

NextBASS:
A New Low-Frequency CMB
Foregrounds Instrument



Jaz R. Hill-Valler
St Edmund Hall
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2019

In memory of my Pappy

Thank you for the many arguments we had when I was a teenager about whether or not Aliens are real. After 8 years of studying space, I can confidently say that I still have no answers for you, but now I know a bit about Bayesian Statistics, so perhaps I'd stand a better chance of winning.

It is befitting that your daughter ended up following something so extraordinary all because of stubbornness, and it's definitely not a coincidence that she ended up mastering a very similar profession to your own.

I would like nothing more than to grab a beer with you and tell you everything I know about the Cosmos. From the Big Bang, past the Quasars and Galaxies, through the Interstellar Dust, beyond the Stars and Exo-Planets. All the way down through our Atmosphere, to the workings of our own, incredible Natural World.

But, I imagine you'd call me a know-it-all and we'd end up chatting about audio engineering.

I miss you everyday and after all these years, I can absolutely agree with you, that analogue is indeed, warmer.

Acknowledgements

As anyone who undertakes a PhD knows, it is unavoidably stressful, isolating but incredibly rewarding. You spend an inordinate amount of time focusing on just one, complicated piece of work, and don't spend much time on the internal workings of your own mind. The only wisdom I can pass on is to always be grateful to others, be proud of yourself, and be humbled by the natural world.

I should like to start by thanking Mike, Angela, Jamie and Luke for convincing me on interview day to join. I've never met a group of people with such an Encyclopaedic knowledge of the bizarre.

To any other residents of the Denys Wilkinson Building, I will always have fond memories of the pub trips where the conversation turned to experimental music and synthesizers.

To Alex, I am pretty sure I still owe you a few beers in the Lamb and Flag. Thanks for all of your endless advice and expertise, you've been an amazing mentor and a good friend throughout my time in Oxford.

Thanks to Josie for all of our much needed stress-relieving tea and doodle therapy sessions. I can't wait to see what creative projects you get involved in next, but I am sure they will involve hummus.

To Sasha, you've offered me so much support over the past three years. You never cease to make me guffaw and have taught me so much about philosophy and finding peace. You always make me feel better and I'm eternally grateful for your companionship.

Finally, and most importantly. I would like to thank my mum, Tina. You have picked up the phone every single time I have needed anything. Every time I've hit a low you've been there and you always listen and give me the best advice. I wouldn't be where I am today without your wisdom.

Abstract

The Next Band All-Sky-Survey (NextBASS) is an instrument capable of measuring the diffuse Galactic Foregrounds continuously in frequency between 7 to 30 GHz, in both temperature and polarisation. Still currently in the design phase, NextBASS will obtain a resolution of $\leq 1^\circ$, using a 6m cross-Dragone telescope with two bands, a single feed operating at 7 – 15 GHz and a 33-feed array of corrugated feed horns for 15 – 30 GHz. NextBASS will achieve a raw sensitivity of $\approx 9.4\mu\text{Kdeg}$ for band-1, and $\approx 1.7\mu\text{Kdeg}$ for in band-2. This sensitivity estimate is based on a conceptual 5 year survey, modelled with the atmospheric conditions and a visible sky fraction mirroring those of the Llano de Chajnantor Observatory. The receiver and digital back-end technologies will use continuous wave stabilisation with a low-bit and fast sampling ADC, and digital-signal-processing performed in an FPGA. Using single pixel forecasting simulations with Bayesian parametric fitting methods, I demonstrate that NextBASS is capable of measuring a synchrotron spectral curvature term. Furthermore, NextBASS has the potential to detect the presence of a 1% polarised AME component, in regions of sky where the AME amplitude is bright comparative to other foregrounds. NextBASS will improve constraints on the diffuse foreground parameters and help remove bias on the cosmic microwave background B -mode amplitude. This thesis focuses on the motivations behind NextBASS by assessing its impact through single pixel forecasting, along with a full proof of concept, detailing the antenna and feed horn designs and developing a receiver model. The final performance of the combined optical and receiver system is capable of achieving beam efficiencies of over 96% in a 1° beam, yielding $T_{\text{SYS}} \leq 34$ K for all frequencies, and a cross-polarisation of less than -35 dB across the band.

Statement of Originality

This is to certify that this thesis was written by myself, and that all intellectual content within this thesis is my own, unless explicitly stated otherwise.

The work in this thesis has not been submitted for a degree to any other university, institution or professional body.

This work is the intellectual property of the author.

Copyright ©2019 Jaz Hill-Valler

Contents

1	The Cosmic Microwave Background and its Diffuse Galactic Foregrounds	1
1.1	Cosmology	3
1.1.1	A brief Cosmological History	4
1.1.2	The Cold Dark Matter Model	4
1.1.3	Cosmic Inflation	5
1.1.4	The Cosmic Microwave Background	7
1.2	CMB Foregrounds	13
1.2.1	Synchrotron	14
1.2.2	Free-free	16
1.2.3	Thermal Dust	16
1.2.4	Anomalous Microwave Emission	17
1.3	CMB Experiments	20
1.3.1	Pre-2020	20
1.3.2	Post-2020	22
1.4	X-BASS and NextBASS	26
1.4.1	Motivation	26
1.4.2	Objectives	27
1.4.3	Planning	28
1.4.4	Instrument Requirements	30
1.5	Thesis Overview	32
2	The impact of X-BASS and NextBASS on CMB Foreground Component Separation	35
2.1	Component Separation	37
2.1.1	Template Fitting	38
2.1.2	Blind and Semi-blind Component Separation	39
2.1.3	Bayesian Parametric Fitting	42

