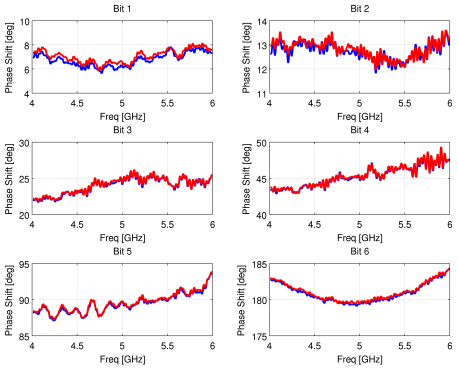

(b) Measured Phase Shifts with temperature. Blue is 33°C and Red is 45°C.

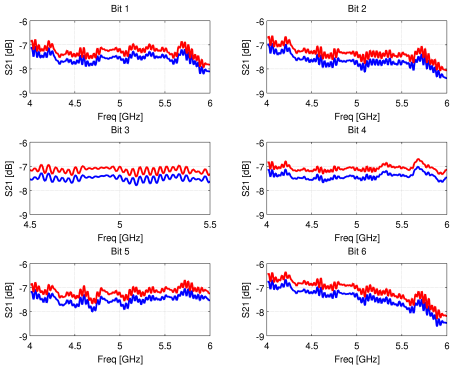

(c) Measured S21 Change with temperature. Blue is 33°C and Red is 45°C.

Figure 4.6: Temperature and phase effect on HMC649 digital phase shifter. Temperature has no appreciable effect on phase, but does e affect the insertion loss. To mitigate this effect we control the phase shifter temperature on a temperature-controlled plate within the cryostat. This plate also controls the physical temperature of the noise diode, and is stable to within 0.5°C.

4.4 Phase Switches

Wideband 180° phase switches were manufactured for the Northern survey (King, 2009).

These phase switches were essential for the analog instrument, allowing us to switch out gain drifts in the detector diodes and post-diode electronics. In the analog receiver, these gain drifts are not removed by the continuous-comparison architecture, as these stages occur after the sky and load signals are separated by the second 180° hybrid.

However for the Southern receiver, we separate the sky and load signals digitally, after sampling the signals using the iADC. The continuous-comparison architecture removes all the low-frequency drifts that may impact the signal before this point, suggesting that we did not need phase switching at all. However, we felt that the phase switching would offer some immunity to other low-frequency signals, and incorporated them into the signal chain during the design phase. Our final testing has confirmed that the phase switching is not necessary for the digital receiver, and we have not implemented it.

4.5 Orthomode Transducer

The orthomode transducer (OMT) converts electromagnetic radiation into a voltage that we can sample. The C-BASS OMT (Grimes et al., 2007) was designed to be compact and coolable inside the cryostat to reduce the effect of any resistive losses on the system temperature. The OMT separates the polarizations using four planar probes. Linear polarizations are produced by combining the opposite out-of-phase signals using a 180° hybrid.

4.6 Linear-to-Circular Convertor and the 180° Hybrids

Measuring linear polarizations is best done by correlating the two circular polarizations, since correlations are immune to the uncorrelated amplifier gain drifts that impact on the intensity measurements. This is described in detail in Section 1.7

The compact OMT introduced in Section 4.5 requires the out-of-phase signals to be combined to produce the linear polarizations. This is done with a 180° hybrid for each polarization. We convert these to circular polarization in order to take advantage of the superior stability of a correlation detection. This is done using a 90° hybrid which combines the two linear polarizations with a 90° phase difference, thereby producing circular polarization. The three hybrids were combined into a single unit located directly after the OMT.

Standalone 180° hybrids were also used to combine the reference load with the sky signal as part of the continuous-comparison archtiecture.

These components were both designed for the C-BASS North survey (King, 2009).

4.7 Power Dividers

The C-BASS architecture used two different power dividers: one microstrip design for the RF band (4.5–5.5 GHz), and another lumped element based splitter for the IF band (DC–1000 MHz).

4.7.1 3 GHz–8 GHz Power dividers

These were designed for the C-BASS North instrument(King, 2009) and were re-used in the Southern system. They use a Wilkinson power divider design, which provides high port-to-port isolation. The power dividers are used to split the noise diode reference signal, and a four-way power divider is used to split the 4.5 GHz LO used to drive the RF band to baseband mixers.

4.7.2 DC–2 GHz Power Divider

These power dividers were developed by myself using the ADP-2-10 package available from Minicircuits and are suitable for use from 5–1000 MHz with approximately 20 dB isolation between the output ports. When we originally used these splitters, we had the DC–500 MHz sub-band on one output and the 500–1000 MHz sub-band on the other. The reflections from the baseband defining filters (see Section 4.9) were coupling back through the input port, and generating interference across the splitter outputs. We originally mitigated this by providing additional attenuation after the splitter, however these were replaced with the IF amplifiers detailed in Section 4.11, which were placed between the power divider and filters, effectively isolating the two components.

4.8 4.5–5.5 GHz Band-pass Filters

The 4.5–5.5 GHz filters were designed for the C-BASS North survey and reused in this system.

These filters define the 1 GHz RF bandwidth of the C-BASS survey. When first deployed we believed they provided sufficient out-of-band rejection, however Figure 2.8 clearly shows that the 4.2 GHz C-Band satellite downlinks were still visible in the Northern survey with a single band-pass filtering stage. In order to avoid this in the Southern survey, we cascaded the two band-pass filters, effectively doubling the out-of-band rejection. The transmission of both a single filter and two filters cascaded in series is given in Figure 4.8.

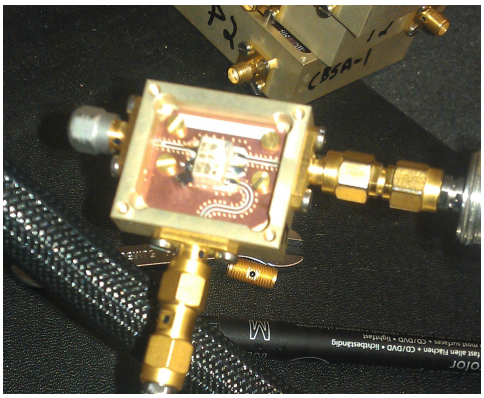

(a) Compact design

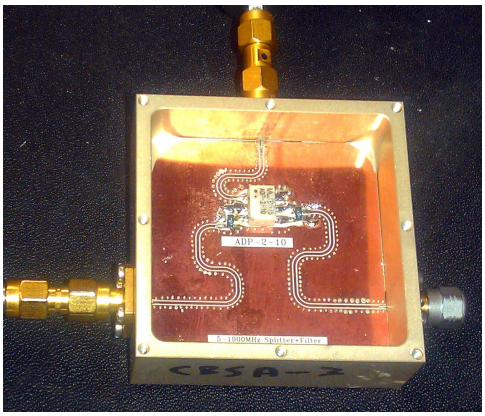

(b) Design including SMT attenuators

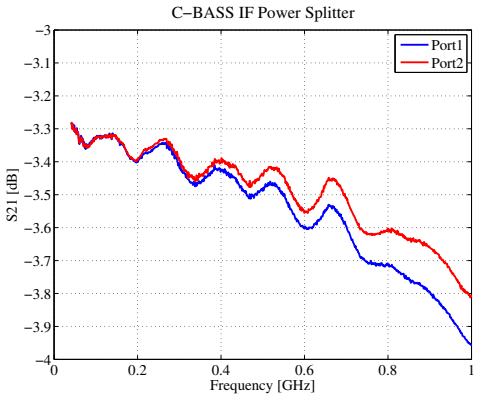

(c) S21 (compact design)

Figure 4.7: The 5 MHz to 1000 MHz power splitter using a Minicircuits ADP-2-10 component. These provide 20 dB port-to-port isolation between the two sub-bands sampled by the iADC. Two designs were built; one compact design to fit in a standard box, and a larger design with the option to add additional surface mount attenuation on the two outputs to increase the port-to-port isolation.

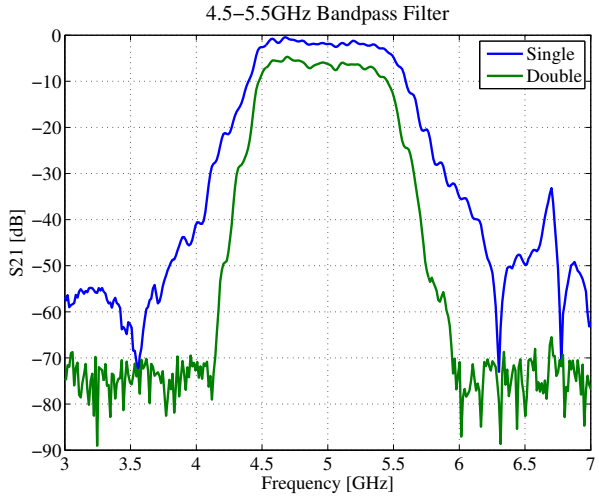

Figure 4.8: RF 4.5–5.5 GHz band-pass filter response. A single band-pass filter has a transmission given by the blue line, while cascading two band-pass filters reduces the out-of-band interference, but increases the in-band attenuation.

4.9 Low Frequency Filters

The C-BASS RF band of 4.5–5.5 GHz is defined using edge-coupled filters early in the RF chain (see Section 4.8). After suitable RF amplification, the signal is downconverted to a baseband of DC–1 GHz using a 4.5 GHz LO. We rely on the 4.5–5.5 GHz band-pass filters to define the DC and 1GHz baseband edges, however we require suitable 500 MHz low pass filter and 500 MHz high pass filter to divide the signal into the two sub-bands for digitization.

The requirements were small form factor, 20dB roll off in the first 20MHz of band overlap and manufacturability.

4.9.1 Filter Design

The difficulty of designing filters in the 1 GHz frequency range is that it straddles two design regions (Matthaei, G and Young, L and Jones, 1980). Higher frequency filters, such as the band-pass filter in Section 4.8, are generally designed using microstrip. The microstrip features scale with wavelength making them unsuitably large for use at low-frequencies. Low frequency designs have tended to use lumped-element capacitors and inductors, but these are notoriously difficult to use at higher frequencies where parasitic inductive, capacitive and resistive effects become more pronounced.

The design we have employed uses a lumped-element based approach, but compensates for the parasitic effects by using published component values from Murata (Murata, 2009). Simulations of the filter response are carried out using the *Microwave Office* software suite.

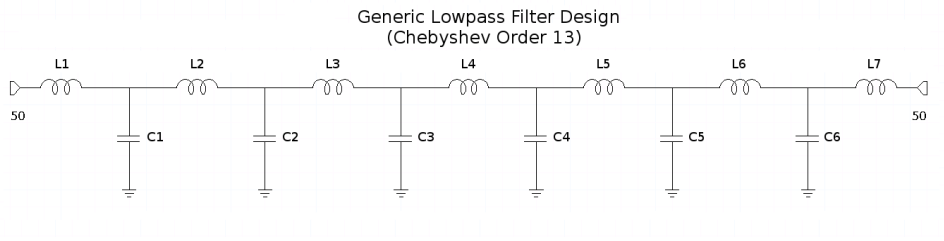

(a) Lowpass filter design schematic.

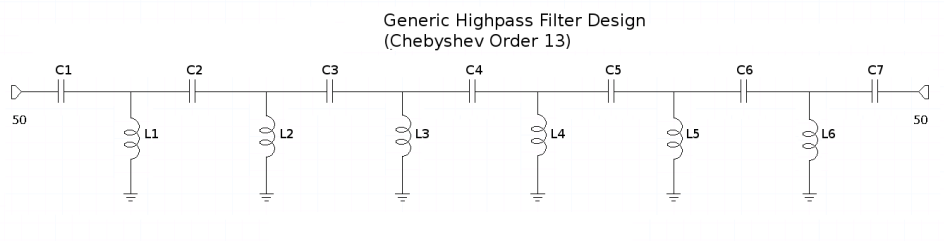

(b) Highpass filter design schematic.

Figure 4.9: Low frequency lumped element filter designs and manufacture. Component values for specific filter frequencies are given in Table 4.2 to Table 4.4.

Circuits based on the schematics in Figure 4.9 are initially simulated to meet specification, and then manufactured using painstaking soldering and boxing techniques developed by Mr Adam Baird in the Department of Physics Central Electronics Group. This process has allowed us to build lumped-element filters for use at unusually high frequencies with rapid turnaround times. The techniques also lend themselves to mass production using pick and place machines on PCB's.

4.9.2 Component Choices

As has been mentioned, parasitic capacitance, inductance and resistance increases at higher frequencies, and can present a significant design problem. These effects need to be carefully considered during design.

Many manufacturers release component libraries, which have typical expected parasitic capacitance/inductive values as a function of frequency. We chose to use Murata components, because they both release these component libraries and have ready availability of components in small quantities.

Surface mount inductors can present especially significant challenges at high frequencies. Parasitic capacitance exists due to electric fields across small potential differences between the wire coils, and the wire itself presents a parasitic resistance. We tried the LQW18/LQW15 wire wound type inductors and the LQG18/LQG15 multilayer inductors (Murata, 2011) before settling on the wire-wound type which was more robust during the high temperatures required

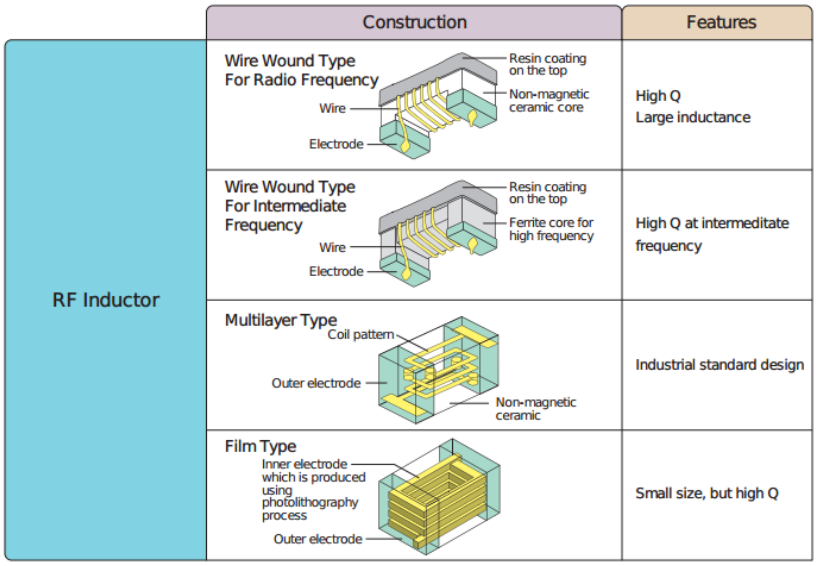

Figure 4.10: Different manufacturing styles of chip inductors (Murata, 2011).

for manufacture. The style of both is shown in Figure 4.10. Both provide ranges from 1-270 nH and are designed for use at high frequencies.

For capacitors we chose the GJM15 Series ceramic capacitors also from Murata (Murata, 2010). These are specifically designed for high frequency operations of up to 10 GHz, and are readily available across the full range of values from 0.1-33pF. They feature low parasitic series resistance (ESR). More details can be found in Murata (2010).

Manufacture Techniques and Packaging

While component choice is an important consideration, so is symmetry in the design. We have used coplanar waveguide to allow access to the ground plane on the top side of the board, and have laid out the components in a symmetrical fashion. As can be seen in Figure 4.11, careful consideration also needs to be given to the way in which the filter boards are packaged. The packaging plays an important role in improving the out-of-band rejection characteristics.

We focused on using a design and manufacture technique that would allow rapid prototyping of these filters, and would allow quick turnaround times for future filter manufacture. The techniques we have developed allow a custom filter to be built in a day. This allows flexibility in designing low-frequency filters and the technique opens up the possibility of building compact filters directly onto printed circuit boards using pick and place for mass production. This has potential for use in large scale radio astronomy deployments such as the Square Kilometer Array.

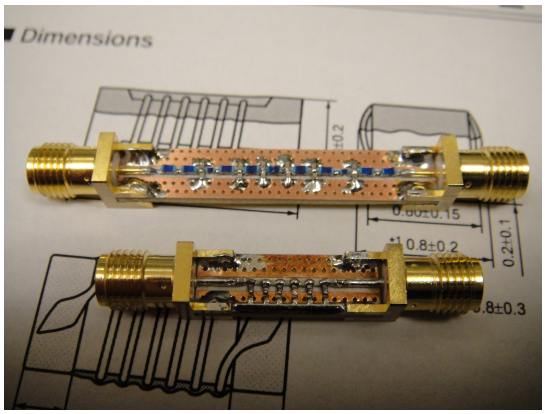

(a) Examples of uncased filter boards.

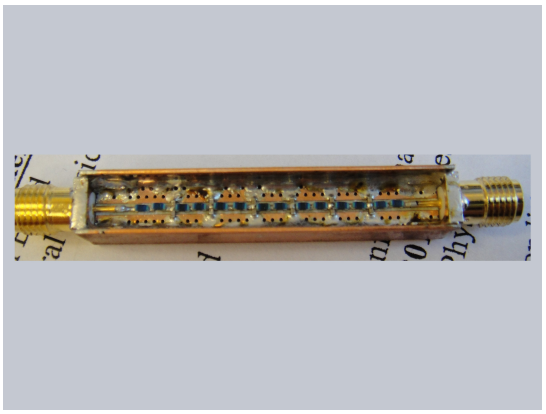

(b) Filter in box, without lid

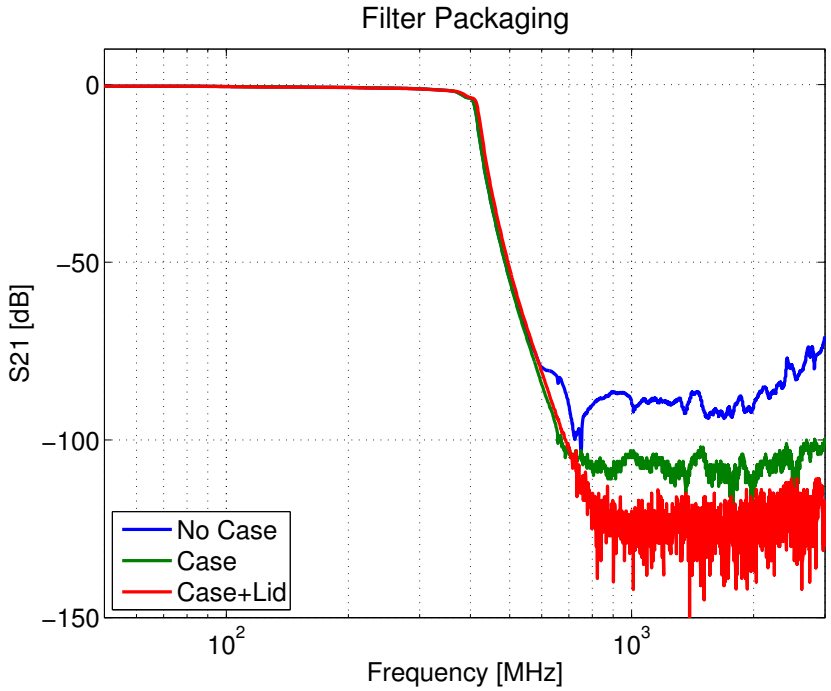

(c) The effect of boxing on the response of a 400 MHz low-pass

Figure 4.11: The different stages of the filter construction process. The plots of filter response show the importance of properly enclosing the filter boards

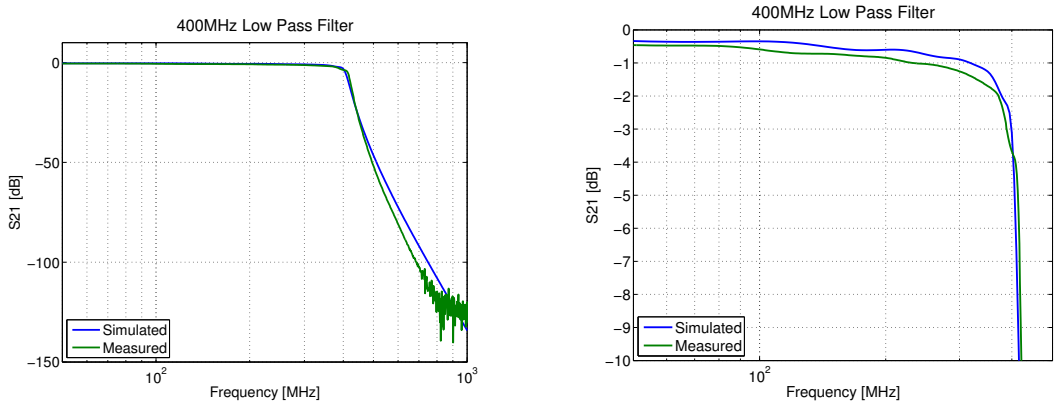

(a) (b)

Figure 4.12: Simulated 400 MHz filter performance against measured performance.