2.1.4	The importance of component separation	42
2.2	NextBASS Simulations	43
2.2.1	Spectral Models	44
2.2.2	Parameter values and priors	48
2.2.3	Simulated Observations	49
2.3	Intensity Simulation Results	56
2.3.1	Straight Synchrotron	56
2.3.2	Curved Synchrotron	57
2.4	Polarisation	59
2.4.1	Basic Polarisation Foregrounds	60
2.4.2	Polarisation Foreground with Synchrotron Curvature	62
2.4.3	Curved Synchrotron Modelling Error	64
2.4.4	1% Polarised AME Component	68
2.4.5	Excluding 1% AME from simulation	72
2.5	Future Work	74
3	The Antenna Design for NextBASS	77
3.1	Dual Reflector Antenna Design Principles	78
3.1.1	Telescope Design Fundamentals	78
3.1.2	On-axis Systems	80
3.1.3	Off-axis Systems	81
3.1.4	Primary Mirror	82
3.1.5	Secondary Mirror	83
3.1.6	Mizuguchi Condition	83
3.1.7	Effective Focal Length	84
3.2	Optical Analysis	84
3.2.1	PO analysis	85
3.2.2	Modelling Spillover	86
3.2.3	Modelling Blockages	88
3.2.4	Overlap Integral	89
3.2.5	Co-polar Beam Pattern	90
3.2.6	Cross-polarisation	90
3.2.7	Calculating Total Power	90
3.2.8	Directivity	92
3.2.9	Peak-to-peak Cross-polarisation	92
3.2.10	Sidelobe level	92

3.2.11	Beam Width	93
3.2.12	Beam Efficiency	93
3.2.13	Aperture Efficiency	94
3.2.14	Beam Ellipticity	94
3.3	X-BASS Optics	95
3.3.1	C-BASS South	96
3.3.2	Photogrammetry	99
3.3.3	X-BASS Performance	101
3.4	NextBASS Optics	106
3.4.1	Performance Targets	106
3.4.2	Detailed Geometry	107
3.4.3	Beam Properties	109
3.4.4	Focal Plane	115
3.4.5	Optical Shielding	120
3.5	Optical Design Review	125
4	The X-BASS and NextBASS Feed Horns	127
4.1	Feed Horn Principle	128
4.1.1	Mode Generation	128
4.1.2	Types of Feed Horn	130
4.2	Broad-band Corrugated Horn Design	132
4.2.1	Waveguide and Mode-Converter	132
4.2.2	Profiling	133
4.2.3	Phase Centre	134
4.2.4	CHAMP Analysis	134
4.3	Far Field Horn Design	136
4.3.1	Gaussian Optics	136
4.3.2	Edge Taper	138
4.3.3	Beam Matching	140
4.3.4	Gaussian beam approximation	140
4.4	An Optimised Horn for X-BASS	141
4.4.1	Recreating C-BASS horn design	141
4.4.2	Optimisation	142
4.4.3	CHAMP results	145
4.4.4	Analysis in HFSS	147
4.4.5	Phase Centre	150

4.4.6	Performance with Optics	150
4.4.7	Edge Tapers	150
4.5	The NextBASS Feed Horns	153
4.5.1	Design Curves	153
4.5.2	Profiles	154
4.5.3	Band-1 corrugated horn design	156
4.5.4	Band-2 corrugated horn design	160
4.6	NextBASS Edge Taper	163
4.7	Feed Array	163
4.7.1	Circular Packing	164
4.7.2	Feed Positioning	164
4.7.3	Manufacturing Methods	165
4.8	Future Feed Development	166
5	A Conceptual 2:1 Bandwidth Polarisation Receiver	167
5.1	Radio Receiver	168
5.1.1	Receiver Front End	168
5.1.2	Analogue Components	169
5.1.3	Digital Backend	171
5.2	Sources of Noise	173
5.2.1	Thermal Noise	174
5.2.2	Pink Noise	174
5.2.3	Quantisation Noise	175
5.3	Methods of Gain Stabilisation	176
5.3.1	Total Power Radiometer	176
5.3.2	Dicke Switch Radiometer	177
5.3.3	Continuous Comparison Radiometer	177
5.3.4	Noise Injection Radiometer	178
5.3.5	Continuous Wave	179
5.4	Polarisation	180
5.4.1	Stokes Parameters	180
5.5	Receiver Architectures	182
5.5.1	Friis Equation	182
5.5.2	Components	184
5.5.3	X-BASS Architecture	188
5.5.4	NextBASS Architecture	188

5.5.5	Cryostat	190
5.6	Filter Design	191
5.6.1	Microwave Filters	192
5.6.2	Design Process	194
5.6.3	Anti-Aliasing Filter	195
5.6.4	7 – 15 GHz Bandpass Filter	197
5.6.5	15 – 30 GHz Bandpass Filter	199
5.7	Further Receiver Development and testing	199
6	An Accurate Estimate of Instrument Sensitivity	201
6.1	Atmospheric Modelling	202
6.1.1	Atmospheric Radio Spectrum	203
6.1.2	Water Vapour	205
6.1.3	Estimating Sky Temperature	208
6.1.4	Location	210
6.2	Receiver System Temperature	216
6.2.1	Receiver chain	216
6.2.2	Cumulative Receiver Power	219
6.2.3	System Temperature	220
6.3	Realistic System Modelling	222
6.3.1	Scanning Strategy	223
6.3.2	Observation Elevation	223
6.3.3	Accurate Brightness Temperature	224
6.3.4	Temperature Integral	224
6.3.5	Antenna Temperature: X-BASS and NextBASS	226
6.4	Calculating Survey Sensitivity	230
6.4.1	Survey Parameters	230
6.4.2	Beam width estimation	231
6.4.3	Calculating Sensitivity	231
6.4.4	Number of Feeds	233
6.5	X-BASS and NextBASS Sensitivities	234
6.6	Summary	237
7	A Summary: The Next Generation of Galactic Foreground Instru-	
	ments	239
A	Markhov Chain Monte Carlo	245