4.9.3 Performance

The filters perform well offering up to 120 dB of out-of-band rejection as can be seen in Figure 4.12 which was made on a VNA with suitable dynamic range to measure this. The plot also shows the good agreement between simulation and measured performance.

The measured filter roll-of is 250dB/octave. For a 500 MHz filter this is a roll-off of 0.5 dB/MHz. Given the 8 MHz spectral resolution, this means that we can expect >4dB attenuation of out-of-band signals per channel. We use this filter to provide anti-aliasing between the two C-BASS sub-bands. This means we we will only need to discard five channels to guarantee 20 dB isolation between the two sub-bands.

4.9.4 Component Values

In this section I list the actual component values for the C-BASS filters implemented using the basic designs in Figure 4.9. These are: a 500 MHz low pass, a 1000 MHz low pass and a 500 MHz high pass.The first simulation used models of wirewound inductors and multilayer ceramic capacitors, MLCCs, to simulate the parasitic inductance and capacitance of each component and include these in the response. This provides a rough guide as to the required component values.

We then carried out a second simulation using available components and using actual measured values of both inductors and capacitors provided by Murata. We chose the components to be as close to the first simulation as possible. In some cases this required using two components so that $L_i \approx L_i a + L_i b$ and $C_i \approx C_i a + C_i b$. In these instances, inductors $L_i a$ and $L_i b$ were placed in series while capacitors $C_i a$ and $C_i b$ were placed in parallel. After achieving satisfactory results in simulation the filters were manufactured as shown in Figure 4.11.

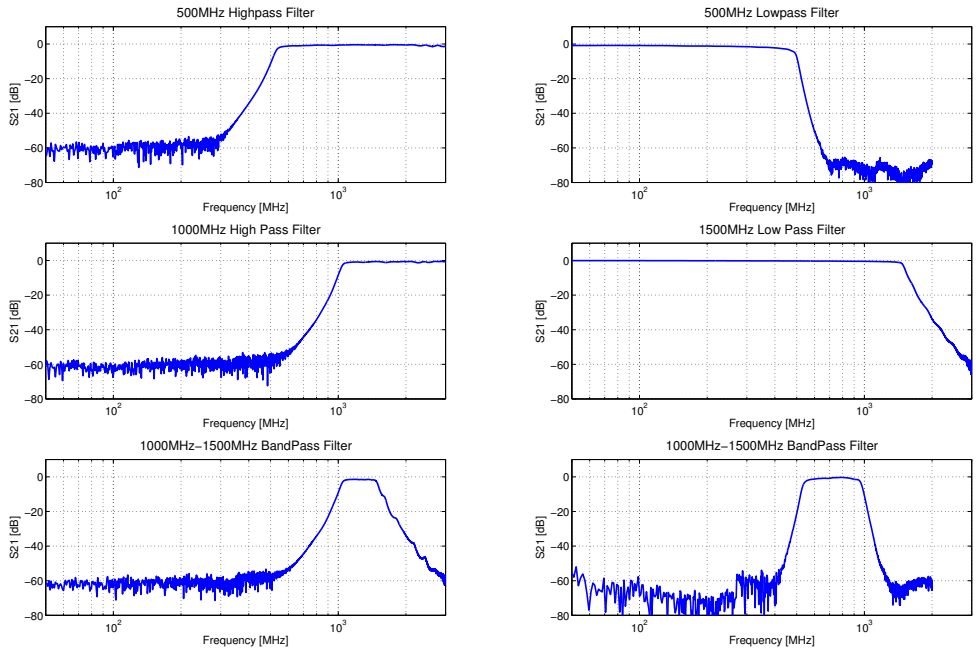

Figure 4.13: Filters produced using the low-frequency filter design technique. Note the log scale on the frequency axis. Using the lowpass and highpass filters in series it is possible to manufacture band-pass filters too.

L_i	L_ia	L_ib	C_i	C_ia	C_ib
12	6.2	4.7	5.1	2.2	3
10	4.7	4.7	3	3.0	-
10	4.7	4.7	3	2.7	-
10	4.7	4.7	3	2.7	-
10	4.7	4.7	3	2.7	-
12	6.2	4.7	3	3	-
-	-	-	5.1	2.2	3

Table 4.2: 500 MHz High pass filter. Values in nH and pF for the inductors, L_i, and capacitors, C_i, respectively.

L_i	L_ia	L_ib	C_i	C_ia	C_ib
10	4.7	4.7	2.0	2.0	2.7
18	9.1	8.2	5.1	2.2	3.0
18	18	-	5.1	2.7	2.7
18	18	-	5.1	2.7	2.7
18	18	-	5.1	2.2	3.0
18	9.1	8.2	4.7	2.0	2.7
10	4.7	4.7	-	-	-

Table 4.3: 1000 MHz Low pass filter. Values in nH and pF for the inductors, L_i, and capacitors, C_i, respectively.

In the Table 4.2-Table 4.4 L_i and C_i are the initial inductor and capacitor component values from the simulation, in nH and pF respectively, laid out as in Figure 4.9. The component values used in the manufactured filters are given by L_ia, L_ib, C_ia and C_ib.

L_i	L_ia	L_ib	C_i	C_ia	C_ib
19	18	-	9.2	4.7	3.0
35	12	22	10.5	8.2	2.2
36	18	18	10.6	6.8	3.9
36	18	18	10.6	6.8	3.9
36	18	18	10.5	8.2	2.2
35	12	22	9.3	4.7	4.7
19	18	-		-	-

Table 4.4: 500 MHz low pass filter. Values in nH and pF for the inductors, L_i, and capacitors, C_i, respectively.

4.10 2–20 GHz Amplifiers

The RF amplifiers we used were designed for C-BASS North (King, 2009). The design uses a pair of cascaded Hittite HMC462 wideband amplifiers and provides 27 dB of gain across the 2–20 GHz band, with a measured noise temperature of 200 K in the C-BASS band. Given the high gain of the Low Noise Factory LNA's, the Friis Equation in Equation 1.50 tells us that the impact on $T_{\rm sys}$ of these subsequent gain stages is negligible.

An unfortunate side-effect of the wideband nature of these amplifiers is amplification of unwanted, out-of-band signals and noise. This has already been discussed with reference to the Northern instrument in Section 2.5, where we required additional band-pass filtering to remove contamination from C-band downlink satellites. For the Southern instrument we added this additional stage of filtering during the design phase.

4.11 DC–3 GHz Amplifiers

The digital backend required gain at the DC–1000 MHz baseband frequencies. We used two different, pin compatible, amplifiers, the Berex 14a and the Minicircuits GVA-84. As can be seen in Figure 4.14c, the two had very different gain characteristics. The GVA-84 has 7 dB higher gain across the band of operation, while the Berex 14a has a wider operating band and flatter gain curve. For C-BASS we used the Berex 14a amplifier which has a flat gain curve across the C-BASS baseband frequency range.

4.11.1 Bandwidth-limiting components

In addition to the amplifier, examinination of the amplifier circuitry in Figure 4.14b shows two other potentially band-limiting components. These are the DC blocking capacitors and the biasing choke. I carefully selected these components to increase the bandwidth of the final amplifier product.

Wideband Capacitors

DC blocking capacitors need to behave like series shorts in the frequency of interest. Capacitors can be modelled as a capacitor, C, in series with a small, parasitic inductance, L and equivalent series resistance, ESR. (Cain, 2010, Fiore, 2001, Urs & Rangu, 2011). The impedance, X_{cap}, is then given by the equation in Equation 4.1. This needs to be low (typically less than 1 ohm) across the frequency range of interest.

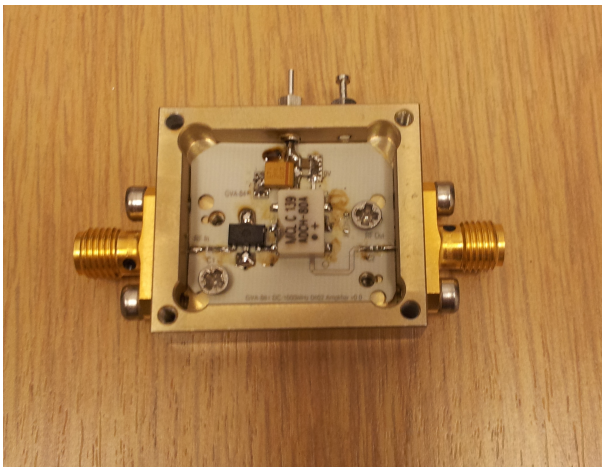

(a) Packaged DC–3 GHz amplifier

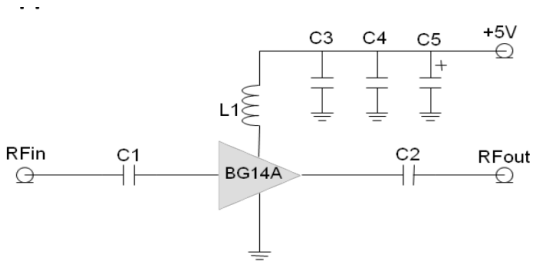

*C1, C2, C3 =100 pF ± 5%; C4 = 1000 pF ± 5%; C5 = 10uF; L1 = 39nH
*less than 20nH improves RF performance at over 1.9GHz.
*40nH or higher value L1 improves RF performance at under 500MHz.
*Optimum value of L1 may vary with board design.
*C1,C2=8200pF, L1=1200nH for 70MHz application,
Vcc=5.2V if 1200nH is used to compensate IR drop across L1.

(b) Recommended amplifier circuitry (Berex, 2010)

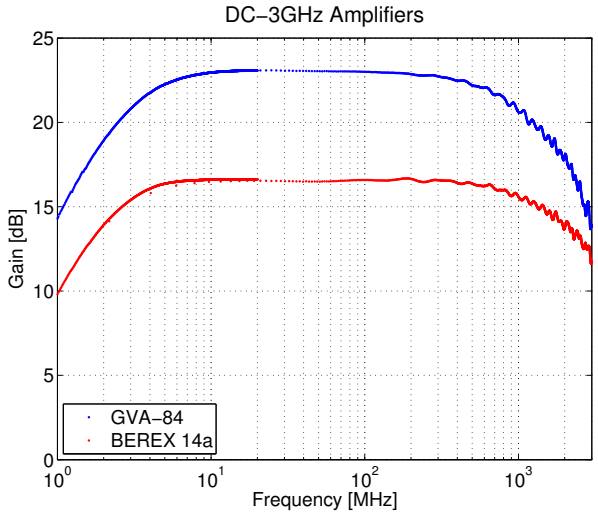

(c) Performance

Figure 4.14: The DC–3 GHz amplifier design and performance

$$X_{cap} = 1/(2\pi fC) + j2\pi fL + ESR \tag{4.1}$$

Using Equation 4.1 large capacitor values will lower X_{cap}. However larger capacitors typically exhibit higher levels of parasitic inductance whose effects increase with frequency. This can limit the effective frequency of large-value capacitors.

A typical 0603 packaged multilayer ceramic capacitor, MLCC, has a typical parasitic inductance of 870 pH (Cain, 2010), largely caused by the geometry of the package. Using Equation 4.1 we can easily see the affect of these parasitics. At 5 MHz a 100 nF capacitor has an impedance due to capacitance of 0.3 ohms, and a negligible 0.03 ohms contributed by typical parasitic inductances, making it safe to ignore these effects. However at 3 GHz the impedance from the parasitic inductance is now a non-negligible 16 ohms, and needs to be taken into consideration.

In order to achieve wideband performance from DC to 3 GHz, wideband DC blocking 100 nF, 0402 capacitors from American Technical Ceramics were used (American Technical Ceramics, 2012). The datasheet claims a self resonant frequency, SRF, of up to 40 GHz implying that the parasitic inductance is negligible up to these frequencies. The capacitors also work very well at low-frequencies as can be seen in Figure 4.14, with the gain only rolling off at frequencies below 5 MHz for the Berex 14a.

Wideband Biasing Choke

The biasing circuitry for the amplifier requires a large inductor in series with the supply voltage. It is well known that large inductors exhibit low-frequency SRF as the parasitic capacitance can create a LC notch filter. We used a wideband amplifier biasing choke (ADCH-80A) from Minicircuits to avoid this.

4.12 Double Sideband Mixers

The receiver design requires that we mix down from our RF band, 4.5–5.5 GHz, to our baseband, DC–1000 MHz. The easiest implementation was to use double sideband mixers with a single LO set at the low end of the RF band, 4.5 GHz, to produce a lower sideband (LSB) of DC–1 GHz and and upper sideband (USB) of 9–10 GHz.

While it is possible to purchase suitable double sideband mixers from commercial suppliers, they are also fairly straightforward to make in house. This section describes the process

4.12.1 Design and Component Choice

The mixer uses the HMC128G8 integrated circuit from Hittite Microwave. This is a double balanced mixer, suitable for use at RF frequencies between 3–8 GHz, providing scaleability for possible future receivers, and comes in a surface-mount package that is very easy to assemble. I designed a PCB using coplanar waveguide that allowed us to mount the mixers in the same-style box as used for the amplifiers.

4.12.2 Performance

I measured the mixer performance using a 14 dBm LO at 4.5 GHz to simulate the C-BASS system. I swept a narrow bandwidth signal across the band, simultaneously recording the detected power on a power meter. As a reference we used an evaluation board supplied by Hittite.

The results are pleasing and are included in Figure 4.15. The C-BASS mixers are very similar in response to the Hittite evaluation board.

4.13 Low Noise Amplifier Power Supply

Low Noise Factory supply bias supply units for their LNAs. However these do not allow remote monitoring and can only bias three amplifiers per board. The schematic is shown in Figure 4.16. As we required remote monitoring of the bias values and C-BASS needed biasing for four amplifiers, we built our own bespoke LNA bias supplies.

While the basic supply was similar to the bias supplies from Low Noise Factory, the design featured a few significant improvements, primarily in the choice of of components. The C-BASS stability criteria requires minimizing long term bias drifts, and we were particularly concerned that ambient temperature fluctuations would affect the amplifier biasing. To mitigate against this we:

1. Changed from a zener diode based voltage reference (thermal stability 50ppm/°C (Texas Instruments, 2011)), to a REF195 E-grade voltage reference (thermal stability 5ppm/°C (Analog Devices, 2011)).

2. Changed from the TL074D opamp (Vo≈3000uV $\Delta V_{OS}/\Delta C$≈18uV/°C (Texas Instruments, 2005)) to the AD8624 opamp (Vo≈10uV $\Delta V_{OS}/\Delta C$≈1.2uV/°C (Analog-Devices, 2011))

A schematic of the bias supply generator is shown in Figure 4.17. The $\pm 10V$ reference voltage generator is shown in Figure 4.18.

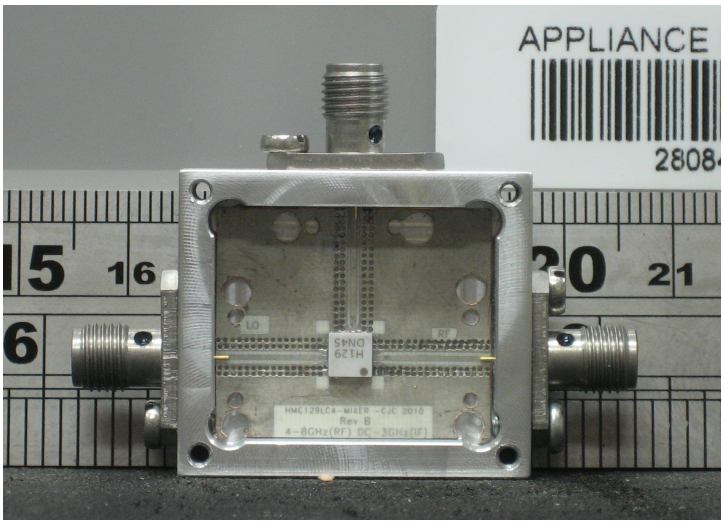

(a) The packaged Hittite Microwave HMC128 mixer

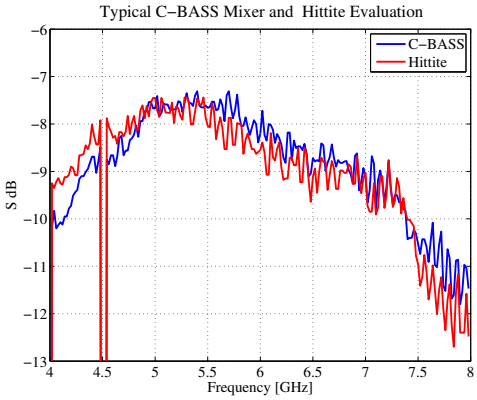

(b) The S21 from 4–8 GHz

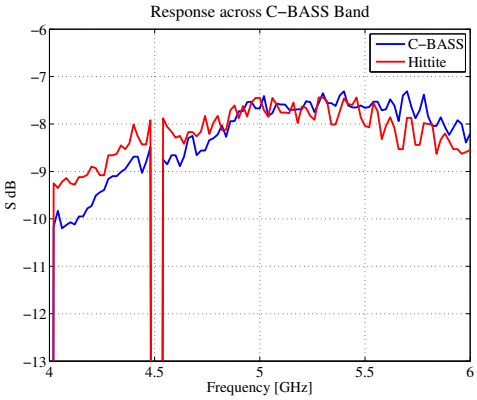

(c) The S21 across the C-BASS Band

Figure 4.15: The HMC128 mixer and the results from the packaged unit. I used a 4.5 GHz mixer frequency with a power level of 14 dBm for these measurements.

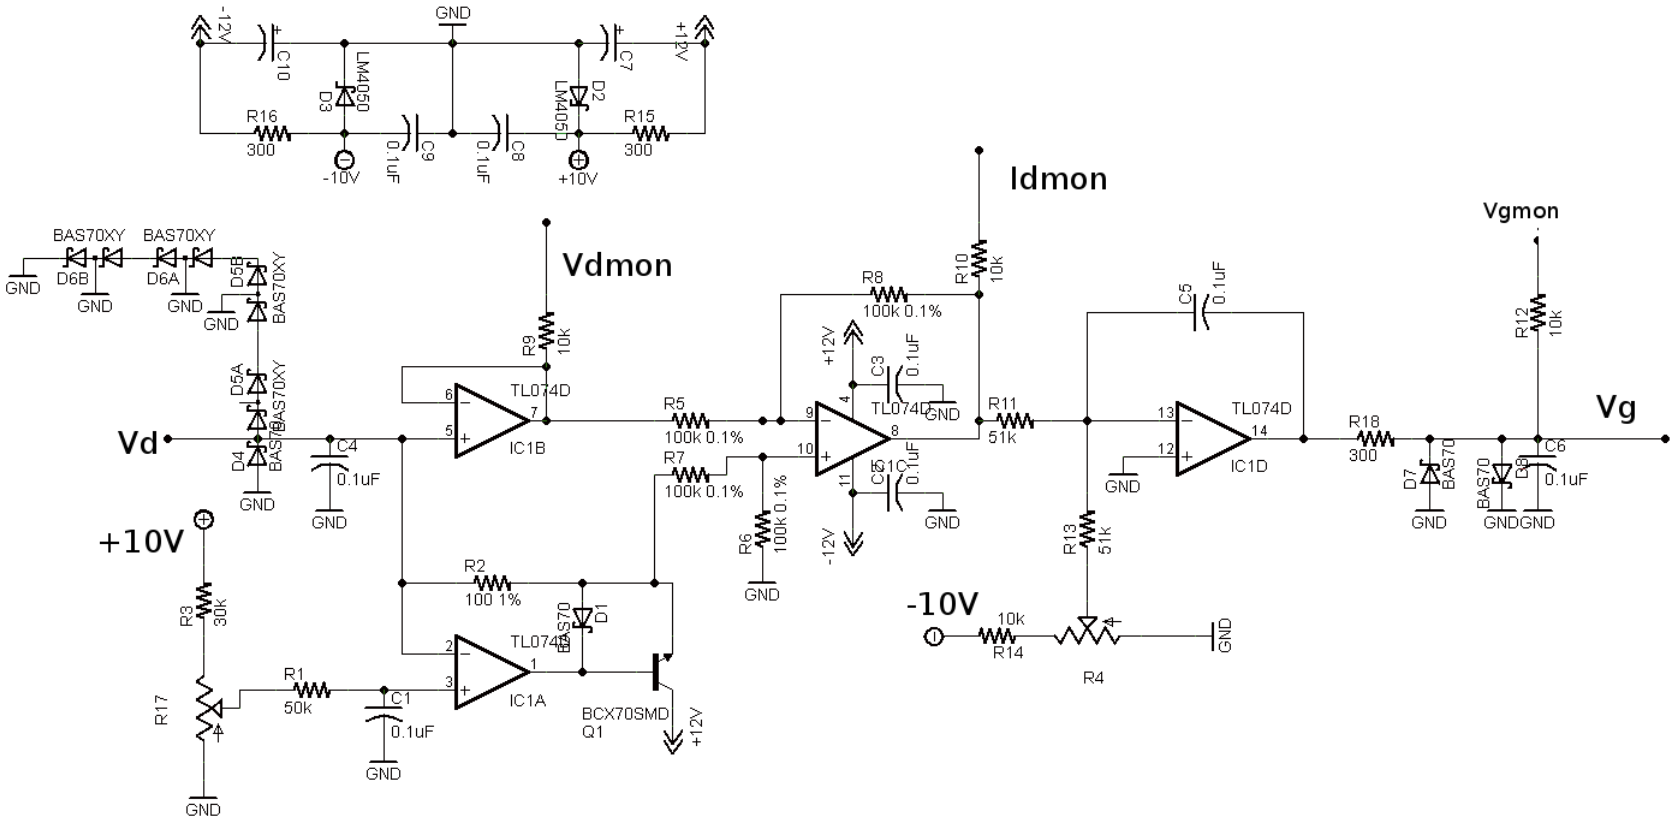

Figure 4.16: The original amplifier power supply schematic. ±10V reference was supplied by LM4050 Zener diode references (thermal stability 50ppm/°C (Texas Instruments, 2011)), and TL074D opamps (Vo≈3000uV $\Delta V_{OS}/\Delta C \approx 18uV/°C$ (Texas Instruments, 2005)) were used for the gain and current stages.

These were replaced by the more thermally stable REF195 voltage reference (Analog Devices, 2011) and AD8624 opamps (Analog-Devices, 2011). Figure 4.18 gives the schematic of the implementation.

4.13.1 Circuit Stability

The completed design showed evidence of high frequency oscillations in the drain voltage. These were not present in the LNF design, and were generated by the AD8624 IC's. The LNAs have 0.1μF input capacitance, and the opamps were not capable of driving this capacitive load.

In order to diagnose and remedy the problem, I simulated the entire circuit in LTSpice. A screenshot of the model is shown in Figure 4.19.

The simulation successfully modelled the oscillations we observed on the drain line, and allowed proper evaluation of different methods of remedial action. Afer testing a variety of fixes using these simulations, we determined that manipulating the amplifier noise gain[3] increased the phase margin sufficiently to stabilize the circuit. The noise gain can be changed by placing a small resisor - or resistor and capacitor in series - across the inverting and non-inverting inputs of the AD8624 opamp. We found that a single 50 Ω resistor placed across the opamp inputs shown in Figure 4.19 worked admirably.

4.13.2 Long term bias stability

The LNA bias supply was redesigned in order to mitigate against any long term variation in the LNA drain current, I_D. I have plotted a 12 hour stretch of measured I_D data in Figure 4.20 together with the ambient temperature, which varies between 4-18°C over this period. There is no easily discernible temperature effect on I_D. I have also plotted the spectrum of the same stretch of data, which shows that the long term variations in the LNA only begin to exhibit an $1/f$ noise spectrum at periods longer than 1×10^{-3} Hz. This is longer than the stabilized C-BASS total intensity fluctuations shown in Figure 1.9 and should not affect the total intensity stabilization.

4.14 Conclusion

This chapter has built on the previous chapter, and provided significant detail on a number of the components used to build the C-BASS receiver. Some of these components were re-

[3] http://www.analog.com/library/analogdialogue/archives/31-2/appleng.html http://www.analog.com/library/analogdialogue/archives/43-09/edch1opamps.pdf

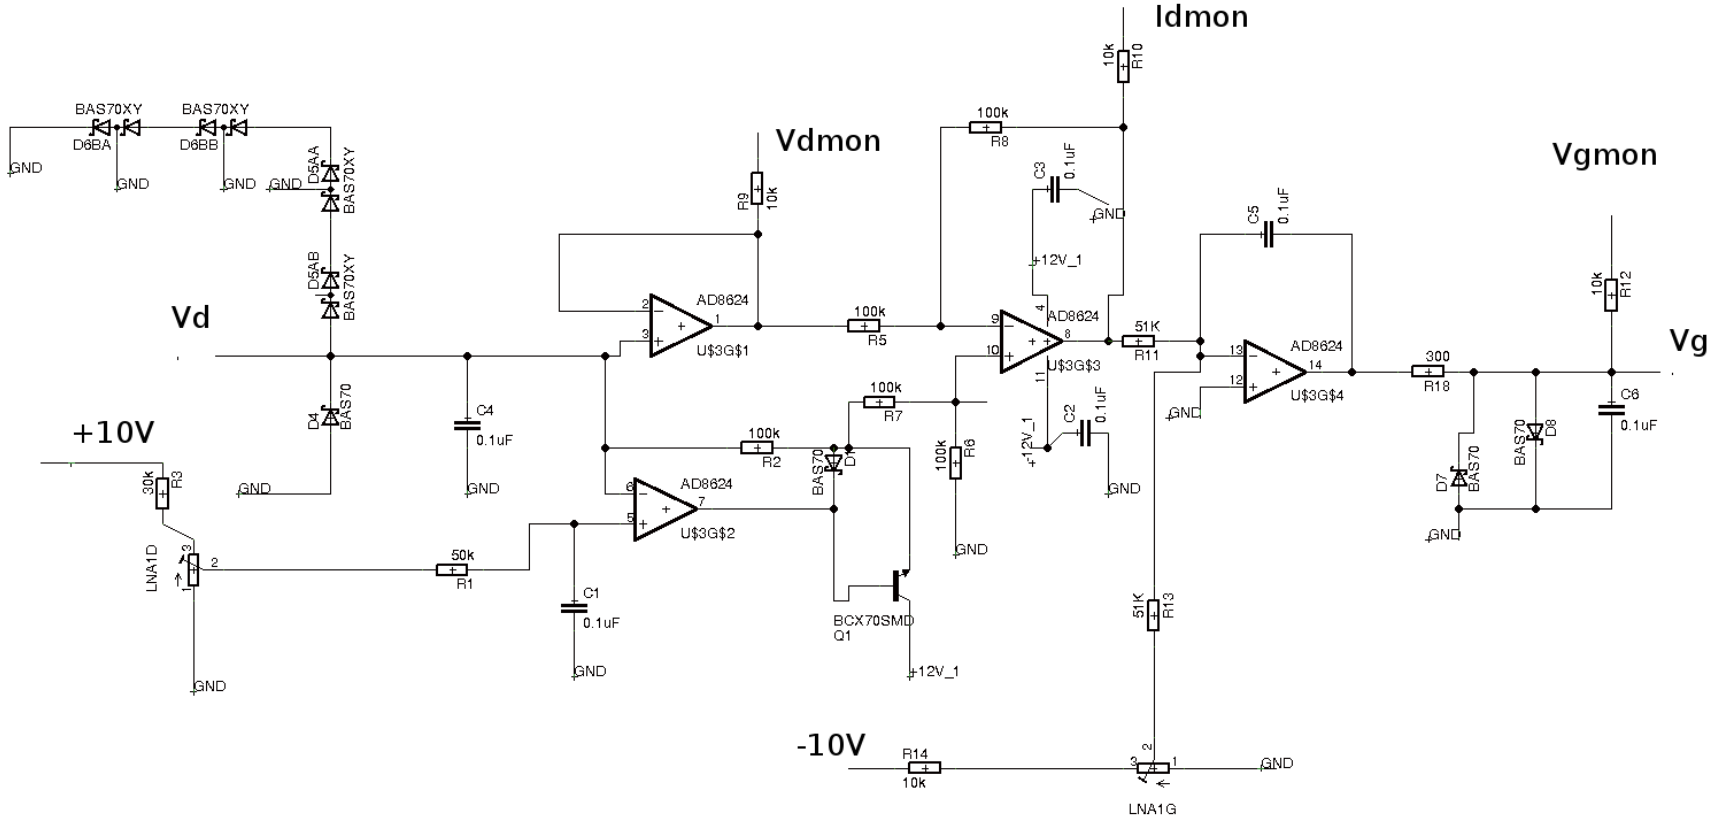

Figure 4.17: New Power supply schematic. The AD8624 opamps (Vo≈10uV $\Delta V_{OS}/\Delta C \approx 1.2$uV/°C) (Analog-Devices, 2011) are less susceptible to temperature induced gain variations. I also used a different, more stable reference voltage generator shown in Figure 4.18 based around the REF195 E-Grade voltage reference (thermal stability 5ppm/°C (Analog Devices, 2011)).

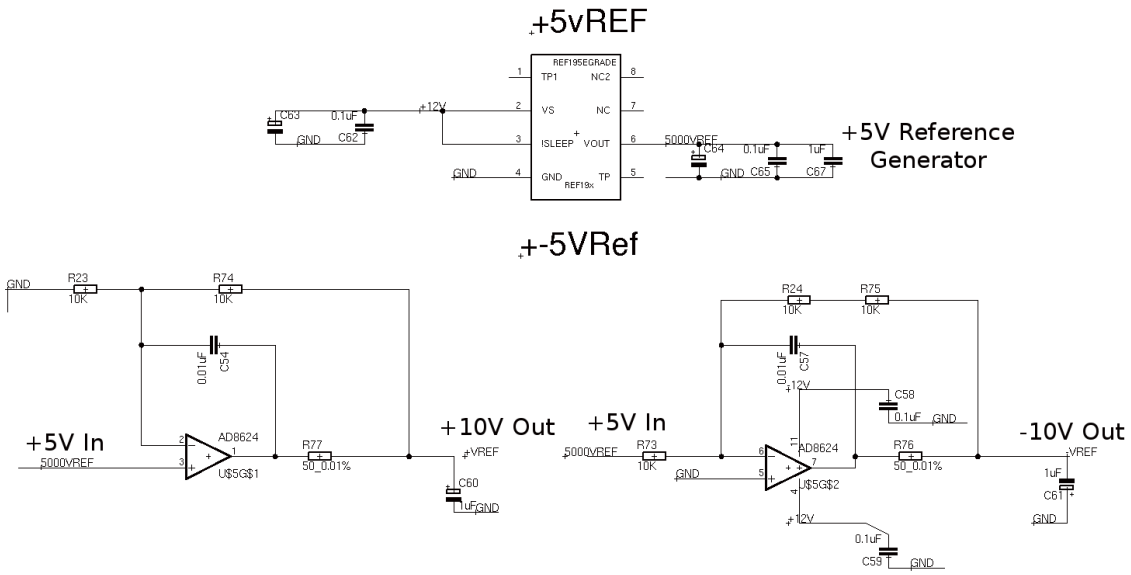

Figure 4.18: The C-BASS LNA bias supply-The +5V reference voltage is provided by a REF195 E-Grade reference IC with a temperature coefficient of less than 5ppm/°C (Analog Devices, 2011).

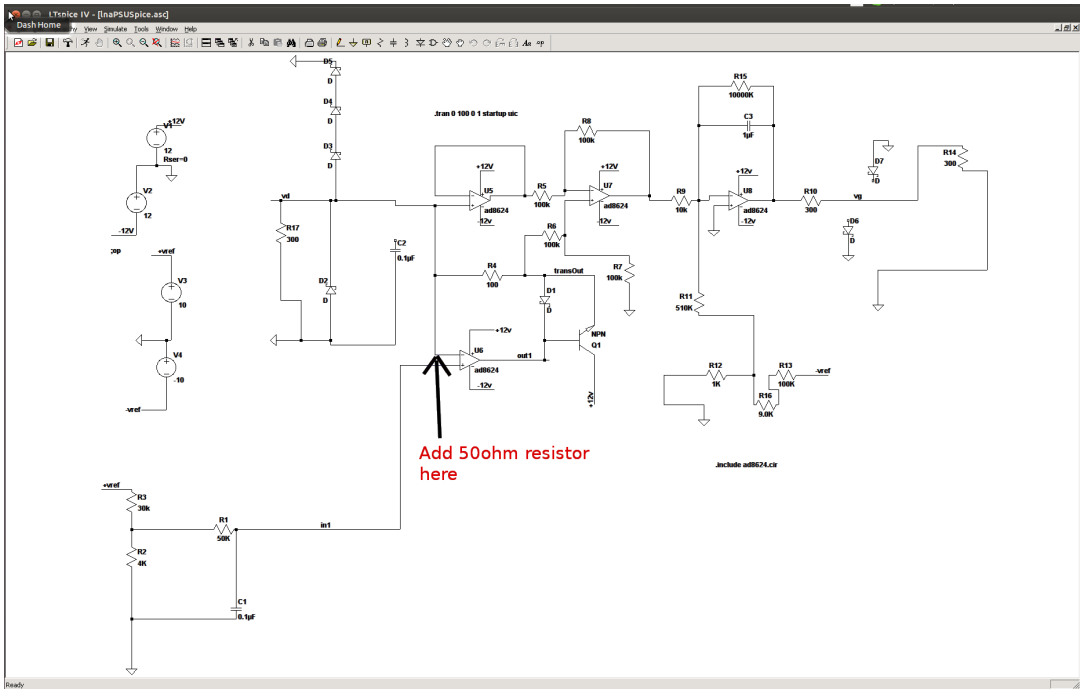

Figure 4.19: Screenshot of the spice model of the LNA bias supply built in LTSpice. A 50 Ω resistor added between the inputs of the opamp shown in the image prevented the Drain voltage from oscillating.

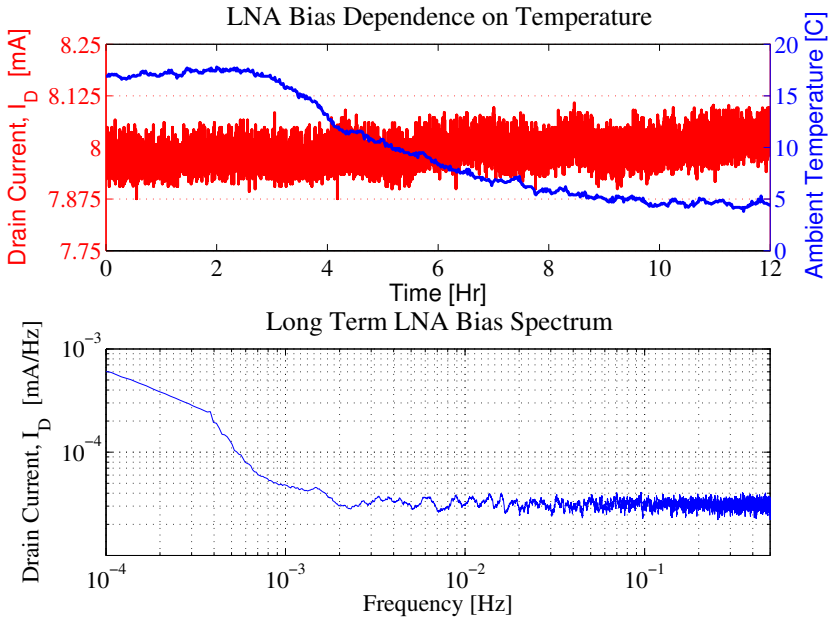

Figure 4.20: The LNA drain current plotted against the ambient temperature. The plot shows the thermal stability of the newly designed LNA bias supply. The bias supply $1/f$ noise begins to dominate at periods longer than 1×10^{-3} Hz. As this is a longer than the stabilized total intensity $1/f$ knee frequency, this will not affect the C-BASS sensitivity.

used from the Northern instrument design, however a number were developed specifically to improve the Southern system. We have focused on characterizing both the stability of the individual components, as well as their immediate suitability to the digital C-BASS instrument.

The next chapter will consolidate the information presented in both this chapter and the previous chapter. I will present lab tests carried out to test the cryogenic receiver, together with the rest of the digital backend. This is in preparation for the results in Chapter 6 where I describe the results of the final receiver deployment.

Chapter 5

Receiver Testing

> In just six weeks from the time the design was started, we had the motor on the block testing its power.
>
> Orville Wright

If only this were the case for C-BASS! After nearly three years of development we were finally able to assemble the complete digital receiver. Before shipping it out to HartRAO we put it through its paces in the lab ensure functionality, stability, and reliability. This chapter describes the suite of tests used to evaluate the assembled receiver as a whole.

5.1 Bandpass Measurements

The analog signal conditioning of the Southern C-BASS receiver was introduced in Figure 3.17. A double-sideband mixer (see Section 4.12) modulates the 4.5 GHz–5.5 GHz C-BASS band with a 4.5 GHz local oscillator, producing a DC–1.0 GHz baseband. The anti-aliasing filters described in Section 4.9 separate the 1000 MHz baseband into the low, DC–500 MHz, sub-band, and the high, 500–1000 MHz, sub-band. The digital backend digitizes the two sub-bands using iADC boards and carries out digital signal processing, (DSP) using two ROACH boards.

The signal path outlined above has four relevant bandpasses summarized in Table 5.1. These are measured using different techniques.

This section presents measurements of the bandpass at each of these stages of signal conditioning.

Description	Frequency	Measurement
RF Band	4.5 GHz–5.5 GHz	Spectrum analyzer
Full Baseband	DC–1 GHz	Spectrum analyzer
Sub-band 1	DC–500 MHz	Spectrum analyzer and digital backend
Sub-band 2	500–1000 MHz	Spectrum analyzer and digital backend

Table 5.1: Different bandpass measurements through the C-BASS signal path.

5.1.1 Spectrum Analyzer Measurements

The first bandpass of interest is the RF band. This is measured after the RF signal conditioning introduce in Section 3.5. This bandpass provides information about the horn, the cryogenic passive components and the LNAs. The importance of these measurements has already been revealed in the discovery of the horn resonance structure in Figure 3.7. To measure the bandpass, I placed a 300 K load at the aperture of the horn and measured the power spectrum using a spectrum analyzer after the RF signal conditioning. The results are in Figure 5.1. In this state the resonances of Figure 3.7 are hidden by the thermal noise, however the plot clearly shows the band pass defining filters and the bandpass shape. Unfortunately the measurement does not have sufficient dynamic range to show the full extent of the filter out-of-band signal suppression.

The spectrum analyzer was also used to measure the baseband signals. These are the signals we digitize with the iADC after mixing down.