B Sky Model	247
C Ground Model	249
Bibliography	250

Chapter 1

The Cosmic Microwave Background and its Diffuse Galactic Foregrounds

“I think, therefore I am... confused”

Benjamin Hoff, The Tao of Pooh

The observations of the Cosmic Microwave Background (CMB) radiation have provided irrefutable evidence for the Big Bang theory describing our Universe’s very beginnings. The isotropic and homogeneous nature of the CMB can be explained by the rapid period of expansion, known as inflation, and the small-scale fluctuations provide a wealth of knowledge on the interactions between primordial photons and matter. From this science, we have constructed the most widely accepted model that describes the behaviour of photons, baryonic matter, dark matter and dark energy in our Universe. This model is known as the Lambda Cold-Dark-Matter model (Λ CDM). Furthermore, the CMB can also be studied in polarisation. The polarisation is generated from the Thomson scattering of photons by free electrons, and is observed as E -mode fluctuations. The gravitational lensing of these fluctuations at later times and on smaller angular scales causes B -modes. Lensed B -modes have been detected and studied by [Hanson et al. \(2013\)](#). However, there is also a purely B -mode signal, whose origins come from gravitational waves produced at inflation by the exponential expansion of space and the interactions between matter and radiation. By detecting these primordial B -modes, we could rule out invalid inflationary theories, we could also explain exactly what happened when the Universe began to reheat itself after inflation, and understand the turbulent mixing of matter and radiation. The CMB

B -mode power spectrum offers a way of probing times before the CMB photons were released at recombination (when neutral atoms first formed).

However, these B -mode signals are extremely weak and difficult to detect. As a result, we look for ‘non-detections’, or when the residual primordial B -mode power spectrum is non-zero. Attempts to measure the B -mode power spectrum have been done by a number of experiments, but only upper limits have been obtained, and these do not quite fit what the theory predicts we should observe. Observationally, there are two major limitations. First of all, the current instruments do not have sensitivities low enough to match the theoretical predictions. This is being overcome by a wave of post-2020 space experiments that will map the sky at sensitivities never before achieved. The second limitation, which concerns the work developed in this thesis, is the presence of diffuse Galactic Foregrounds. These foregrounds dominate the sky at all frequencies and consist of emissions from a variety of sources. The two most dominant are synchrotron radiation at low frequencies, and thermal dust emission at high frequencies. By mapping the sky at these two ‘lever-points’, we can create extremely accurate foreground models, which can be used to clean the polarised sky maps obtained by the new wave of satellite missions.

Observations of thermal dust cannot be done from the ground because the atmosphere is opaque at these frequencies. However, at lower frequencies, the radio window means all frequencies where synchrotron dominates the sky, can be observed from the ground. Most of the space experiments will be able to make measurements of the dominant thermal dust as they cover higher frequencies. Currently, there have only been a few measurements at frequencies which the synchrotron dominates. The C-BASS data point ([Jones et al. \(2018\)](#)) at 5 GHz and S-PASS at 1.4 GHz ([Carretti et al. \(2019\)](#)), have made sufficient measurements of synchrotron, demonstrating the importance of foreground cleaning and the importance of synchrotron in regards to understanding the behaviour of electrons in the Galaxy, for example, Faraday depolarisation ([Carretti et al. \(2019\)](#)).

Yet the synchrotron radiation is still dominant up to frequencies of ≈ 40 GHz. Low-frequency channels are difficult to implement on satellite experiments because the receivers are large and take up the majority of the payload weight, meaning fewer detectors become available for the sensitive, high-frequency channels. The lack of low-frequency data on these foregrounds is a just cause to investigate new ways of detecting and modelling these dominant diffuse emissions, and hence provide cleaner maps of the CMB.

This thesis contains the preparatory work required for a robust conceptual design of a new Galactic Foregrounds instrument. Following the success of the *Planck* satellite, the CMB community must make efforts to increase the breadth of frequency coverage and improve sensitivity requirements in order to extract the CMB B -mode power spectrum. In this Chapter, I will describe the science of the CMB, its foregrounds and how the literature supports the claim that a new low-frequency instrumentation would be beneficial to both CMB cleaning and foreground science.

I begin with an introduction to modern cosmology followed by a summary of the CMB. I then describe the mechanisms behind the diffuse Galactic Foreground components and why they are interesting as a stand-alone science case. I then list the small and large scale past, present and post-2020 experiments that will be measuring cosmological parameters and looking for the primordial B -mode signal.

To finish the chapter, I outline the planning that has gone into the proposed NextBASS instrument consisting of a description for two potential experiments, the ‘X-Band All Sky-Survey’, X-BASS and the ‘Next Band All-Sky Survey’, NextBASS. The X-BASS Survey will cover between 7 – 15 GHz and is set out as a path-finder experiment that will be able to constrain the synchrotron spectral parameters in observations. NextBASS is an ideal instrument, covering between 7 – 30 GHz in continuous frequency coverage, with a sensitivity which matches the foreground demands for post-2020 CMB satellites. This instrument will be able to make a map of the polarised sky and isolate region-by-region appropriate foreground models, as well as being sensitive enough to confirm the presence of any currently unknown polarised foregrounds that may cause a shift in the B -mode power spectrum if not modelled correctly. The first two chapters of this thesis are concerned with the motivations as described in the literature and analysing the statistical performance of NextBASS. The following three chapters describe the conceptual design of the optics and receiver technology development to show the technology readiness of these experiments. The final chapter describes the modelling that has gone into estimating the sensitivity of NextBASS covering atmospheric modelling and survey strategies.