Measuring the baseband passband with a spectrum analyzer presents a challenge since we nominally use 4.5 GHz as our modulating signal to the mixer. We cannot measure the roll-off at the low end of the band, since the spectrum analyzer cannot measure negative frequencies. The two sub-bands measured using a 4.5 GHz LO are shown in Figure 5.2a.

In order to measure the low frequency side I moved the baseband up in frequency by changing the modulating signal from 4.5 GHz to 4.0 GHz. This has the effect of shifting the baseband up by 500 MHz, from DC–1000 MHz to 500–1500 MHz. This is graphically illustrated in Figure 5.2b. It is worth noting that the power splitters used to separate the two sub-bands are only rated up to 1 GHz. This accounts for the noticeable slope at the higher frequencies.

These spectrum analyzer measurements demonstrate that the system works as designed. The two sub-bands are well separated by the anti-aliasing filters, and the general passband is flat enough that the digital backend can be used to remove any remaining structure. The C-BASS band is well defined by the RF 4.5-5.5 GHz bandpass filters.

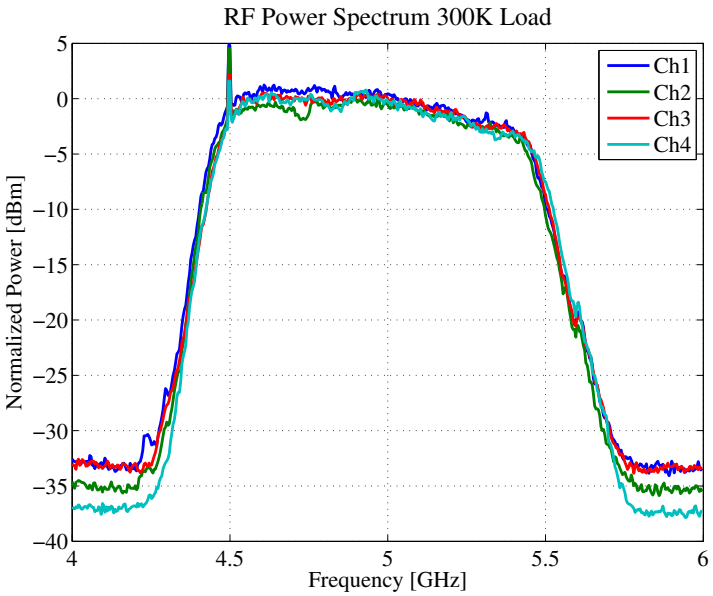

Figure 5.1: Output power with a 300 K load on the horn. Plots have had a gain adjustment applied to illustrate the similarility of the passband shapes. The 4.5 GHz LO used to modulate the RF down to baseband is clearly visible in the plot.

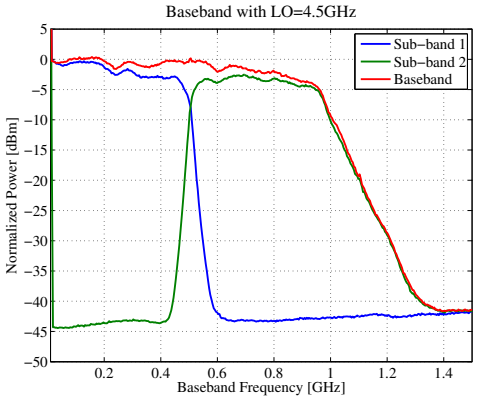

(a) The two digitized bands, DC–500 MHz and 500–1000 MHz

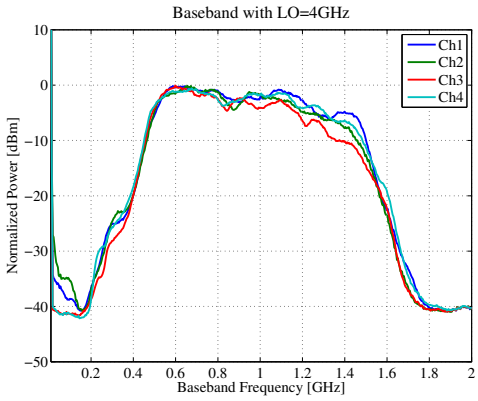

(b) Full Band measured by changing to 4GHz LO. This shows the low side rolloff. The steep slope after 1GHz is due to the power splitter frequency response.

Figure 5.2: Measurements of the baseband sub-bands digitized by the iADC.

5.1.2 Digital Backend Passband Measurement

The bandpass can also be measured by injecting a test continuous wave, or C.W. , signal through the 30 dB couplers inside the cryostat (see Figure 3.17). This narrow-bandwidth signal has a known, stable power level. The C.W. frequency is swept through the band, and the detected output power is simultaneously recorded by the digital backend. This provides a direct measurement of the transmission as a function of frequency, as well as the channel-to-channel isolation. We can see both the effectiveness of the 4.5-5.5 GHz band defining filter, and the anti-aliasing filters used to separate the two sub-bands. See Section 4.8 and Section 4.9 for further details on these components

The results are plotted in Figure 5.3. We discard 13 channels, or $\sim$100 MHz, on the edge of each band to reduce out-of-band noise from leaking into the band. These results show that the 500 MHz low pass/high pass filters designed in Oxford work very well. However the filter roll-off between the two sub-bands, defined by the anti-aliasing filters of Section 4.9, is superior to the roll-off at the edge of the band, defined by the 4.5-5.5 GHz filters of Section 4.8. The out-of-band roll-off could be improved either by increasing the roll-off of the 4.5-5.5 GHz RF bandpass filter, or by shifting the DC–500 MHz sub-band up in frequency by using a different mixing frequency, allowing the low-frequency anti-aliasing filters to be used to define the C-BASS band. The plot also shows the effect of 'scalloping loss'. This is the signal loss that occurs between frequency channels as a result of the non-square shape of the individual channel responses (see Figure 3.21), and results in the slightly fuzzy appearance of the bandpass. This could be improved by increasing the number of taps in the PFB; however since the plot shows that the 'scalloping loss' is small ($< 0.2dB$ r.m.s.) compared to the overall signal level, this will not significantly affect the system sensitivity.

5.1.3 Possible Improvements

A pertinent question given the results of this section is whether there is any way of improving or optimizing the bandpass for C-BASS measurements? One option that has been given serious consideration would be to use two different LO signals to generate the sub-bands. If we used a 4 GHz and 4.5 GHz for the low sub-band and high sub-band respectively, we could locate both sub-bands between 500–1000 MHz, rather than the current DC–500 MHz and 500–1000 MHz split. Aside from standardizing the required filters, the band definition would be significantly improved as we would be able to use the superior low-frequency filters, rather than the less effective RF bandpass filters. This option, while not yet complete at the time of writing, is likely to be implemented in the near future.

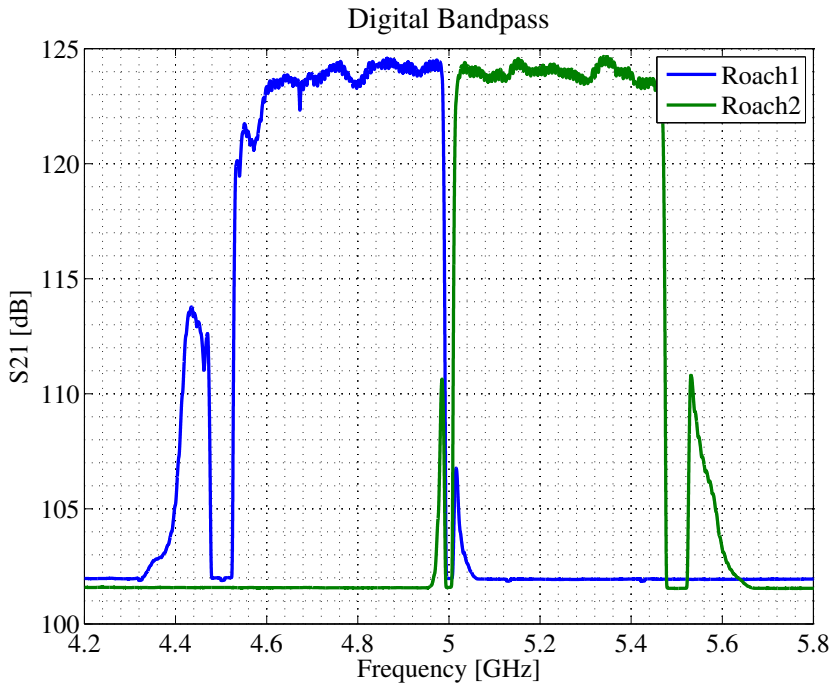

Figure 5.3: The C-BASS digital bandpass, measured by injecting a signal though the 30 dB couplers, sweeping the frequency across the C-BASS band and recording detected power. The low frequency filters (Section 4.9) used to separate the two sub-bands have steeper roll-off than the 4.5-5.5 GHz bandpass filters (Section 4.8) . The zeroed regions on the edge of the band, were applied on the FPGA to reduce aliasing. The 'fuzziness' of the detected power level across the frequency band is due to the 'scalloping loss' between the backend frequency channels. This is because of the non-square shape of the channel frequency response in Figure 3.21.

5.2 Low Noise Amplifiers

Measuring the noise temperature of modern, 3 K, cryogenic low noise amplifiers (LNAs) is difficult as the noise from other components can easily swamp the LNA noise itself. One way of avoiding this is to use the cold attenuator method (Fernandez, 1998). In this method, a noise diode is coupled to the input of the cryogenically cooled LNA through a large-value attenuator, typically of order 20 dB. Both the LNA and the attenuator are typically cooled to between 4 K and 20 K to take full advantage of the superior LNA noise at these temperatures. When the noise diode is on, the LNA sees the hot temperature of the noise diode (typically 15000 K) attenuated by 20 dB, and the physical temperature of the attenuator in the cryostat. However when the noise diode is off, the LNA sees the much lower cold temperature of the noise diode (typically room temperature, $\sim$290 K) again attenuated by 20 dB, and the physical temperature of the attenuator. We made use of this method to measure the noise temperature of all four LNAs simultaneously, using the test setup in Figure 5.4. The noise diode signal was split by a four-way Wilkinson power divider, and the four outputs sent through a cryogenic 20 dB attenuator into the LNA. The attenuators were at the physical temperature of the cold plate, or 10 K. The output power in both states was measured to determine the individual LNA Y-factors and is described in the following section. The test setup is shown in Figure 5.4.

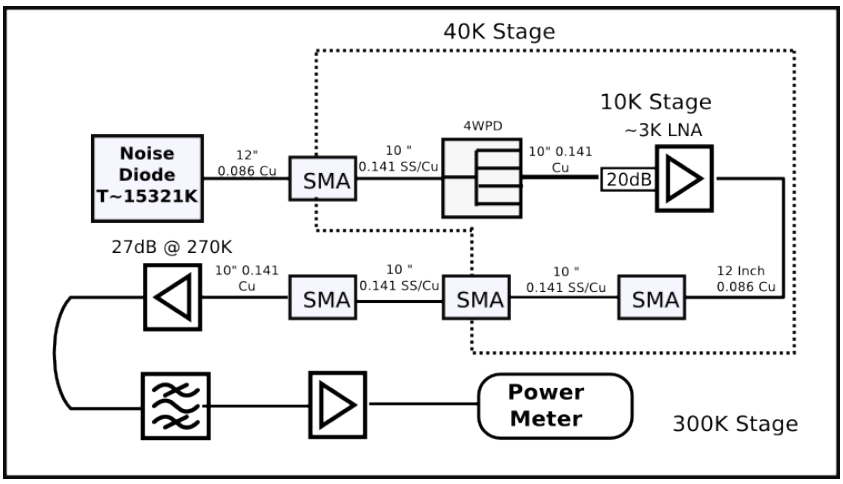

Figure 5.4: Measuring the LNA noise temperature. A four-way power divider is used to split the noise diode calibration signal to simultaneously measure all four amplifiers. The results are contained in Table 5.3

5.2.1 Y Factor

A number of different methods exist to measure the temperature of an RF component, T_{noise}. In these calculations we use the Y-factor method. We present a known temperature to the

T_{hot}[K]	T_{cold}[K]	S21[dB]	L	Component
15300±50	290±1	1.3±0.1	0.74±0.02	Noise Diode
11560±340	290±1	6.2±0.1	0.24±0.01	First cable Sections and SMA
2800±160	80±5	0.7±0.1	0.85±0.02	Four Way Power Divider
2400±300	70±5	20.5±0.2	0.01±0.04	Second Cable Section
32±8	10±0.1			20 dB attenuator

Table 5.2: The effective noise diode temperature at the input to the LNA. The temperature seen after a lossy component is given by $T = LT_{in} + (1 - L)T_{phys}$. From this we estimate a $T_{hot} = 32 \pm 8$ K and $T_{cold} = 10 \pm 0.1$ K. The small uncertainty on the cold value is due to the well-constrained temperature of the cold attenuator. Component uncertainties quoted are estimated from noise on typical measurements and temperature uncertainties are from propagating the component uncertanties.

receiver chain, vary it, and record the receiver output power for a variety of temperatures. T_{noise} is simply the noise we would detect with zero input power i.e. they intercept of a plot of measured power vs input temperature. To estimate this we need a minimum of two measurements done at different temperatures, T_{hot} and T_{cold}. The power we would detect in these two states is then given by:

$$P_{hot} = G(T_{hot} + T_{rx}) \tag{5.1}$$

$$P_{cold} = G(T_{cold} + T_{rx}) \tag{5.2}$$

Taking the ratio of the two powers i.e Y=$\frac{P_{hot}}{P_{cold}}$ we solve for the y-axis intercept using

$$T_{noise} = \frac{T_{hot} - YT_{cold}}{Y - 1} \tag{5.3}$$

The AS5254 noise diode has an ENR of 17 dB in the 4.5–5.5 GHz band (see Table 4.1). This is equivalent to a hot noise diode temperature of 15300 K, while the cold temperature is given by the physical temperature, nominally 290°C. The signal is then attenuated by the various components in the RF path, shown in Figure 5.4, with the running noise diode temperature summarized in Table 5.2. Given these losses, we can calculate both the hot temperature (∼32 K), and the cold temperature (∼10 K) of the noise diode at the input to the LNA. Using these we are able to calculate T_{LNA} using Equation 5.3.

Measuring the $Y = P_{hot}/P_{cold}$ ratio we obtain the results summarized in Table 5.3:

5.3 Long-term Stability

The C-BASS continuous-comparison architecture is complicated. The reason for this complexity is to improve the long term stability of the receiver, and in particular to reduce the $1/f$ gain fluctuations that affect the the total intensity measurements. In the C-BASS LNAs the $1/f$ noise begins to dominate over the white noise at frequencies lower than the $1/f$ knee frequency, $f_{knee}=$ 1 Hz as previously shown in Figure 1.9. The azimuth scanning strategy outlined in Section 1.6 requires minimizing the $1/f$ noise associated with an azimuthal scan; since the C-BASS scan speed will be 4 °/s, a full 360° scan will require 90 s to complete and we need stabilize the measurements such that $f_{knee}< \frac{1}{90}$ Hz. The continuous-comparison technique achieves this, but has the effect of adding noise to the measurement. This typically reduces the sensitivity by a factor of $\sqrt{2}$. See Section 1.8.3 for a full discussion on the sensitivity implications of the continuous-comparison architecture.

The rest of this section will discuss the tests carried out to demonstrate the effectiveness of this approach, and in particular its implementation on the digital backend.

5.3.1 Stability Tests

The location of the receiver while undergoing lab testing precluded using the sky as a load. We thought of using a 77 K liquid nitrogen cooled load, however it was not possible to heat the reference load to 77 K to balance this. Instead we decided to bypass the optics of the system, and in preference terminated the input of the receiver with 50 ohm resistive loads inside the cryostat at the temperature of the coldhead ($\sim$10 K). This allowed us to simulate the final configuration providing a $\sim$10 K sky noise temperature and a balanced $\sim$10 K reference load noise temperature.

Terminated Input with Unstabilized Cold Head Temperature

To carry out the tests we terminated four linear-to-circular convertor inputs with 50 ohm loads. The 50 ohm loads were left to track the unstabilized temperature of the cold head as seen in

Channel	Y[dB]	T_{LNA}[K]
1	4.00±0.01	4.4±1.2
2	4.15±0.01	3.6±1.0
3	4.15±0.01	3.6±1.0
4	4.05±0.01	4.1±1.1

Table 5.3: Measured LNA Noise temperatures. These are largely consistent with the typical datasheet noise temperature of 2.6 K.

Figure 5.5a.

It rapidly became apparent that this was not a good way to do this test. Judging by the correlation between measured power and coldhead temperature in Figure 5.5b, the temperature of the input resistor appear to track the unstabilized cold plate temperature fluctuations very closely. Unfortunately this effect is indistinguishable from the $1/f$ gain variations. An improvement, described in the following section, was made by stabilizing the physical temperature of the input resistor.

Temperature-controlled input

The new load (described in Section 4.1) has four temperature controlled resistors, making it possible to provide two temperature-controlled references in addition to the two stabilized continuous-comparison references. While this is not sufficient to provide stabilized loads to the four linear to circular inputs, we can bypass the linear-to-circular convertor, and provide two stabilized sources at the input to the noise diode couplers. This is satisfactory, since the passive linear-to-circular does not contribute significantly to long term variations in the final configuration.

As a result we rearranged the cryostat into the configuration in Figure 5.6a. In this configuration the variations in the load temperature are reduced to negligible levels, a better configuration for accurately measuring the receiver gain variations. The improvement is clearly shown in Figure 5.6b.

With the input temperature stabilized we can properly measure the effectiveness of the continuous-comparison architecture. We see that the total intensity f_{knee} is improved to $\sim$11 mHz, corresponding to a period of 90 s. This is a significant improvement over the unstabilized amplifiers which exhibit a f_{knee} of 1 Hz and maintain the stability of the receiver over a 90 s azimuth scan. As expected the polarization measurements do not show any evidence of $1/f$ gain variations.