1.1 Cosmology

This section describes a summary of modern Cosmology. This includes the evolution of the universe and the theory of cosmic inflation. I will then describe how observing the CMB can produce a power spectrum, which provides us with observation evidence of these cosmological theories and the parameters we use to describe the Universe.

1.1.1 A brief Cosmological History

Our Universe is best described by the Standard Big Bang Model. In this model, the universe begins as a singularity followed by a period of inflation (D’Inverno (1992)). Periods of cooling and ionisation then shaped the formation and evolution of the universe up to the horizons we now see today. At about 10^{-36} second old, the universe underwent a period of accelerated expansion, known as cosmic inflation, this lasted until about $\approx 10^{-32}$ seconds (Yardly (2015)). The early universe is responsible for emergence of the fundamental forces that govern the Universe. After $\approx 50,000$ years old, the Universe cools and matter becomes more dominant than radiation. Around 400,000 years after the Big Bang, the Universe became cool enough for neutral matter to form. This epoch is known as the epoch of recombination, where photons released during the decoupling of matter and radiation are now detectable as the Cosmic Microwave Background. The Universe was left with large, dense clouds of hydrogen, slowly collapsing to form stars and galaxies. Since neutral hydrogen only emits in radio waves, these have since been red-shifted to unobservable wavelengths. This period was known as the ‘Dark Ages’. Eventually the first stars and galaxies formed, briefly re-ionization space at around 400 million years old. After 1 billion years, the Universe as we know it today, was formed (Burke and Graham-Smith (2010)).

1.1.2 The Cold Dark Matter Model

The Λ CDM model is the most established description of Cosmology. It consists of six free parameters which can accurately describe cosmological observations such as the relative abundances of primordial elements, and the clustering of galaxies on large scales. It also can accurately fit the CMB power spectrum. These parameters are: $\Omega_b h^2$, the density of baryonic matter; $\Omega_c h^2$, the density of cold dark matter; $100\theta_{MC}$, the approximate angular scale of the first acoustic peak of the CMB power spectrum; τ , the optical depth of reionization; n_s , the spectral index of the primordial scalar perturbations and the log-value of the power of primordial curvature perturbations at 0.05Mpc^{-1} . Table 1.1 has the Λ CDM parameters best fitted by Planck Collaboration et al. (2018).

Parameter	Value
$\Omega_b h^2$	0.02242 ± 0.00014
$\Omega_c h^2$	0.11933 ± 0.00091
$100\theta_{MC}$	1.04101 ± 0.00029
τ	0.0561 ± 0.0071
$\ln(10^{10} A_s)$	3.047 ± 0.0014
n_s	0.9665 ± 0.0038

Table 1.1: The Λ CDM parameters and constraints from the likelihood analysis carried out by [Planck Collaboration et al. \(2018\)](#) using TT, TE, EE, low E, lensing and BAO measurements.

1.1.3 Cosmic Inflation

Inflation was an exponential, accelerated expansion of space and was driven by the means of an inflationary field. Today the universe is still expanding but at a much slower rate. Alexander Friedmann ([Friedmann \(1922\)](#)), developed Einstein's equations of general relativity to show the universe is still expanding, in contrast to Einstein's static universe ([Einstein \(1915\)](#)). The Friedmann-Robertson-Walker metric best describes the homogeneous, isotropic and expanding universe and is given by [D'Inverno \(1992\)](#):

$$ds^2 = a(t)^2 ds_3^2 - c^2 dt^2, \quad (1.1)$$

where $a(t)$ is the scale factor and ds_3^2 is a 3-dimensional metric that describes the Universe's geometry. We can form the Friedmann equations from the analytical solutions to Einstein's field equations. These are expressed as:

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3} \left(\rho + \frac{3P}{c^2} \right) + \frac{\Lambda c^3}{3}. \quad (1.2)$$

$$\left(\frac{\dot{a}}{a} \right)^2 = \frac{8\pi G}{3} \rho - \frac{kc^2}{a^2} + \frac{\Lambda c^2}{3} \quad (1.3)$$

where ρ is the energy density, P is the pressure, k is the curvature, G is the gravitational constant and Λ is the cosmological constant.

From these equations, the Hubble parameter, which relates an object's distance to its recession velocity, can be defined as $H_0 = \frac{\dot{a}}{a}$. Measurements of the Hubble constant can be made from astrophysical measurements of Doppler shifted Galaxies using Type 1a Supernovae. This has been done by experiments such as the Hubble Space Telescope ([Riess et al. \(2009\)](#)). The Hubble constant can also be estimated by

fitting cosmological models to the CMB, most up-to-date value of H_0 from these types of measurements has been made by the *Planck* satellite, giving $H_0 = 67.66 \pm 0.42 \text{ kms}^{-1}\text{Mpc}^{-1}$ (Planck Collaboration et al. (2018)).

To usefully describe the Universe that we see, we need to parametrise it in such a way that each parameter corresponds to an observable. The metric, ds_3^2 , describes three types of curvature depending on the total energy density. If space is flat, k takes the value of 0; if space is a sphere with a positive curvature, $k = 1$; and if space is hyperbolic with a negative curvature, $k = -1$. We can then define a critical density which would be just sufficient to make the universe flat as $\rho_c = \frac{3H^2}{8\pi G}$. The total density parameter, Ω_{tot} , which describes the energy density of the Universe, can be written as a fraction of initial density of the universe, ρ_t over the critical density, ρ_c such that $\Omega_{tot} = \rho(t)/\rho_c$. Since Ω_{tot} describes the total energy from all sources, it can be written as the summation (D’Inverno (1992)):

$$\Omega_{tot} = \Omega_m + \Omega_R + \Omega_k + \Omega_\Lambda, \quad (1.4)$$

where Ω_m = all dark and baryonic matter contributions, Ω_R = radiation density, Ω_k = spatial curvature density and Ω_Λ is the dark energy density, causing the accelerated expansion.