5.4 System Linearity, Conversion Factors, and T_{rx}

In this section we investigate the effect of increasing the input temperature by a known amount. Consider the system presented in Figure 5.6a. We can control the temperature of the resistors. Changing the temperature of these resistors while measuring the detected backend power allows us to check a number of important system characteristics. Firstly, it allows us to confirm the linearity of the backend measurements. Secondly, it allows us to measure the scale-factor between the input temperature and backend measurement units (B.U). Thirdly we can use this

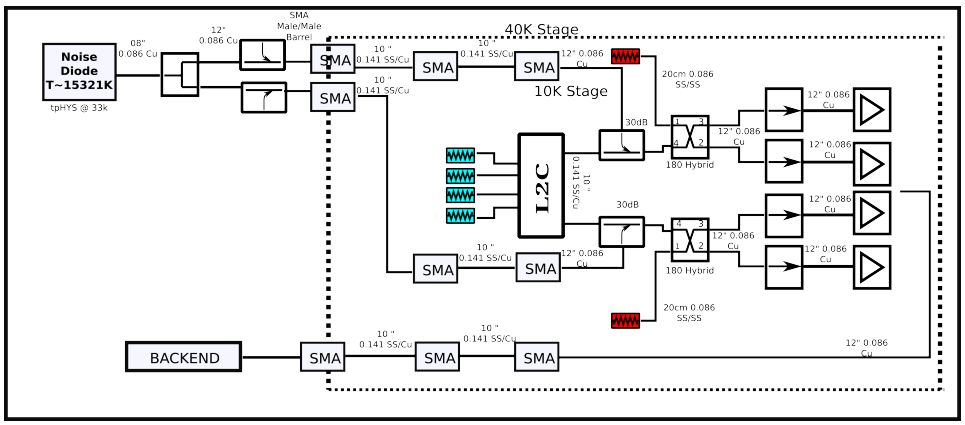

(a) Blue coloured resistors were (nominally) at the temperature of the unstabilized cold head, while the red coloured resistors were kept at a controlled 11.000±0.001 K using the Cryocon Temperature controller

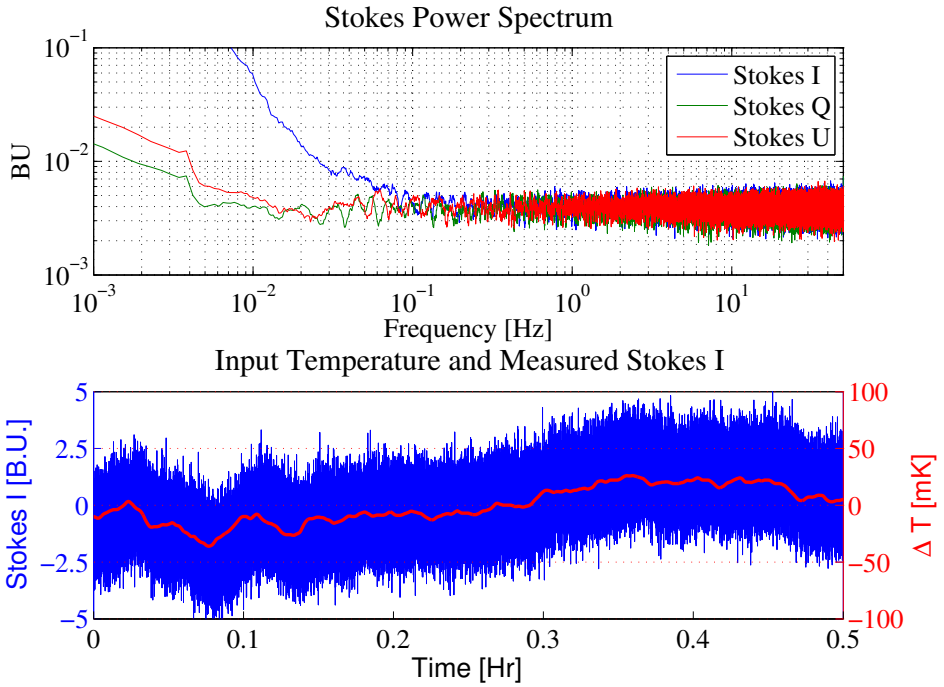

(b) The first plot shows the Stokes I, Stokes Q and Stokes U detections. The $1/f$ noise begins to dominate the white noise below about 50 mHz in Stokes I. The load temperature and the detected Stokes I are visibly correlated. We cannot measure the receiver stability with this setup since the $1/f$ knee frequency will be dominated by long term temperature changes of the input.

Figure 5.5: First attempt to measuring intrinsic receiver stability.

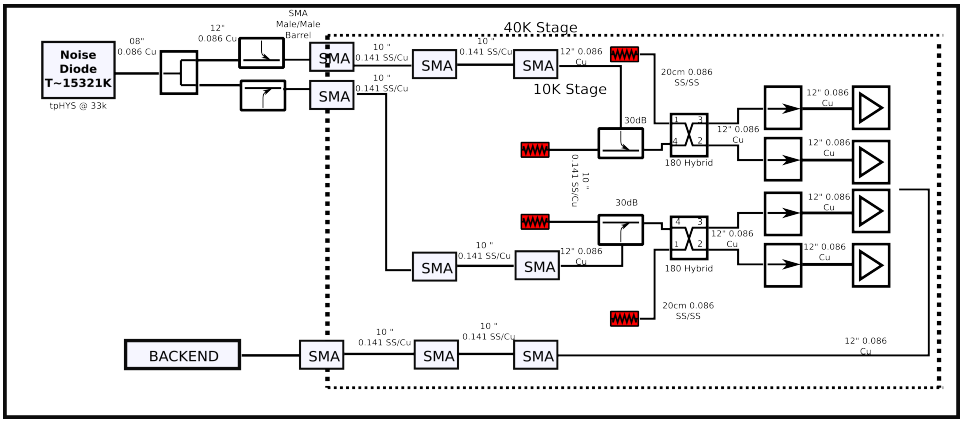

(a)

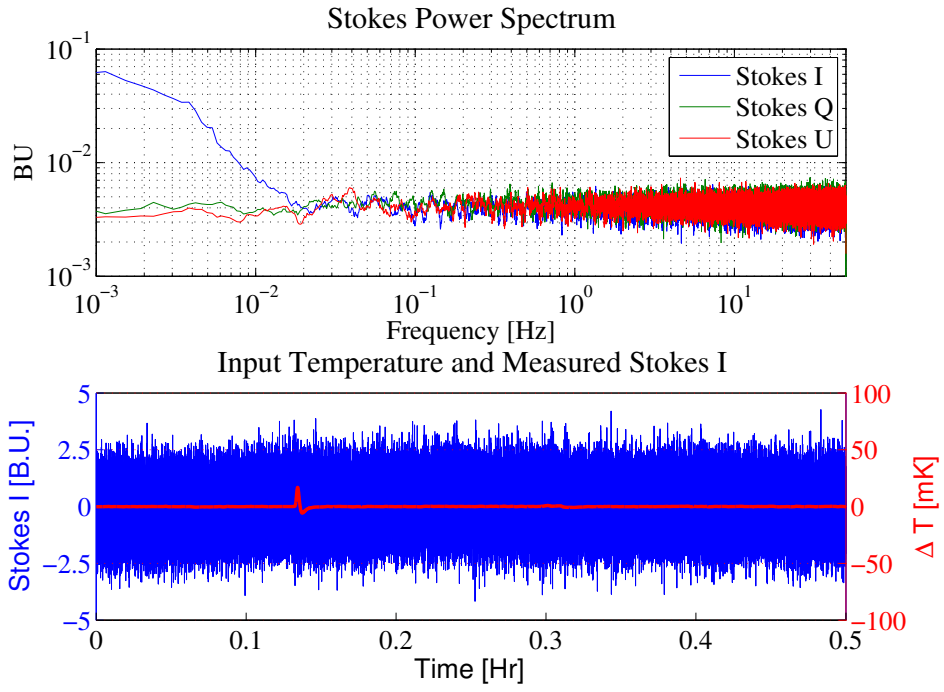

(b)

Figure 5.6: Second attempt to measuring intrinsic receiver stability. The red coloured resistors are temperature controlled to 12.00±5 mK using the Cryocon temperature controller with feedback from a Cernox temperature sensor (see 4.1.1. While the absolute accuracy of sensor calibration is only quoted to 5 mK, the relative stability of the temperature controlled loop is within 1 mK r.m.s. This removes the cold load variation effect on the detected Stokes I signal seen in Figure 5.5b. This is a better reflection of the intrinsic receiver stability with f_{knee}∼6 mHz. Polarization does not suffer from any $1/f$ noise as expected.

Component	T_{phys}	Gain (dB)	G[Lin]	$\prod G$	T_{rx}[K]
Coupler	12.0±0.5 K	-0.2±0.01	0.950±0.002	0.950±0.002	0.50±0.02
Hybrid	12.0±0.5 K	-0.5±0.05	0.891±0.010	0.851±0.011	1.91±0.10
Isolator	12.0±0.5 K	-0.5±0.05	0.891±0.010	0.758±0.019	3.43±0.17
LNA	T_{amp}=3.6±1.0 K	40±0.1	10000±230	7600±400	11.18±1.16
Cable	40±5 K	-1.2±0.1	0.8	5700±410	11.18±2.59

Table 5.4: Expected $T_{\rm rx}$ for the Southern system without the optics.

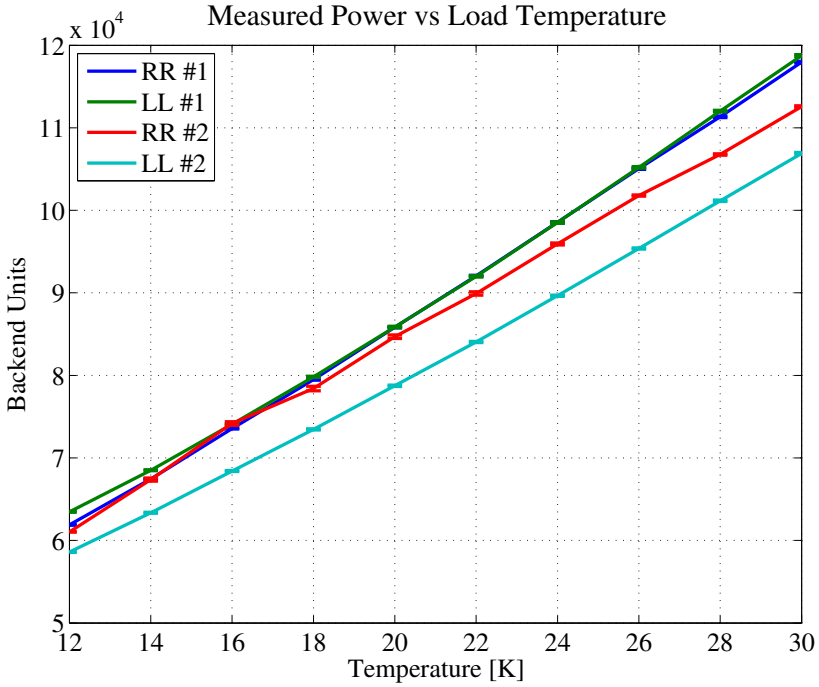

Figure 5.7: Detected power against the input resistor temperature. Conversion between B.U. and K is ∼2700:1.

test to measure $T_{\rm rx}$ using the Y factor method outlined in Section 5.2.1. We estimate the $T_{\rm rx}$ to be ∼11.18±2.59 K using the Friis Cascaded noise formula shown in Table 5.4.

The results are plotted in Figure 5.7. The system behaves linearly between 10 K and 30 K and we can also estimate that the conversion factor between B.U. and Kelvin is ∼2700 B.U.:1K, and that $T_{\rm rx}$ estimated using the Y-factor method is between 8.1±1.4K and 9.5±1.7K consistent with our estimate of 11.18±2.59K in Table 5.4.

5.5 Examining the thermal noise

The thermal noise level of the receiver provides the most direct measurement of the sensitivity of any total power radiometer. The higher the thermal noise, the less sensitive a receiver.

Channel	T_{hot}	T_{cold}	Y	$T_{\rm rx}$ [K]
RCP #1	30±0.5	15±0.5	1.65±0.1	8.1±1.4
LCP #1	30±0.5	15±0.5	1.65±0.1	8.1±1.4
RCP #2	30±0.5	15±0.5	1.61±0.1	9.5±1.7
LCP #2	30±0.5	15±0.5	1.62±0.1	9.5±1.6

Table 5.5: The $T_{\rm rx}$ of the four 500 MHz channels measured using Y-factors in the test setup shown in Figure 5.6a.

Typically we aim to have the thermal noise about 5 times smaller than the minimum desired detectable signal.

The r.m.s. of the thermal noise is given by the well known radiometer equation,

$$\sigma_{RMS} = \frac{T_{sys}}{\sqrt{\tau \Delta\nu}}, \tag{5.4}$$

where T_{sys} is the System Temperature, τ is the integration time, and $\Delta\nu$ is the bandwidth.

5.5.1 Measured Thermal Noise

The effective $T_{\rm sys}$ of the stabilized measurement, T_{com} given by the arrangement in Figure 5.6a is given by,

$$T_{com} = \sqrt{(T_{sky} + T_{rxSky})^2 + (T_{load} + T_{rxLoad})^2}. \tag{5.5}$$

In this configuration, $T_{sky} = 12{\pm}0.1\ K$, $T_{load} = 12{\pm}0.1\ K$ and we estimate $T_{\rm rx}$ ∼8.7±1.5K using the four measurements made in Section 5.4. Using these values together with Equation 5.5, we estimate $T_{com}= 29.3 \pm 4.5\ K$. The C-BASS integration period is $\tau= 0.01\ s$, and the effective bandwidth (after discarding 13 of the 128 channels as described in Section 5.1.2) is $\Delta\nu= 900\ MHz$.

With these numbers we use Equation 5.4 to estimate (using Equation 5.4) $\sigma_{RMS}= 9.8 \pm 1.5\ mK$ which compares favourably with the noise in Figure 5.8b of $\sigma_{RMS}= 10.2\ mK$. The noise also reduces as expected with increasing bandwidth. This is shown in Figure 5.8 with the noise reducing by a factor of $\sqrt{2}$ with a doubling of the bandwidth, $\Delta\nu$. These data are sampled at 100 Hz; the convention is to quote sensitivity in mK$\sqrt{s}$ which would be a factor of $\sqrt{100} = 10$ better than the numbers indicated on the plots.

Sensitivity and Required Integration Time

To estimate the sensitivity to Stokes I, Q, and U, I plotted the respective measurement thermal noise in Figure 5.9 and summarized them in Table 5.7.

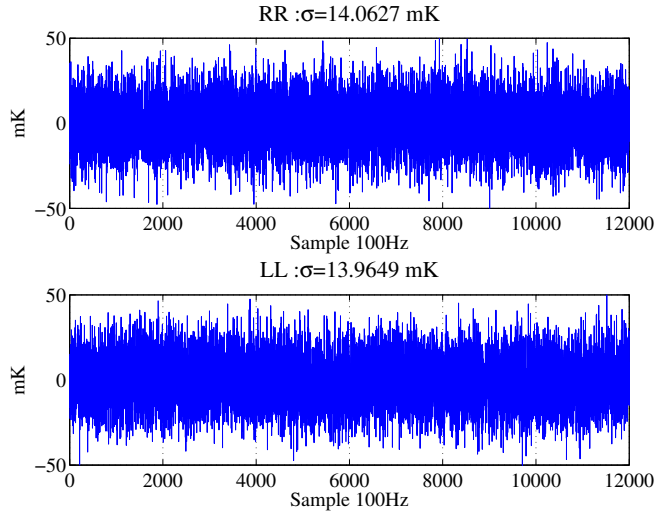

(a) 450 MHz bandwidth

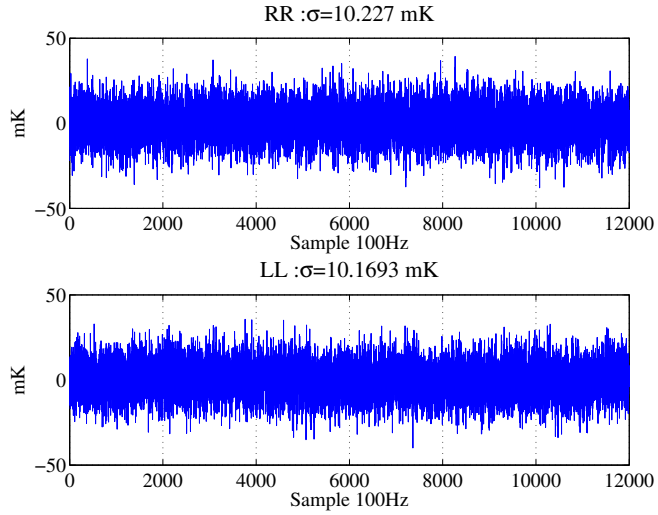

(b) 900 MHz bandwidth

Figure 5.8: Thermal noise of the C-BASS Stokes I data with bandwidth. Here we compare 450 MHz bandwidth against 900 MHz, and show that the noise integrates down by $\sqrt{2}$ as expected from Equation 5.4. We are approximating the sky by a 12 K resistor as seen in Figure 5.6a. The conversion to Kelvin is done using the 2700B.U/K derived in Section 5.4. To get the thermal noise sensitivity in the convention mK$\sqrt{s}$, simply divide by 10 as suggested by Equation 5.4

Channel	Bandwidth [MHz]	Noise [mK]	T_{com}
RCP	450	14.08	44K
LCP	450	13.97	44K
RCP	900	10.23	31K
LCP	900	10.18	31K

Table 5.6: Thermal noise from the receiver with the input to the coupler terminated by a 12K temperature controlled 50 ohm resistor (stabilized by reference against a cold load) as introduced earlier with Figure 5.6a. As expected we see that the thermal noise decreases by $\sqrt{2}$ when we increase the bandwidth by a factor of 2.

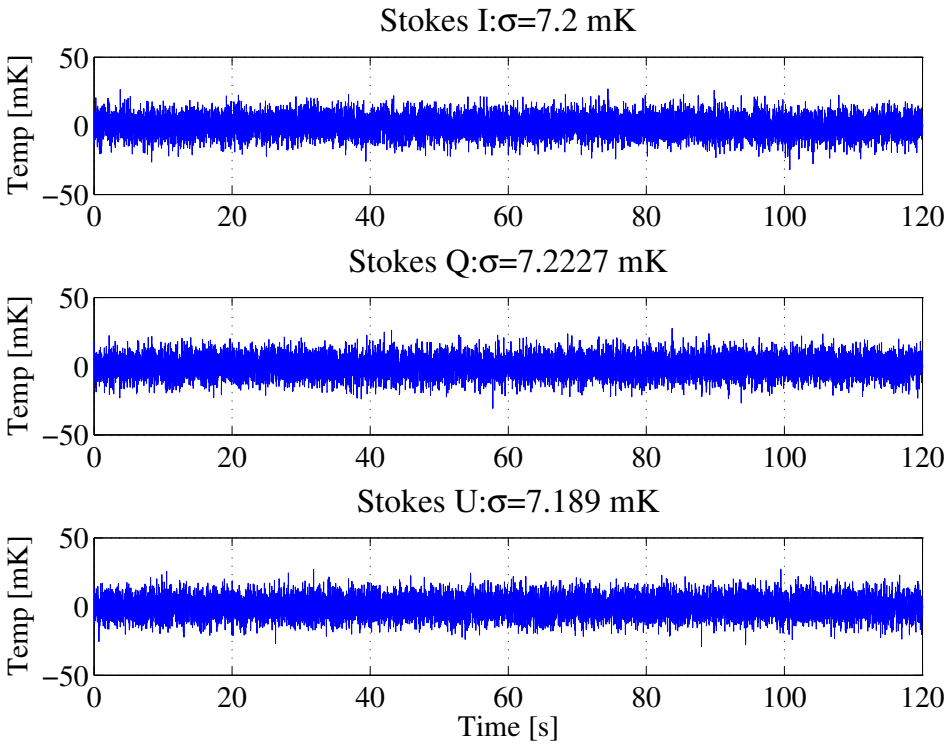

Figure 5.9: Thermal noise on the three C-BASS Stokes measurements in the configuration in Figure 5.6a. This is with a 100 Hz sampling period. A 1 Hz sampling period would be 10 times lower.

Measurement	Sensitivity [mK$\sqrt{s}$]
Stokes I	0.720
Stokes Q	0.722
Stokes U	0.719

Table 5.7: Thermal noise on the three C-BASS Stokes measurements in mK$\sqrt{s}$

The C-BASS sensitivity from the plots and table is 0.72 mK$\sqrt{s}$ with the resistor temperature of 12 K. We will require 15 seconds of integration time per beam to achieve the C-BASS sensitivity requirement of 0.13mK/pixel with these levels. This is easily achievable with the integration times estimated in the double-elevation scanning strategy of in Section 1.6.

5.6 Polarization Tests

We were able to test the response to a polarization both by using an optical source and by rotating the polarization of the calibration noise diode. This is possible because of the digital phase shifters, described in Section 4.3, placed in the noise diode path.

5.6.1 Optical Polarization Tests

It is relatively easy to generate and rotate a polarized signal. Parallel conductors spaced at a fraction of the wavelength can be used as a polarizer, and these can be rotated in front of a radiation source such as the sky. We manufactured two varieties of polarizers, shown in Figure 5.10a, with a conductor spacing of $\sim$0.5 mm and placed these between a 300 K load and the aperture of the horn. Rotating the polarizer rotates the polarization angle of the input signal.

Of course, rotating the polarizer is difficult without a custom-built jig. Nonetheless the results of this test, shown in Figure 5.10, do show that the digital backend can detect the polarization angle of a signal.

5.6.2 Noise Diode Polarization Tests

The Southern receiver includes facility to rotate the polarization angle of the noise diode, an improvement over the Northern receiver where the noise diode calibration signal is seen as an entirely Stokes Q signal. This is made possible by the digital phase shifters in the noise diode injection paths.

Figure 5.11 shows the Stokes measurements while I gradually changed the phase shift applied to the noise diode signal in one of the noise diode injection paths, while keeping the other steady. The noise diode signal is rotated from Stokes Q into Stokes U and back again. This allows us to calibrate the polarization measurements in both of the two measurements directly against the noise diode. The measurement has had smoothing applied to remove the quantization of the 6 bit digital phase shifters.

(a) Two polarizing grids made out of welding rods (Left) and aluminium tape (Right).

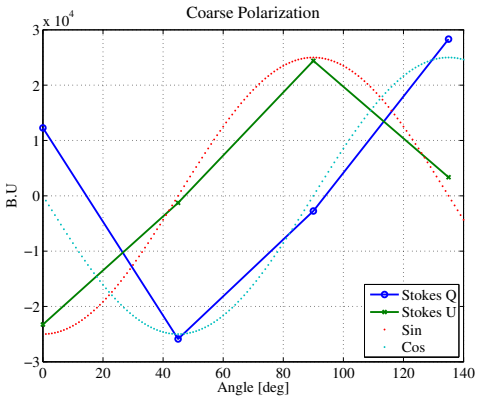

(b) Detected Stokes Q and Stokes U (in digital backend units) plotted against the (coarse) angular rotation of the polarizing optical assembly. I have overplotted sin and cos curves to demonstrate that this (roughly) works as expected.

Figure 5.10: Polarization response to rotating polarized signal. We tested the backend response to a polarized signal by placing the polarizing grid seen in Figure 5.10a, between the 300K eccosorb load and the horn aperture. Rotating the polarizing grid rotates the polarization angle of the detected radiation. The results are seen in Figure 5.10b.

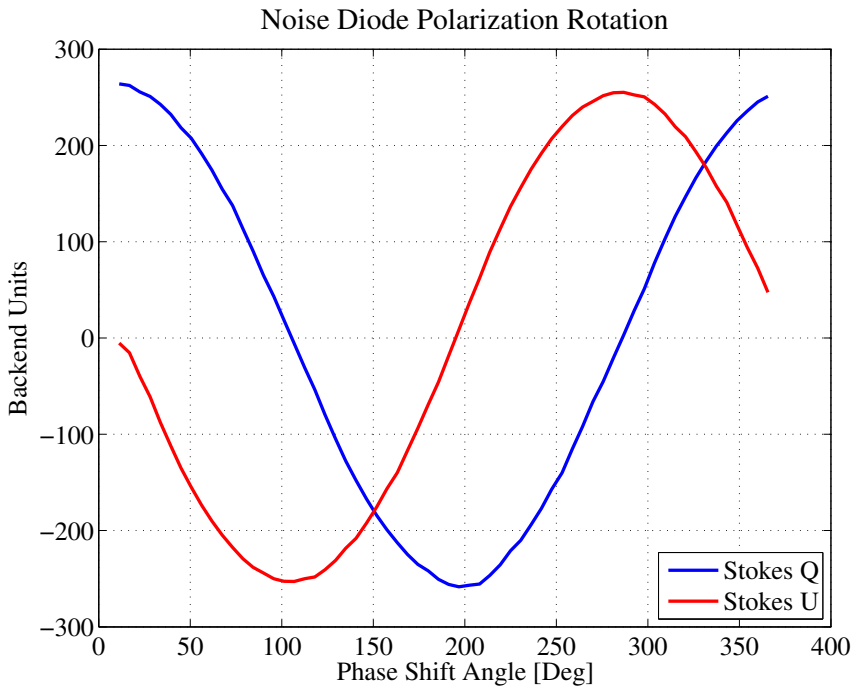

Figure 5.11: Polarization response to rotating noise diode calibration signal. We also tested the backend response to a polarized signal by rotating the polarization of the noise diode signal. We applied different phase shifts to one of the noise diode arms and measured the detected Stokes Q and Stokes U.

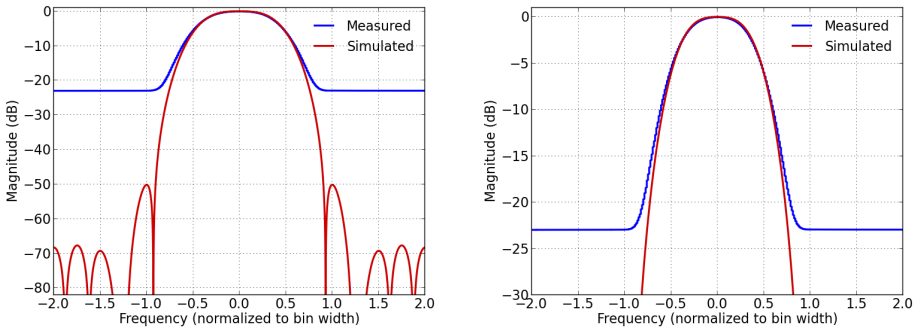

(a) (b)

Figure 5.12: Measured digital cross-channel isolation. Measured by sweeping a frequency through the digital backend and monitoring the detected power on adjacent channels. We expect greater than 50 dB isolation. Measurements, limited by the dynamic range of the digital backend, show a minimum of ∼22 dB.

5.7 Cross-channel Isolation

I have described the polyphase FFT algorithm in Section 3.7.6. This implementation of the FFT algorithm allows a windowing function to be efficiently applied to the data in order to improve channel isolation. For C-BASS we use a Hamming window function with a 4-tap polyphase filterbank. The measured and expected channel response is given in Figure 5.12. This shows the significant sidelobe suppression (greater that 45 dB) achieved by windowing the data with the Hamming window. Unfortunately the ∼22 dB dynamic range of the digital backend limited the level to which we were able to measure. The FFT stages would overflow with very large signals and as 22 dB of dynamic range is sufficient for C-BASS I did not pursue this further. However, the plots show that there is a minimum of 22 dB cross-channel isolation.

5.8 Conclusion

I have described a number of tests of the receiver in this chapter. These have all demonstrated that the digital C-BASS receiver functions as expected. Of particular importance are the tests of the receiver gain stabilization and sensitivity, all of which meet the instrument requirements. Probably the most pleasing of the results are the polarization detections demonstrated by shifting the phase of one of the noise diode calibration signals in Figure 5.11.

In the following chapter, I will present preliminary measurements made with the C-BASS receiver on the sky during tests at HartRAO. These will demonstrate the functionality of the servo control system, and will also begin to quantify the current receiver sensitivity on the sky.

Chapter 6

Commissioning the Southern Instrument

> Well, I think we tried very hard not to be overconfident, because when you get overconfident, that's when something snaps up and bites you.
>
> Neil Armstrong

Commissioning of the Southern receiver is currently underway at HartRAO. A great deal remains to be done before moving the receiver to the final Northern Cape site; this work is ongoing. This chapter provides details on the current status of the commissioning.

6.1 Optical Pointing

The first antenna pointing model was made by taking photographs of stars. We observed the positions of stars over a number of nights using an inexpensive Mintron 12V6HC-EX CCD camera suitable for taking astronomical images. This was coupled with a f4.5-5.6 300 mm Nikon zoom lens and installed on the telescope together with CCTV cable running from the camera location to the control PC (see Figure 6.1) where images could be taken using the Sensoray 611 PCI frame grabber. Offsets in star positions extracted from the images can be fitted against antenna pointing direction and used to calculate a pointing model. The entire process was largely automated for C-BASS North by Stephen Muchovej and an identical process was followed for C-BASS South after the servo control had been integrated with the high levels control system. The optical pointing provides a good test of this integration. The procedure is as follows:

1. Drive to nominal star position
2. Grab frame from optical camera

Figure 6.1: CCD camera installed on the telescope.

3. Find star position in picture frame using image processing software

4. Calculate antenna angle offset from frame centre using frame angular size

5. Correct antenna position offset to centre the star

6. Record offsets as a function of mean antenna positions

The usual process is to calculate a rough pointing model using the CBI pointing model outlined in Shepherd, M. (2001) by following an object from rise to set across the sky. This provides a crude pointing model in Figure 6.2a. After applying the crude model, it is usually possible to drive to sources anywhere in the sky and have them within the $2^\circ \times 2^\circ$ field-of-view[1]. Observations of pointing offsets distributed across the sky allow us to refine the pointing model in Figure 6.2b.

6.2 Receiver Temperature Measurements

The expected $T_{\rm rx}$ for the Southern receiver with horn is $T_{\rm rx}$ =15.9 K (see Table 1.4). The $T_{\rm rx}$ without the horn was measured in Section 5.4 as $T_{\rm rx}$ $\sim$8.7$\pm$1.5K. This section describes three different measurements of $T_{\rm rx}$ with the horn in place.

[1]FOV (degrees) = (CCD Size(mm) x 57.3) / focal Length (mm) where CCD size 12mm and focal length is 300mm

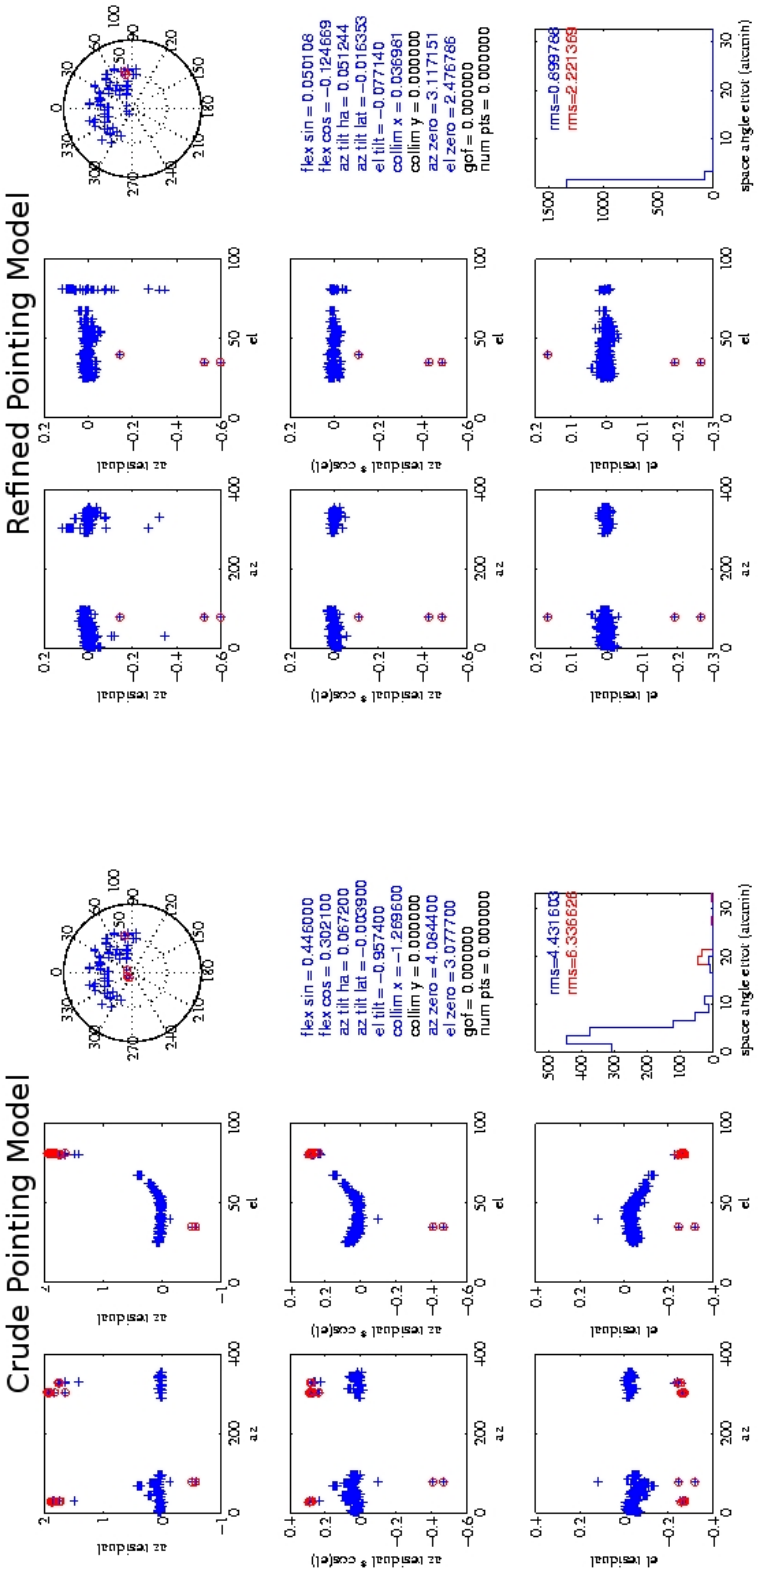

(a) Crude pointing model

(b) Refined pointing model

Figure 6.2: The optical pointing model fitted to the C-BASS South Antenna

Frequency [GHz]	4.5-5.0	5.0-5.5
Y(dB)	7.8±0.2	6.8±0.2
T_{rx}(K)	20.6±2.0	30.2±1.0

Table 6.1: Measured Y-factors using the power meter and T_{hot} =298±1 K, T_{cold} =6±1 K, and T_{load} =12±0.5 K. The $T_{\rm rx}$ in the lower half of the band is about 10 K lower than in the upper half of the band.We used a modification of the standard Y-factor method using T_{hot}=($T_{eccosorb} + T_{load}$)/2, and T_{cold}=($T_{sky} + T_{load}$)/2.

In all three tests we measured the Y-factor by comparing a hot, 300 K Eccosorb RF absorber against a cold, 6 K sky.

When measuring $T_{\rm rx}$ we measure a combination of both the sky power and the load power and we need to take this into account. For the C-BASS receiver the modified input temperatures are given by T_{hot}=($T_{eccosorb} + T_{load}$)/2, and T_{cold}=($T_{sky} + T_{load}$)/2.

6.2.1 Power Meter

The power meter allows us to measure the $T_{\rm rx}$ in the upper and lower 500 MHz bands independently, and gives the results in Table 6.1. We note that the 4.5-5.0 GHz band has a lower $T_{\rm rx}$ than the 5.0-5.5 GHz band.

These results show the $T_{\rm rx}$ is higher than the equivalent measurement carried out on the Northern system in Table 2.2, where the $T_{\rm rx}$ =13.6 K. We are still investigating this difference.

6.2.2 Spectrum Analyzer and Digital Backend

The spectrum analyzer and digital backend can measure $T_{\rm rx}$ as a function of frequency. We again need to account for the addition of the reference load noise as in the power meter measurements above. The results are given in Figure 6.3.

These estimates of $T_{\rm rx}$ are consistent with the power meter measurements. The $T_{\rm rx}$ in the upper half of the band is higher than the lower half of the band. The cause of this is yet unknown, but will be investigated in the coming weeks. The frequency structure in the bottom half of the frequency range is because of the horn S11 response shown previously in Figure 3.8. This has been removed by remachining the horn corrugations.

6.2.3 Receiver Sensitivity on Sky

All three forms of measurement seem to be in general agreement; $T_{\rm rx}$ is ∼20 K in the lower half of the band and ∼30 K in the upper half of the band.

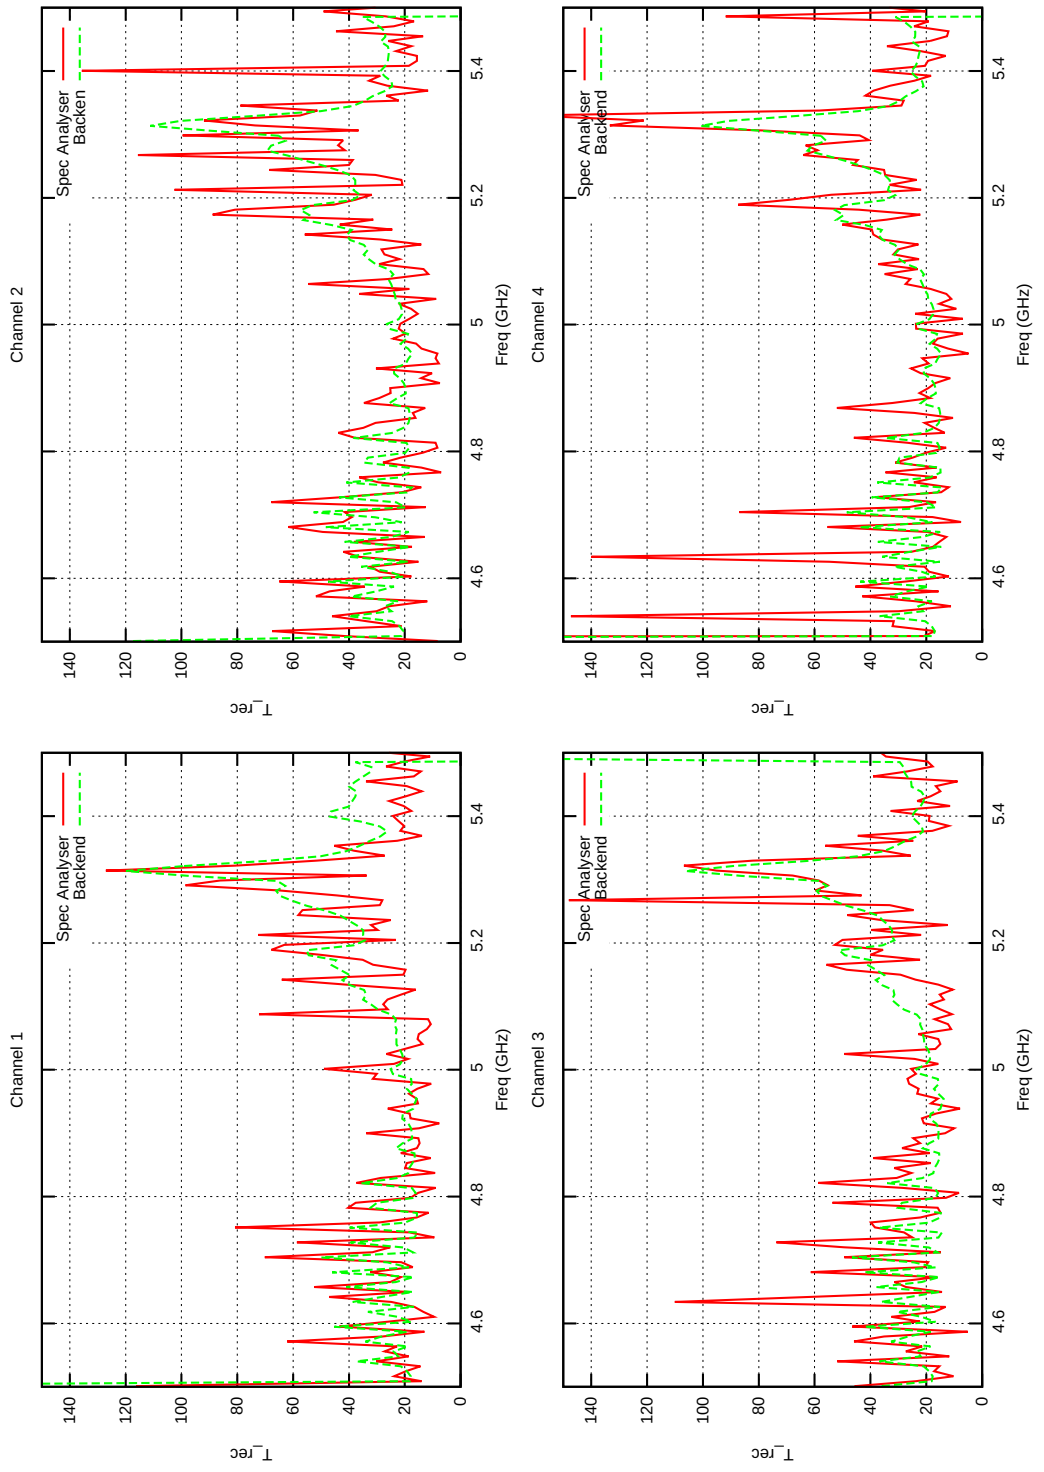

Figure 6.3: Spectrum analyser (red) and digital backend (green) measured T_{rx} . These are consistent with the power meter measurements.