Differences in the values for these densities describe different universes of very different evolution. To understand the value of these parameters, we must consider the theory of inflation. The universe as we observe it today, is geometrically near flat, homogeneous and isotropic across the sky. These properties cannot be explained if it was not for the accelerated expansion of the universe, known as the period of inflation.

Early CMB experiments such as BOOMERanG (de Bernardis et al. (2000)) and the Wilkinson Microwave Anisotropy Probe (*WMAP*), (Bennett et al. (2013)), determined that space follows a nearly Euclidean flat geometry. If space truly is flat, then the universe must have reached a finely tuned critical density, that is extremely difficult to maintain because the initial density of the universe, ρ_t , could have been any value. If the curvature of the universe, differs even slightly from $k = 1$, the total energy density, Ω_{tot} , would either be so large that the universe would curve back in on itself, resulting in a ‘Big Crunch’, or Ω_{tot} would be so low that the universe would have expanded so quickly that space would be sparse, and there would not be strong enough gravitational interactions to form the galaxies we observe today.

Inflation is driven by an inflationary field that has a constant density, which stays constant as space expands. As a result, inflation offers a solution to the ‘flatness’

problem by forcing Ω_{tot} to stay at unity, maintaining a density close to ρ_c (Bennett et al. (2013)).

Inflation also explains the homogeneous and isotropic nature of the CMB. For any two regions of sky, no matter their distance apart, they will be in thermal equilibrium. The laws of thermodynamics mean this is only possible if they are in thermal contact. As a result, the isotropic and homogeneous nature of space can only be explained as all of space being in thermal contact with itself at some point in history, and then accelerating apart at such a rate, that all regions of space starting at the same temperature have now cooled to the 2.71 K that we see today (D’Inverno (1992)).

Another problem that arises in modern cosmological models, is the issue regarding the absence of magnetic monopoles. The existence of magnetic monopoles is predicted in many theories of particle physics and cosmology, however there has been no such detection (Rajantie (2012)). Magnetic monopoles were likely to have formed during the extreme temperatures and pressures of the early universe. During inflation, the accelerated expansion would have driven them so far apart, giving them a negligible number density, making a detection statistically unlikely.

1.1.4 The Cosmic Microwave Background

The Cosmic Microwave Background (CMB) is the imprint of residual photons from the Big Bang, it is isotropic and homogeneous. It was first predicted by Ralph Alpher and Robert Herman in 1948 (Alpher and Herman (1948)), however this was not broadly discussed in physics at the time. Robert Dicke made a similar prediction in the early 1960’s and set out to build an instrument to make the first measurement of the CMB. However, in 1964 Arno Penzias and Robert Wilson at Bell Telephone Laboratories were trying to show that the absolute temperature measurement of the sky was 0K, however what they actually measured was around 4.2 K (Penzias and Wilson (1965)). As a result, Penzias and Wilson won the Nobel prize for the discovery. The most accurate measurement of the CMB today is at a temperature of 2.72548 ± 0.00057 K (Fixsen (2009)).

The early Universe was a hot dense plasma, which as the universe expanded, cooled to a temperature where matter and photons could decouple and neutral matter could form. This epoch of recombination is where the observable universe begins. According to the theory of inflation, any interactions that happened before the inflationary period should be smoothed out. However, anisotropies are observed on varying angular scales. These are due to quantum fluctuations in the inflationary

field, which were inflated up to macroscopic size, are visible as small temperature fluctuations in the black-body spectrum of the CMB.

To investigate the anisotropies on differing angular scales, we need to express the CMB temperature fluctuations as a decomposition of its spatial properties in spherical harmonics. The net temperature fluctuations can be expressed as $T(\hat{n})/T = \frac{T(\hat{n})-T_0}{T_0}$, where \hat{n} is the direction of the sky such that $\hat{n} = (\theta, \phi)$ and T_0 is the CMB temperature without any fluctuations. To express this in terms of angular scale, we can use spherical harmonics with coefficients $a_{\ell m}$, where, $Y_{\ell m}$ is a series of Legendre Polynomials $P_{m\ell}(\cos\theta)$ (Burke and Graham-Smith (2010)), giving the temperature fluctuations as:

$$T(\hat{n})/T = \sum_{\ell=0}^{\ell=\infty} \sum_{m=-\ell}^m a_{\ell m} Y_{\ell m}(\hat{n}), \quad (1.5)$$

where the coefficients are given by:

$$a_{\ell m} = \int_{\theta=-\pi}^{\pi} \int_{\phi=0}^{2\pi} (T_{\hat{n}}/T) Y_{\ell m}^*(\hat{n}) d\Omega, \quad (1.6)$$

and $Y_{\ell m}$ is the series of Legendre Polynomials, such that:

$$Y_{\ell m} = \sqrt{\frac{(2\ell+1)(\ell+m)!}{4\pi(\ell-m)!}} P_{\ell}^m(\cos\theta) e^{im\phi}. \quad (1.7)$$

We define the power spectrum coefficients, C_{ℓ} , to be the variance of the spherical harmonic coefficients such that $\langle a_{\ell m} a_{\ell m}^* \rangle = \delta_{\ell\ell'} \delta_{mm'} C_{\ell}$. Following thus:

$$C_{\ell} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} a_{\ell m} a_{\ell m}^*. \quad (1.8)$$

There are $(2\ell+1)$ multipoles for each value of ℓ (Burke and Graham-Smith (2010)). We can see from equations 1.6 and 1.8 that the power spectrum, C_{ℓ} can be expressed in terms of the temperature fluctuations squared, $(T(\hat{n})/T)^2$. Therefore, we can plot the CMB temperature power spectrum as the temperature fluctuations squared, vs. multipole. Figure 1.1 shows the angular power spectrum $\frac{\ell(\ell+1)}{2\pi} C_{\ell}$ vs. ℓ for the CMB temperature anisotropies, measured by a combination of experiments. The angular scale can be estimated from $\theta = \frac{180}{\ell}$, such that observing the CMB dipole at $\ell = 1$ corresponds to the 180° , or half the sky.