The sensitivity of the receiver depends on the system temperature, $T_{\rm sys}$, and we need to take the additional noise from the atmosphere and CMB into account. At 5 GHz we expect these to contribute $\sim$6 K. This allows us to estimate the effective $T_{\rm sys}$, T_{com} using the appropriate equation for the C-BASS radiometer derived earlier in thesis and repeated here,

$$T_{com} = \sqrt{(T_{sky} + T_{rxSky})^2 + (T_{load} + T_{rxLoad})^2} \tag{6.1}$$

If we assume a 12 K temperature for the cold load and the T_{rxLoad} =10 K, this gives us an estimated T_{com}=34-40 K

We derived the sensitivity of the continuous-comparison architecture back in Equation 1.47 and Equation 1.48. Using this estimate of T_{com} we can calculate the r.m.s. noise on Stokes I, Stokes Q, and Stokes U as between 1.1 mK$\sqrt{s}$ and 1.3 mK$\sqrt{s}$, much higher than the lab tests results in Table 5.7. To reach the target of 0.1 mK noise level we would require $\sim$140 s of integration time. This is much longer than budgeted for in the survey strategy.

It is important that we identify the cause of the excess noise and remove it.

6.3 Sky power spectra

The continuous-comparison architecture of the C-BASS receiver mitigates gain fluctuations and stabilizes the Stokes I measurements. Tests of this functionality were conducted in during the lab commissioning in Section 5.3. In this section I describe the same tests repeated measuring the sky.

These tests are in Figure 6.4. The stability on the long timescales is significantly improved using the continuous-comparison stabilization, although the $1/f$ knee frequency is worse than it was in lab tests, due to variations in the atmosphere. The polarization measurements do not require stabilization as the correlation operation removes the uncorrelated gain fluctuations.

6.4 Solar Observations

The sun is not resolved by the 30° HPBW of the horn, so with the receiver on the telescope, we scanned the horn through the Sun to check that the beam pattern matched the beam pattern measured in Figure 3.6. This observation measures the horn beam pattern and also provides a calibration source for the system in this state.

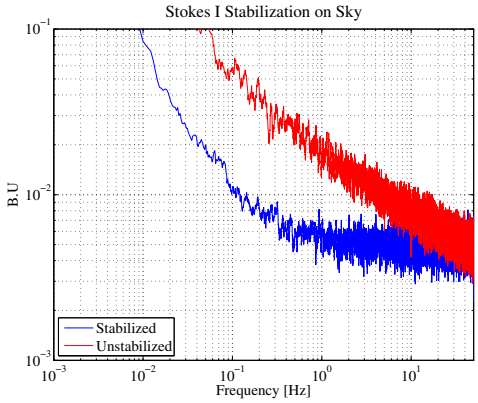

(a) Stokes I stabilization

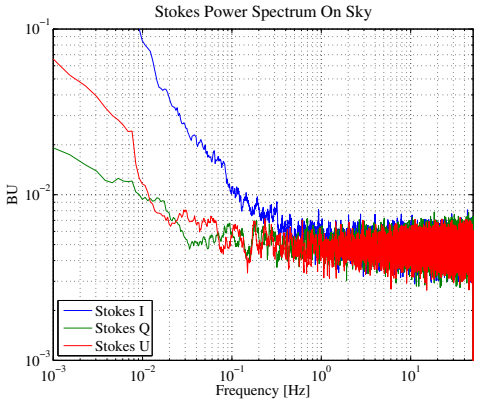

(b) Total intensity and polarization power spectrums

Figure 6.4: C-BASS noise power spectrum measurements on sky receiver and feed-horn only. The the continuous-comparison architecture improves the Stokes I stability significantly, although the $1/f$ knee is not as low as it was in lab tests due to the fluctuating atmosphere. Polarization is fairly stable. The power spectra resemble those in Figure 5.5b with an unstabilized 50 ohm load

6.4.1 Horn Beam Patterns

A series of azimuth scans were conducted through the sun at 36° elevation. In order to calculate the sky angle, θ_{sky}, we need to correct for the projection effect of the elevation angle, θ_{El},

$$\theta_{sky} = \theta_{Az} cos(\theta_{El}) \tag{6.2}$$

After doing this we arrive at the plot in Figure 6.5. This is a plot of right circular polarization (RCP) intensity against sky angle offset. Encouragingly, the measured beam shape is very close to the simulation.

6.4.2 System Equivalent Flux Density

A useful metric in receiver sensitivity is the System Equivalent Flux Density or $SEFD$. The $SEFD$ is defined in Equation 6.3 and captures both the T_{sys} and A_e and allows comparison of system sensitivities. It is the unit of flux that doubles the system temperature, T_{sys} .

$$SEFD = 2k_b T_{sys}/A_e \tag{6.3}$$

In these scans I only used the bottom half of the band, and did not employ the continuous-comparison mode of the radiometer, so T_{sys} is given by,

$$T_{sys} = T_{sky} + T_{rx} \tag{6.4}$$

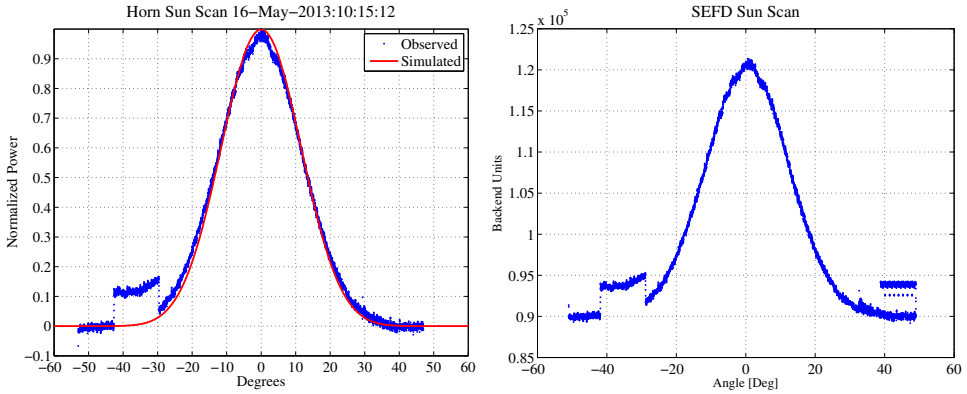

(a) Baseline subtracted and normalized

(b) With baseline, unnormalized

Figure 6.5: An azimuth cross scan through the Sun, (a) as measured with the digital backend and (b) with baseline subtracted and beam normalized to the maximum power. The angle is in sky angle coordinates, and we see a noise diode event at angle -45°.

Location	Flux [kJy]
Learmonth	1770
San Vito	1860
Sag Hill	1690
Palehua	1880
Average	1800 ±67

Table 6.2: Solar flux measurements (Space Weather Prediction Center, National Oceanic and Atmospheric Administration, 2013) taken on the 16 May 2013

In the configuration described, $T_{sys} = 6 \pm 1\ K + 20.6 \pm 2.0\ K = 26.6 \pm 3\ K$ (see the T_{rx} measured in Table 6.1). From the effective area of the horn, $A_e = 0.0146 \pm 0.0003\ m^2$ (see Table 1.5), and substituting into Equation 6.3, we expect $SEFD = 5.1 \pm 1.3$ MJy.

One of the useful properties of the $SEFD$ as a metric is that it can be measured directly using a source of known flux, $S(\nu)$. If we define the power detected when observing the source as P_{on} and the power detected when not observing the source as P_{off} then

$$\frac{P_{on}}{P_{off}} = \frac{SEFD + S(\nu)}{SEFD} = \frac{S(\nu)}{SEFD} + 1 \tag{6.5}$$

Rearranging this it is possible to describe the $SEFD$ by

$$SEFD = S(\nu)(\frac{P_{on}}{P_{off}} - 1)^{-1}. \tag{6.6}$$

I used Equation 6.6 to measure the $SEFD$ of the receiver with the bare horn using the Sun

	P_{on}/P_{off}	S_ν	$SEFD_{meas}$	$SEFD_{exp}$
Horn	1.34±0.01	1800±67 kJy	5.1±1.3 MJy	5.26±0.23 MJy
Secondary		1800±67 kJy	-	4.6±1.2 KJy

Table 6.3: The C-BASS $SEFD$. The measured $SEFD_{meas}$ of C-BASS is comparable to the expected $SEFD_{exp}$ from the T_{sys} measurements. We will measure the $SEFD$ for the full optical assembly soon to complete the table.

cross scans and the solar flux measurements at 4995 MHz summarized in Table 6.2 (Space Weather Prediction Center, National Oceanic and Atmospheric Administration, 2013). The $SEFD$ measurements are in Table 6.3, which also includes a consistent estimate of the $SEFD$ using Equation 6.3 with appropriate values of A_e and $T_{\rm sys}$ for the bare horn. We have not yet measured the $SEFD$ of C-BASS with the secondary optics in place, but this is estimated to be $\sim$4.6 kJy using Equation 6.3 and assuming the effective area in Table 1.5 and a similar $T_{\rm sys}$.

6.5 Conclusions

In this chapter I have shown that the servo control and high level control system are functioning correctly, and that we are both able to do an optical pointing model of the C-BASS dish, and drive the receiver through the Sun while recording the backend output data.

The stabilization on Stokes I measurements provided by the continuous-comparison architecture appears to be functioning correctly when used to stabilize the sky measurements, however atmospheric turbulence increases the $1/f$ knee frequency anyway.

I have also shown the results of $T_{\rm rx}$ measurements of the receiver with the horn in place. The different methods used to make these measurement produce consistent results, however these are higher than we expect from the $T_{\rm rx}$ measurements without the horn. We are still investigating this. Scans through the Sun produce consistent beam shapes with those measured earlier in the thesis, and also measure and $SEFD$ which is consistent with the $T_{\rm rx}$ measurements mentioned earlier.

Chapter 7

Thesis Conclusions

> After climbing a great hill, one only finds that there are many more hills to climb.
>
> Nelson Mandela

This thesis has presented my contributions to the C-BASS project and, where clearly stated, the work of others in the C-BASS collaboration. My work has largely been to design, build, integrate and test the Southern telescope, but it has also included technical work to support the Northern instrument.

I have presented a method of converting former telecommunications to radio astronomy platforms that includes measurement of the primary reflector structure and a customizable and inexpensive servo controller. This has been presented as it was implemented for C-BASS but it could be used for similar projects in the future. The digital receiver complements this, as it can be trivially changed to cover a larger, or different, frequency range by scaling the sampling hardware, or changing the mixer frequency.

I will conclude this thesis by summarizing the scope of this work and place the current status of the C-BASS project into context.

7.1 The C-BASS Science Case and Instrument

This thesis began by describing the CMB, and in particular the importance of CMB polarization studies. Galactic foregrounds are likely to limit the effectiveness of the CMB polarization surveys — I described the physics underlying these emissions and outlined the importance of sky surveys in understanding them. A few methods of foreground removal were also presented.

The C-BASS survey addresses a number of other important science cases. These include: 1) improving our understanding of the magnetic fields distribution throughout our Galaxy, and

2) aiding in the detection of new regions of AME. The recently published results in Carretti et al. (2013), reporting massive magnetized outflows from the Galactic centre, show the importance of large scale surveys of our own Galaxy.

I then introduced the survey requirements of C-BASS. In particular I focused on the C-BASS receiver architecture and the difficulty associated with the survey sensitivity goals. I presented data which illustrated the effectiveness of the C-BASS instruments in meeting these challenges.

I concluded the first chapter by describing the C-BASS telescopes and in particular their novel optical design.

7.2 The Northern Survey

The Northern C-BASS survey was discussed in Chapter 2. The ambit of this thesis did not extend to a full description of the Northern Survey or the current survey results. I did, however, describe the lessons learned during the Northern instrument deployment. A number of these were used in design and manufacture of the Southern instrument. In particular I learned the importance of designing the receiver with serviceability in mind. I have designed the Southern receiver with easy accessibility to the RF connections and the vacuum flanges and it is also possible to run the Southern receiver completely decoupled from the telescope, which is not possible on the Northern system.

My primary-work on the Northern instrument was the design and manufacture of a set of narrow-band cryogenic notch filters to remove in-band RFI from the analog backend. Despite their success, I conclude that a digital receiver, such as the digital backend described in this thesis, would be more effective.

The Northern survey is expected to be completed in the next 6 months. Current projections indicate that the survey will reach the target sensitivity within this time frame.

7.3 The Southern Telescope

The rest of the thesis outlined the development of the Southern telescope. I divided this roughly as follows: Chapter 3 provided a system-level overview of the instrument, Chapter 4 was a detailed description of the design and manufacture of individual components used in the receiver, Chapter 5 provided results of tests conducted in the lab before deployment, and Chapter 6 described the preparatory instrument commissioning at the HartRAO test site.

In addition to describing the C-BASS receiver itself, Chapter 3 also described the process employed to convert an unused telecommunications antenna into a platform suitable for de-

ploying the C-BASS instrument. This involved accurately measuring the dish surface, double-checking these measurements with radio observations, and developing (and deploying) an antenna servo-controller using inexpensive and readily available hardware. These techniques are easily replicable and can be used for similar projects in the future. I described the cryogenic front-end, and especially the improvements made over the Northern receiver. I provided a detailed overview of the Southern C-BASS receiver, presenting a diagram of the entire C-BASS signal path and described the CASPER signal processing hardware used for the digital backend and calculate the theoretical channel-channel isolation provided by the new backend. This is important for RFI excision. The chapter concluded by describing the higher level control system — adapted from similar code used on the Northern instrument — and the establishment of the Northern Cape site to accomodate the C-BASS instrument when finally deployed.

Having described the high-level instrumentation of the Southern C-BASS, Chapter 4 dealt with the individual components. Almost all are bespoke C-BASS designs. Some of these were re-used from the Northern instrument, but a significant number were developed specifically for the Southern instrument. The components designed (by myself) for the Southern instrument include: an improved temperature control mechanism for the C-BASS reference load, a temperature stabilization platform inside the cryostat to stabilize the noise diode calibration signals, low frequency (DC–2 GHz) power dividers, low frequency (DC–2 GHz) anti-aliasing filters, low frequency (DC–3 GHz) baseband amplifiers, double sideband mixers, 3–8 GHz digital phase shifters, and a stabilized LNA bias supply. These components can be easily repurposed for future receiver developments. Extensive test data are provided for all of these.

Before the instrument left Oxford, we spent time conducting a series of wide-ranging lab tests. These are outlined in Chapter 5. I began with extensive characterization of the C-BASS bandpass. Since the C-BASS instrument requires extreme sensitivity, measurements were made of the $1/f$ suppression achieved by the continuous-comparison architecture. I monitored the output power while varying the input temperature of the temperature-controlled load. This evaluated the linearity of the system, measured the receiver temperature, and quantified the conversion between backend measurement units and antenna temperature. The chapter includes explicit examination of the r.m.s. noise measured by the receiver, an important metric that presents the final limit on the instrument sensitivity. I concluded the chapter by showing that the receiver is sensitive to polarized signals – both through the horn and through the phase-shifted noise diode signal. I also demonstrate that the digital backend provides a minimum of 22 dB isolation between digital channels.

In Chapter 6 I briefly described preliminary tests conducted on the receiver at the first-stage on-sky commissioning conducted at HartRAO. The antenna pointing was evaluated using an optical camera integrated with the high level control system. These tests double as a test of

the servo-controller introduced in an earlier chapter. The receiver temperature – including the optics – was then measured using a power meter, spectrum analyzer, and the digital backend. Following this, I presented the continuous-comparison architecture applied to total-intensity sky-data, and included a power spectrum of the polarization detection. The final part of the chapter included a few tests with the receiver in-place on the antenna. Unfortunately, the full optics are still not installed but plots taken using just the horn were shown. Azimuth scans taken through the Sun were used to estimate the system temperature, and obtained results consistent with the receiver temperature measurements.

7.4 Receiver Improvements

I have mentioned areas of possible future improvements through the course of this thesis. Probably the most useful of these would be to use a 4.0 GHz mixing frequency to move the DC–500 MHz sub-band up to 500–1000 MHz. This will allow us to standardize both sub-bands and also to use the low frequency filters to entirely define the frequency band. These have a much steeper roll-off and would reduce any potential for out-of-band interference. This would remove the need for the RF band-defining filters entirely and would make scaling or changing the backend frequency coverage as easy as changing the frequency of the mixer.

7.5 The Future of C-BASS

The Northern survey is ongoing and will be completed in the near future.

The Southern observations will begin in earnest once we have completed the preliminary commissioning. We hope to move the receiver to the Northern Cape site in March 2014. The only current concerns are the excess-noise in the measurements with the optics in place and the fine structure present in the low-half of the horn passband. We believe we have fixed the fine structure by remachining the horn. We hope to evaluate this fix and work on the excess noise in the coming weeks.

Once these are understood, the commissioning at HartRAO will proceed. The first step will be to put the full optical assembly in place and check that this does not add significantly to the $T_{\rm sys}$. An important measurement for the HartRAO commissioning will be a high dynamic-range measurement of the telescope beam pattern. This was not possible on the Northern telescope measurements due to the mechanical limits of the telescope frame and will help in our understanding of the final foreground maps. We will make these beam pattern measurements by mounting a 5 GHz rectangular horn on a nearby tower and scanning the antenna through this source. The measurement will provide both an excellent measurement of the beam pattern

and a good calibration source. We hope to be able to do useful scientific observations from the HartRAO site but a major aim will also be to soak-test the receiver to ensure the systems function reliably when we move to the more-isolated Northern Cape.

C-BASS has a well developed data reduction pipeline with excellent map-making software. Throughout the commissioning period, we will be integrating the Southern receiver output with the data reduction pipeline and we should be able to produce preliminary Southern sky maps before the move to the Northern Cape. This will provide a good (and very useful) indicator that the systems are functioning as we expect.

The final Northern Cape site construction has been completed and is ready for the instrument to be installed. The site will provide an excellent location for making the highly sensitive C-BASS all-sky maps. The elevated position and clear horizon minimizes the need for complex ground subtraction templates and will allow us to achieve the optimal survey sensitivity. The excellent survey weather conditions will allow us to obtain maximum time on the sky, with little down time. Completion of the Southern survey is anticipated for early 2016.

With the release of the *Planck* results earlier in the year, the importance of low frequency surveys for foreground subtraction has become clear yet again. The C-BASS survey fills an important frequency gap and will provide critical information to CMB polarization surveys. In addition, the mapping of the low frequency synchrotron distribution provided by C-BASS is an excellent way of improving our understanding of the Galactic magnetic field.

A final word is that the impact of the C-BASS survey does not need to end with the final maps. A great deal of the technology outlined in this thesis can be reused on other experiments. I hope that this thesis will provide both a useful roadmap to achieving this.