Figure 1.1: The CMB temperature angular power spectrum showing which multipoles correspond to which angular scales. At the lowest multipoles the errors are dominated by cosmic variance, i.e the fact that we observe only one realisation of the universe. Taken from <http://sci.esa.int/jump.cfm?oid=51555> based on Planck 2013 CMB results (Planck Collaboration et al. (2015)).

The angular power spectrum shows a number of acoustic peaks and a damping tail, known as the Silk damping tail. Each feature corresponds to a fundamental part in the Universe’s evolution. The first peak represents the scale at which the oscillation has just reached its first compression at recombination. The second peak is the first rarefaction. Subsequent odd and even peaks are multiple compressions/rarefactions respectively. The angular power spectrum supports adiabatic density perturbations, which are only possible with inflation, such that there is an equal amount of baryonic and photon perturbations (Yardly (2015)).

The first acoustic peak is at an angular scale $< 1^\circ$ and indicates the geometry of the Universe. For a closed universe, the first acoustic peak would be located at an angular scale larger than for a flat universe, conversely, for an open universe, this would be at a smaller angular scale. All peaks are due to harmonic oscillations between Baryon mass and radiation pressure. The potential wells depend on Baryon masses at points of maximum compression. These wells are larger for larger masses, this affects primarily the height of the first, and odd number peaks. The rarefaction is not mass dependent and changes the height of the second, and other even peaks. The ratio between the first and second peak amplitudes therefore indicates the ratio of Baryon density to the critical density.

Other features arising from the number of Baryons can affect the frequency of oscillations. This means there are degeneracies between cosmological parameters such that oscillations have higher frequencies at higher multipoles meaning a change in height of the first peak can be due to the number of Baryons, not just the mass.

The amplitudes of the higher acoustic peaks determine the total density of non-relativistic matter in the Universe. This is a probe into the transition from radiation to matter domination. Finally, the Silk damping tail is affected once more by the density of Baryons, in particular, damping is an indicator of the behaviour at the surface of last scattering. If the distance travelled by these scattered photons becomes longer than the wavelength, the damping oscillations become more severe. This damping tail can give insight into the age of the Universe (Yardly (2015)).

It is not just the total intensity amplitude of the CMB fluctuations that tell us about inflation. These interactions also left polarisation imprints on the CMB photons. Thomson scattering of photons due to free electrons has caused the CMB to be linearly polarised to a 10% level. This polarisation level was first detected by the Degree Angular Scale Interferometer (DASI) in 2001 (Kovac et al. (2002)).

There are two types of polarisation fluctuations: scalar, which are due to density/temperature fluctuations and tensor patterns, which can only be a result of gravitational waves. The scalar fluctuations leave an E -mode polarisation imprint on the CMB, and the tensor patterns cause a B -mode imprint. The E and B modes are of even and odd parity respectively. The scattering of photons creates the dominant E -mode.

Gravitational waves then scramble E -modes and produce some B -mode signals along with a tiny amount of E -mode. However this E -mode signal is lost beneath the larger E -mode signals from Thomson scattering. It is predicted that during inflation, gravitational waves were produced, these are too faint to detect directly but they may have left a tensor B -mode polarisation pattern in the CMB photons on large angular scales. Any B -modes at smaller scales can be introduced by weak gravitational lensing of E -modes. These start to dominate at multipoles of $\ell \approx 150$. Figure 1.2 shows the polarisation imprint of both the E and B modes. These are such that they represent the polarisation vectors on the sky, depending on the angular scale of the observation and its direction. The receiver measures orthogonal polarisation Stokes parameters, Q and U ; these are then transformed into spin-weighted spherical harmonics. The correlations between the Q and U signals hence describe the B and E -modes (Zaldarriaga (1998)).




Figure 1.2: The E and B polarisation patterns specify the linear polarisation field at any given point on the sky. The angular scale of this pattern will depend on the cause of the E or B mode. The E mode is a curl-free mode such that the polarisation vectors are radial around the cold spots and tangential around the hot spots. The B -mode polarisation is divergence-free and has a curl. The choice of E and B shares similarities with electrostatics such that $\nabla \times \mathbf{E} = 0$ and $\nabla \cdot \mathbf{B} = 0$ (Baumann et al. (2008)).

The B -mode signal is so faint, that foreground measurements are very important for trying to measure it. In 2014, BICEP2 Collaboration (2014) reported they had detected the B -mode signal. However, when including *Planck* data, (BICEP2/Keck, Planck Collaboration et al (2015)) showed that what they had detected were tensor-mode signals due to polarised thermal dust emission. This was an important step in the path to understanding foreground emissions and how they impact measurements of the CMB polarisation.

Since a direct detection of primordial gravitational waves is extremely difficult, if not impossible, we can instead measure the tensor-to-scalar ratio. The tensor-to-scalar ratio, r , is the ratio of tensor fluctuation amplitudes to scalar amplitude. This is described by the ratio of the amplitudes of the relative power spectra at $\ell = 4$, since this multipole is at angular scales that describe the behaviour of the recombination bump (Abazajian et al. (2016)). An estimate of r is hence:

$$r \approx \frac{P_t(k_*)}{P_s(k_*)}. \quad (1.9)$$

where $P_t(k_0)$ is the tensor power spectrum, and $P_s(k_0)$ is the scalar power spectrum at a pivot scale k_* .