Appendix A

Radiometer Equation

The definition of T_{total} given in Equation 1.29 is a standard χ^2 distribution of v_{total} with one degree of freedom. This has a probability distribution (Abramowitz & Stegun (1964) Eq. 26.4.1) given by:

$$P(T) = \frac{1}{\sqrt{2\pi}\sigma} T^{-\frac{1}{2}} e^{-T/2\sigma^2} \tag{A.1}$$

$$(T >= 0) \tag{A.2}$$

We can evaluate the variance using the standard definition namely:

$$\sigma_T^2 = < T^2 > - < T >^2 \tag{A.3}$$

To evaluate this we need to calculate $< T^2 >$ and $< T >$.
Firstly we note the following relationship:

$$\int_0^{+\infty} x^n e^{-ax} dx = \frac{\Gamma(n+1)}{a^{n+1}} \tag{A.4}$$

where $\Gamma(n)$ is the Gamma function such that:

$$\Gamma(n) = (n-1)! \tag{A.5}$$

We continue by calculating the expectation value of the detected quantity, $< T >$

$$< T > = \int_0^{+\infty} TP(T) dT \tag{A.6}$$

$$= \frac{1}{\sqrt{2\pi}\sigma} \int_0^{+\infty} T T^{-\frac{1}{2}} e^{-T/2\sigma^2} dT \tag{A.7}$$

We now make use of Equation A.4 to rewrite this using $n = \frac{3}{2}$ and $a = \frac{1}{2\sigma^2}$:

$$\frac{1}{\sqrt{2\pi}\sigma}\int_0^{+\infty} T^{\frac{1}{2}} e^{-T/2\sigma^2} dT = \frac{1}{\sqrt{2\pi}\sigma}\frac{\Gamma(\frac{1}{2}+1)}{(\frac{1}{2\sigma^2})^{\frac{1}{2}+1}} \tag{A.8}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}\frac{\Gamma(\frac{3}{2})}{(\frac{1}{2\sigma^2})^{\frac{3}{2}}} \tag{A.9}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}(\frac{1}{2}\pi^{\frac{1}{2}})2^{\frac{3}{2}}\sigma^3 \tag{A.10}$$

$$= \sigma^2 \tag{A.11}$$

This important result tells us that the average noise power equals the variance of the noise voltage. This average power is known as the *System Temperature*, T_{sys} (already introduced in Equation 1.8) and is a metric of the receiver sensitivity.

We follow a similar process to calculate the quantity $< T^2 >$

$$< T^2 > = \int_0^{+\infty} T^2 P(T) dT \tag{A.12}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}\int_0^{+\infty} T^2 T^{-\frac{1}{2}} e^{-T/2\sigma^2} dT \tag{A.13}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}\int_0^{+\infty} T^{\frac{3}{2}} e^{-T/2\sigma^2} dT \tag{A.14}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}\frac{\Gamma(\frac{3}{2}+1)}{(\frac{1}{2\sigma^2})^{\frac{3}{2}+1}} \tag{A.15}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}\frac{\Gamma(\frac{5}{2})}{(\frac{1}{2\sigma^2})^{\frac{5}{2}}} \tag{A.16}$$

$$= \frac{1}{\sqrt{2\pi}\sigma}(\frac{3}{4}\pi^{\frac{1}{2}})2^{\frac{5}{2}}\sigma^5 \tag{A.17}$$

$$= 3\sigma^4 \tag{A.18}$$

We now estimate the variance with Equation A.3

$$\sigma_T^2 = < T^2 > - < T >^2 \tag{A.19}$$

$$= 3\sigma^4 - (\sigma^2)^2 \tag{A.20}$$

$$= 2\sigma^4 \tag{A.21}$$

Bibliography

Abramowitz, M. & Stegun, I. A. (1964). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, 9th Edition.

American Technical Ceramics (2012). ATC550L104 Datasheet. (pp. 9001).

Analog-Devices (2011). Low Power , Precision Rail-to-Rail Output Op Amp.

Analog Devices (2011). *REF19x Series*. Technical report, Analog Devices.

Armitage-Caplan, C., Dunkley, J., Eriksen, H. K., & Dickinson, C. (2011). Large-scale polarized foreground component separation for Planck. MNRAS, 418, 1498–1510.

Armitage-Caplan, C., Dunkley, J., Eriksen, H. K., & Dickinson, C. (2012). Impact on the tensor-to-scalar ratio of incorrect Galactic foreground modelling. MNRAS, 424, 1914–1924.

Barkats, D., Bischoff, C., Farese, P., et al. (2005). First Measurements of the Polarization of the Cosmic Microwave Background Radiation at Small Angular Scales from CAPMAP. ApJ, 619, L127–L130.

Barkats, D., Bischoff, C., Farese, P., et al. (2005). Cosmic Microwave Background Polarimetry Using Correlation Receivers with the PIQUE and CAPMAP Experiments. ApJS, 159(1), 1–26.

Baumann, D., Jackson, M. G., Adshead, P., et al. (2009). Probing Inflation with CMB Polarization. volume 1141 of *American Institute of Physics Conference Series* (pp. 10–120).

Bennett, C., Hill, R., Hinshaw, G., et al. (2003). First-year Wilkinson microwave anisotropy probe (WMAP) observations: foreground emission. ApJS, 148, 97.

Berex (2010). BG14A BG14A 5-4000 MHz Cascadable InGaP HBT Gain Block.

Bhattacharyya, A. K. (2001). High-Q resonances due to surface waves and their effects on the performances of corrugated horns. *IEEE Antennas and Propagation*, 49(4), 555–566.

Burbidge, G. R. (1956). On Synchrotron Radiation from Messier 87. ApJ, 124, 416.

Cain, J. (2010). *Parasitic Inductance of Multilayer Ceramic Capacitors*. Technical report, AVX.

Carretti, E. (2010). Galactic foregrounds and CMB polarization. *arXiv preprint arXiv:1008.4983*, (pp. 1–13).

Carretti, E., Crocker, R. M., Staveley-Smith, L., et al. (2013). Giant magnetized outflows from the centre of the Milky Way. *Nature*, 493(7430), 66–69.

Carretti, E., Haverkorn, M., McConnell, D., et al. (2009). PGMS: to study the Galactic magnetism out of the Galactic plane. volume 36 of *Revista Mexicana de Astronomia y Astrofisica Conference Series* (pp. 9–16).

Carretti, E., Haverkorn, M., McConnell, D., et al. (2010). The Parkes Galactic Meridian Survey: observations and CMB polarization foreground analysis. MNRAS, 405, 1670–1689.

Condon, J. J., Cotton, W. D., Greisen, E. W., et al. (1998). The NRAO VLA Sky Survey. AJ, 115, 1693–1716.

Cordes, J. & Lazio, T. (2002). NE2001. I. A new model for the galactic distribution of free electrons and its fluctuations. *arXiv preprint astro-ph/0207156*, (pp. 1–21).

Dahl, J. (2003). *Time aliasing methods of spectrum estimation*. PhD thesis, Brigham Young University.

Davis, Jr., L. & Greenstein, J. L. (1951). The Polarization of Starlight by Aligned Dust Grains. ApJ, 114, 206.

de Bernardis, Ade, P., Bock, J., et al. (2000). A flat Universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404(6781), 955–9.

de Vos, M., Gunst, A., & Nijboer, R. (2009). The LOFAR Telescope: System Architecture and Signal Processing. *Proceedings of the IEEE*, 97(8), 1431–1437.

Delabrouille, J. & Cardoso, J. (2009). *Data Analysis in Cosmology*, volume 665 of *Lecture Notes in Physics*. Springer.

Deng, G. & Faig, W. (2001). An evaluation of an off-the-shelf digital close-range photogrammetric software package. *Photogrammetric engineering and remote sensing*, 67(2), 227–233.

Dicke, R. H. (1946). The Measurement of Thermal Radiation at Microwave Frequencies. *Review of Scientific Instruments*, 17, 268–275.

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. (1965). Cosmic Black-Body Radiation. ApJ, 142, 414–419.

Draine, B. T. & Lazarian, A. (1998). Electric Dipole Radiation from Spinning Dust Grains. ApJ, 508, 157–179.

Draine, B. T. & Lazarian, A. (1999). Magnetic Dipole Microwave Emission from Dust Grains. ApJ, 512, 740–754.

Dunkley, J., Amblard, A., Baccigalupi, C., et al. (2009a). Prospects for polarized foreground removal. *AIP Conference Proceedings*, 222(1), 222–264.

Dunkley, J., Spergel, D. N., Komatsu, E., et al. (2009b). Five-Year Wilkinson Microwave Anisotropy Probe (Wmap) Observations: Bayesian Estimation of Cosmic Microwave Background Polarization Maps. ApJ, 701(2), 1804–1813.

Elder, F. R., Gurewitsch, A. M., Langmuir, R. V., & Pollock, H. C. (1947). Radiation from electrons in a synchrotron. *Phys. Rev.*, 71, 829–830.

Eriksen, H. K., Banday, A. J., Górski, K. M., & Lilje, P. B. (2004). On Foreground Removal from the Wilkinson Microwave Anisotropy Probe Data by an Internal Linear Combination Method: Limitations and Implications. ApJ, 612, 633–646.

Eriksen, H. K., Dickinson, C., Lawrence, C. R., et al. (2006). Cosmic Microwave Background Component Separation by Parameter Estimation. ApJ, 641(2), 665–682.

Fernandez, J. (1998). A noise-temperature measurement system using a cryogenic attenuator. *TMO Progress Report*, (pp. 1–9).

Finkbeiner, D. P., Davis, M., & Schlegel, D. J. (1999). Extrapolation of Galactic Dust Emission at 100 Microns to Cosmic Microwave Background Radiation Frequencies Using FIRAS. ApJ, 524, 867–886.

Fiore, B. R. (2001). *Capacitors in Broadband Application*. Technical report, American Technical Ceramics.

Fixsen, D. J. (2009). The temperature of the Cosmic Microwave Background. ApJ, 707(2), 916–920.

Futatsumori, S. & Hikage, T. (2008). ACLR Improvement of a 5-GHz Power Amplifier Using High-Temperature Superconducting Reaction-Type Transmitting Filters. *Proc. European Microwave Conference*, (October), 1145–1148.

Galindo, V. (1964). Design of Dual-Reflector Antennas with Arbitrary Phase and Amplitude Distortions. *IEEE Transactions on Antennas and Propagation*, 12(4), 403–408.

Garcia-Garcia, J., Bonache, J., Gil, I., et al. (2005). Comparison of electromagnetic band gap and split-ring resonator microstrip lines as stop band structures. *Microwave and Optical Technology Letters*, 44(4), 376–379.

Gawronski, W. (2001). Antenna control systems: from PI to H. *Antennas and Propagation Magazine, IEEE*, 43(1), 52–60.

Gawronski, W. (2008). *Modeling and Control of Antennas and Telescopes*. Mechanical Engineering Series. Boston, MA: Springer US.

Gold, B., Bennett, C. L., Hill, R. S., et al. (2009). Five-Year Wilkinson Microwave Anisotropy Probe Observations: Galactic Foreground Emission. ApJS, 180(2), 265–282.

Górski, K. M., Hivon, E., Banday, A. J., et al. (2005). HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. ApJ, 622, 759–771.

Grimes, P., King, O., Yassin, G., & Jones, M. (2007). Compact broadband planar orthomode transducer. *Electronics Letters*, 43(21), 1146–1147.

Hall, J. S. (1949). Observations of the Polarized Light from Stars. *Science*, 109, 166–167.

Hanany, S., Ade, P., Balbi, A., et al. (2000). MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of 10°. ApJ, 545, L5–L9.

Hansen, F. K., Banday, A. J., Eriksen, H. K., Górski, K. M., & Lilje, P. B. (2006). Foreground Subtraction of Cosmic Microwave Background Maps Using WI-FIT (Wavelet-Based High-Resolution Fitting of Internal Templates). ApJ, 648, 784–796.

Harris, F. (1978). On the Use of Windows with the Discrete Fourier Transform. *Proceedings of the IEEE*, 66, 51–83.

Haslam, C. G. T., Klein, U., Salter, C. J., et al. (1981). A 408 MHz all-sky continuum survey. I - Observations at southern declinations and for the North Polar region. A&A, 100, 209–219.

Hiltner, W. (1949). Polarization of light from distant stars by interstellar medium. *Science*, 109(2825), 9–10.

Holler, C. M., Taylor, A. C., Jones, M. E., et al. (2013). A Circularly Symmetric Antenna Design With High Polarization Purity and Low Spillover. *IEEE Transactions on Antennas and Propagation*, 61(1), 117–124.

Howes, G. (2004). *Galactic Dynamics with Magnetic Fields*. PhD thesis, University of California, Los Angeles.

Hu, W. & White, M. (1997). A CMB polarization primer. *New Astronomy*, 2(4), 323 – 344.

Jansson, R. (2010). *The Magnetic Field of the Milky Way*. PhD thesis, New York University.

Johnson, J. (1928). Thermal agitation of electricity in conductors. *Physical Review*, 541(1918).

Jonas, J. (2009). MeerKAT The South African Array With Composite Dishes and Wide-Band Single Pixel Feeds. *Proceedings of the IEEE*, 97(8), 1522–1530.

Jonas, J. L., Baart, E. E., & Nicolson, G. D. (1998). The Rhodes/HartRAO 2326-MHz radio continuum survey. MNRAS, 297, 977–989.

King, O. (2009). *C-BASS: C-Band All-Sky Survey*. PhD thesis.

Leach, S. & Cardoso, J. (2008). Component separation methods for the Planck mission. *Astronomy & Astrophysics*, 615, 597–615.

Leitch, E. M., Kovac, J. M., Halverson, N. W., et al. (2005). Degree Angular Scale Interferometer 3 Year Cosmic Microwave Background Polarization Results. ApJ, 624(1), 10–20.

Leitch, E. M., Readhead, a. C. S., Pearson, T. J., & Myers, S. T. (1997). An Anomalous Component of Galactic Emission. ApJ, 486(1), L23–L26.

Luhmann, T. (2010). Close range photogrammetry for industrial applications. *ISPRS Journal of Photogrammetry and Remote Sensing*, 65(6), 558–569.

Lyons, R. G. (2004). *Understanding Digital Signal Processing (2Nd Edition)*. Upper Saddle River, NJ, USA: Prentice Hall PTR.

Maino, D., Burigana, C., Maltoni, M., et al. (1999). The Planck-LFI instrument: Analysis of the 1/f noise and implications for the scanning strategy. *Astronomy and Astrophysics Supplement Series*, 140(3), 383–391.

Matthaei, G and Young, L and Jones, E. (1980). *Microwave Filters, Impedance-Matching Networks, and Coupling Structures*. Artech House Publishers.

Meeks, M. (1976). *Methods of Experimental Physics*. Academic Press.

Miller, A. D., Caldwell, R., Devlin, M. J., et al. (1999). A Measurement of the Angular Power Spectrum of the Cosmic Microwave Background from L = 100 to 400. ApJ, 524, L1–L4.

Mitchell, K. J. & Condon, J. J. (1985). A confusion-limited 1.49-GHz VLA survey centered on α = 13 H 00 M 37 s, δ = + 30 deg 34 arcmin. *The Astronomical Journal*, 90, 1957.

Moore, G. (1965). Cramming more components onto integrated circuits. *Electronics*, 38(8).

Murata (2009). Murata Components Library for AWR Microwave Office.

Murata (2010). *GJM High Frequency Ceramic Capacitors*. Technical report, Murata.

Murata (2011). *Murata Chip Inductors Datasheet*. Technical report, Murata.

Nordling, L. (2012). Recycled dishes form telescope network. *Nature*, 488(7413), 571.

Nyquist, H. (1928). Thermal agitation of electric charge in conductors. *Physical review*, 3(1918).

Oxford Cryosystems (2013). Oxford Cryosystems two stage coldheads. `http://www.oxcryo.com/wp/wp-content/uploads/2010/11/two-stage-coldheads.pdf/`.

Page, L., Hinshaw, G., Komatsu, E., et al. (2007). Three-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations Polarization Analysis. ApJS, 170(2), 335–376.

Parsons, A., Backer, D., & Siemion, A. (2008). A scalable correlator architecture based on modular FPGA hardware, reuseable gateware, and data packetization. *Publications of the Astronomical Society of the Pacific*, (pp. 1207–1221).

Parsons, A., Werthimer, D., Backer, D., et al. (2010). Digital Instrumentation for the Radio Astronomy Community. *The Astronomy and Astrophysics Decadal Survey*.

Parsons, A. R., Backer, D. C., Foster, G. S., et al. (2010). The Precision Array for Probing the Epoch of Re-ionization: Eight Station Results. AJ, 139, 1468–1480.

Penzias, A. A. & Wilson, R. W. (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. ApJ, 142, 419–421.

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. (2013). Planck 2013 results. I. Overview of products and scientific results. *ArXiv e-prints*.

Readhead, A. C. S., Myers, S. T., Pearson, T. J., et al. (2004). Polarization observations with the Cosmic Background Imager. *Science (New York, N.Y.)*, 306(5697), 836–44.

Rees, M. J. (1968). Polarization and Spectrum of the Primeval Radiation in an Anisotropic Universe. ApJ, 153, L1.

Reich, W. (2006). Galactic polarization surveys. *ArXiv Astrophysics e-prints*.

Sachs, R. K. & Wolfe, A. M. (1967). Perturbations of a Cosmological Model and Angular Variations of the Microwave Background. ApJ, 147, 73.

Schmid, H. (2008). Circuits at the Nano Scale. (pp. 95–115).

Schott, G. (1912). Electromagnetic radiation and the mechanical reactions arising from it. *Cambridge University Press*.

Scott, D. & Smoot, G. (2004). Cosmic Background Radiation Mini-Review. *ArXiv Astrophysics e-prints*.

Scott, P. F., Saunders, R., Pooley, G., et al. (1996). Measurements of structure in the cosmic background radiation with the cambridge cosmic anisotropy telescope. *The Astrophysical Journal Letters*, 461(1), L1.

Seiffert, M., Mennella, A., Burigana, C., et al. (2002). 1/f noise and other systematic effects in the Planck-LFI radiometers. A&A, 391(3), 1185–1197.

Shepherd, M. (2001). Pointing the CBI . `http://www.astro.caltech.edu/~mcs/CBI/pointing/`.

Smoot, G. (1998). Galactic free-free and h-alpha emission. *arXiv preprint astro-ph/9801121*, (pp. 1–10).

Smoot, G., Bennett, C., Weber, R., et al. (1990). COBE Differential Microwave Radiometers - Instrument design and implementation. ApJ, 360, 685–695.

Smoot, G. F., Bennett, C. L., Kogut, A., et al. (1992). Structure in the COBE differential microwave radiometer first-year maps. ApJ, 396, L1–L5.

Space Weather Prediction Center, National Oceanic and Atmospheric Administration (2013). Solar Radio Data. `http://www.swpc.noaa.gov/ftpdir/lists/radio/7day_rad.txt`.

Tegmark, M., Eisenstein, D. J., Hu, W., & de OliveiraCosta, A. (2000). Foregrounds and Forecasts for the Cosmic Microwave Background. ApJ, 530(1), 133–165.

Testori, J. C., Reich, P., & Reich, W. (2008). A fully sampled λ21 cm linear polarization survey of the southern sky. A&A, 484, 733–742.

Texas Instruments (2005). *TL074D Datasheet*. Technical Report March, Texas Instruments.

Texas Instruments (2011). *LM4050 / LM4050Q Precision Micropower Shunt Voltage Reference*. Technical report, Texas Instruments.

Tim Westcott (2000). PID without a PhD. `http://www.embedded.com/design/embedded/4211211/PID-without-a-PhD/`.

Tinbergen, J. (1996). *Astronomical Polarimetry*. Cambridge University Press.

Urs, B. & Rangu, M. (2011). Estimation of parasitic elements of multilayer ceramic capacitors. *2011 IEEE 17th International Symposium for Design and Technology in Electronic Packaging (SIITME)*, (pp. 261–264).

Wawrzynek, J. & Brodersen, R. (2005). BEE2 A High-End Reconfigurable Computing System. *IEEE Design and Test of Computers*, 22(2), 114–125.

Welch, J., Backer, D., Blitz, L., et al. (2009). The Allen Telescope Array: The First Widefield, Panchromatic, Snapshot Radio Camera for Radio Astronomy and SETI. *Proceedings of the IEEE*, 97(8), 1438–1447.

Wilson, T., Rohlfs, K., & Hüttemeister, S. (2009). *Tools of radio astronomy*. Springer.

Wolleben, M., Landecker, T. L., Reich, W., & Wielebinski, R. (2006). An absolutely calibrated survey of polarized emission from the northern sky at 1.4 GHz. Observations and data reduction. A&A, 448, 411–424.