The Gaussian theory that describes the fluctuations seen in the CMB requires that four power spectra are needed to measure the fluctuations. These are the autocorrelation between T , E and B and the cross-correlation of E and T . Any cross-correlations between B and the other parameters cancel out (Zaldarriaga (1998)). The E and B -mode power spectra, in terms of the auto- and cross-correlations are given in Figure 1.3. The value of r has an upper limit and since the E -modes have been well measured, the tensor power spectrum could have any amplitude below the upper limit of $r=0.1$. At higher multipoles, the B -mode spectrum starts to mirror the peaks of the E -mode spectrum. This is due to the E -modes crossing over into the B power as gravitational waves go through recombination. The bump at lower ℓ is the reionisation bump. The amplitude of the scalar fluctuations is given by A_s in Table 1.1.




Figure 1.3: The TT , EE and BB CMB Power Spectra, taken from Errard et al. (2016). The TT , EE and BB correspond to the auto-correlations of the temperature, E polarisation and B -mode polarisation power spectra respectively. The amplitude of the B -mode power spectrum can be scaled depending by the cosmological parameter that describes the scalar fluctuation amplitude. The tensor-to-scalar ratio is given by the ratio of the pure tensor to pure scalar power spectrum. This is proportional to their amplitudes once scaled to a pivot scale, defined by The Planck Collaboration et al. (2016) as $k_* = 0.05 Mpc^{-1}$. There is a leak at higher multipoles from the E -mode signal, into the B -mode signal as the gravitational lensing from galaxies causes the E -modes to have some B -mode signal.

Planck has made the most precise measurements of the tensor-to-scalar ratio to date. Upper bounds on r from Planck Collaboration et al. (2018), in conjunction

with data from Baryonic Acoustic Oscillations (BAO), BICEP2, and *Keck* data are currently $r_{0.002} < 0.07$ to a 95% confidence level. To make a detection consistent with slow roll inflation models, r must be detected to $r \leq 10^{-2}$ to a high confidence level (Planck Collaboration et al. (2018)).

1.2 CMB Foregrounds

From our atmosphere to the observable horizon, everything in between can be considered a foreground. On small scales, galaxy clusters contribute to foregrounds, but on the larger scales, diffuse galaxy emission is the main contributor. The four main components of galactic foregrounds are synchrotron radiation, free-free emission, the anomalous microwave emission (AME), peaking around 23 GHz and finally thermal dust, which dominates at frequencies above 60 GHz. I will be discussing each of these foregrounds in more detail in turn (Kiyotomo (2014)). Figure 1.4 shows the foreground frequency spectra in both intensity and polarisation. The vertical bars indicate the current frequency space coverage by various experiments. This indicates that there is a gap between the C-BASS frequency, at 5 GHz and the *Planck* frequencies.

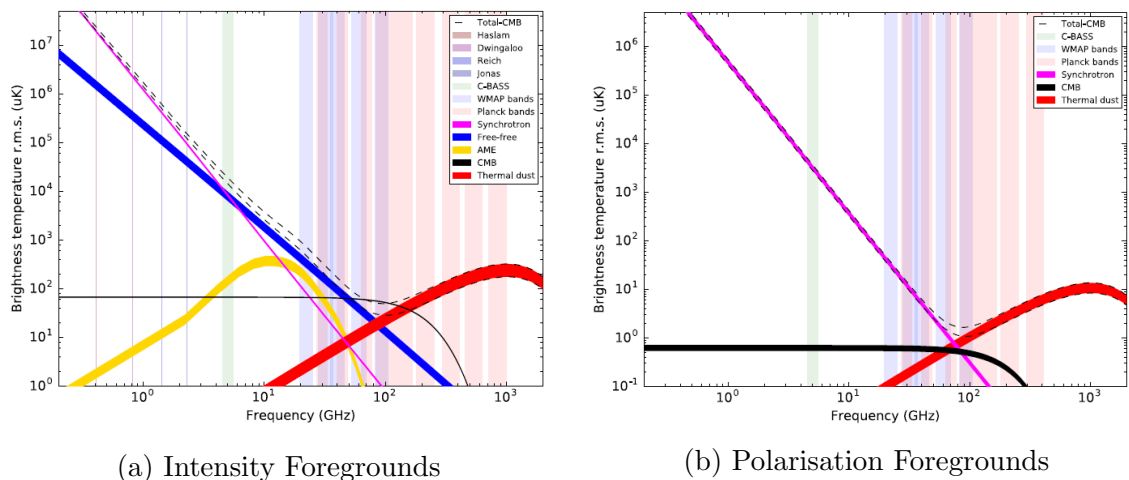


Figure 1.4: The frequency spectra of the CMB foregrounds in temperature and polarisation at a 1° full-width-half-maximum resolution. The E -mode polarisation amplitude has been calculated by the Planck Collaboration et al. (2016a) best fit power spectrum. The lower bounds for each component come from the rms fluctuations for the Planck Collaboration et al. (2016a) model for an 80% mask of the sky, the upper limit comes from the mask that includes 90% of the sky. (Taken from Jones et al. (2018), note the synchrotron power spectrum is represented by a single spectral index of $\beta = -3.1$).

1.2.1 Synchrotron

Synchrotron radiation arises from electrons moving at relativistic velocities in magnetic fields. These are the electron component of cosmic rays in the galaxy. In radio imaging synchrotron is visible as spurs or filaments around the galactic plane. It was first detected as a strong linear polarisation at optical wavelengths in the Crab Nebula. An electron undergoes acceleration perpendicular to the magnetic field and the resulting emission is linearly polarised. The percentage of polarisation depends on the level of tangling of the magnetic field along the line of sight which can be determined by observing the Q and U Stokes parameters (Burke and Graham-Smith (2010)). The spectral index is steeper than for free-free emission making synchrotron distinguishable in both intensity and polarisation. The emission model of the synchrotron spectrum relies on a fixed spectral index that scales the amplitude with frequency. The spectrum followed the same shape in both temperature and polarisation. Synchrotron has been observed between 0.1 GHz and 30 GHz and can be well described by a spectral index of $\beta_S = -3.1 \pm 0.05$ (Kogut (2012)).

Fuskeland et al. (2014) found that the synchrotron spectral index varies over Galactic latitude with a steeper spectral index toward the Galactic centre and anti-centre than towards the arms. They found this steepening of the spectral index to change from $\beta_s = 2.98$ to $\beta_s = 3.12$ from the Galactic plane to a high Galactic latitude. This measurement was carried by using maximum likelihood analysis methods with the 9 year *WMAP* sky maps in the K (22.45 GHz) and Ka (32.64 GHz) bands. Figure 1.5 shows the spatial variations of the synchrotron spectral index as generated by Jones et al. (2018) with both the C-BASS North Temperature data and data from the *Planck* frequencies.




Figure 1.5: Spectral index map obtained by the C-BASS collaboration, combined with *Planck* data, showing how the synchrotron spectral index varies spatially across the sky. This image is from the most recent C-BASS North data, taken from [Jones et al. \(2018\)](#).

Synchrotron radiation acts as a probe into low-energy cosmic rays. At low frequencies, the frequency spectrum of synchrotron emission reaches the mean kinetic temperature of the emitting electrons, causing self-absorption. At high temperatures, electrons can re-absorb the synchrotron radiation released by other electrons, causing the frequency spectrum to flatten as absorption and emission mechanisms cancel along the line of sight. Below the frequency at which this happens, the synchrotron behaves as a power law with a fixed spectral index. At higher frequencies, the electrons emitted at the greatest energies radiate fastest and hence the population of emitting electrons is much lower. Due to energy losses from self-absorption, the synchrotron spectral index is steeper at higher energies.

As a result, the model that best describes the behaviour of the synchrotron radiation can be changed to a power law with a varying spectral index. This is done by adding a curvature term, C_s , to the power law which accounts for the emissions from higher energy cosmic rays. To investigate the steepening of the spectral index as a function of frequency, [Kogut \(2012\)](#) combined data from ARCADE-2 with radio surveys between 22 MHz to 23 GHz (*WMAP*). Synchrotron radiation was found to have a mean spectral index of $\beta = -2.64 \pm 0.03$ and a curvature of $C = -0.052 \pm 0.005$ at a reference frequency of 0.31 GHz. Observations at higher frequencies would result in a steeper spectral index and a positive curvature.

1.2.2 Free-free

Free-free radiation is equally widespread in the Galaxy with a maximum contribution along the galactic plane. Free-free is also known as thermal-bremsstrahlung and arises from collisions between thermal electrons and protons plus other ions in the warm ionised ISM. Polarisation can be considered negligible due to the random and isotropic nature of Thomson scattering in optically thick regions. [N.Macellari et al. \(2011\)](#) showed that free-free emission is un-polarised with an upper limit on the free-free polarisation fraction of 3.4% to 2σ .

Free-free radiation is well understood as it can be traced with the $H\alpha$ line emission, which is dominant across the Galaxy. [Marta et al. \(2016\)](#) used a $H\alpha$ emission template to evaluate the spectral index of free-free radiation across the sky. Using cross-correlation analysis, they studied 15 regions of sky and estimated the emissivity of free-free as a ratio of template brightness. They had to correct for the foreground dust absorption and confirmed that there was zero correction needed for high galactic latitudes. They found that the spectral index of free-free radiation is $\beta_{ff} = -2.14$ for an electron temperature of $T_e = 4000K$.

1.2.3 Thermal Dust

At frequencies greater than 70 GHz, thermal emission from interstellar dust grains dominates foreground emissions in both polarisation and intensity. These dust grains mostly comprise of graphites, silicates and PAHs (Polycyclic Aromatic Hydrocarbons). These are typically heated to between $T_d = 10 - 30K$ ([Burke and Graham-Smith \(2010\)](#)). These dust grains are not spherically symmetric and become aligned by the galactic magnetic field which results in a net linear polarisation of the starlight. *Planck* produced the first all-sky map of the polarised emission from dust at 353 GHz finding a thermal dust to be 20% polarised ([Planck Collaboration et al. \(2015\)](#)). These polarisation maps offer a way of studying the Galactic magnetic field. The maps show filament features in I that have Q and U counterparts. Galactic interstellar dust is made of grains of multiple sizes and different materials hence changing polarisation properties. *Planck* probed the large grains that emit in the sub-mm which extinguish and polarise light in the visible spectrum ([Planck Collaboration et al. \(2015\)](#)).

Interstellar dust is created in stellar winds and the circumstellar ejecta of dying stars where the heavy elements in the gas condense to solid grains - these grow through accretion and coagulation and can be destroyed by radiation, grain-grain collisions and gas-grain collisions. As a result, studying the thermal dust emissions across the

Galaxy can demonstrate the differing behaviours of dust around stars at different stages of stellar evolution.

1.2.4 Anomalous Microwave Emission

The Anomalous Microwave Emission (AME), also known as spinning dust emission, is visible in frequency space as a bump peaking at 23 GHz. The AME is expected to only have a 1% polarisation due to the alignment mechanisms of small grains creating an overall cancellation in net polarisation. The cause of AME emission is widely debated, the most favoured carriers are spinning dust grains and magnetic grain materials, however neither theory has been proven. Observations show that the carrier may change with interstellar media, further increasing the model complexity. Interstellar dust grains are asymmetric and naturally have permanent electric dipoles, once excited these grains have significant emission at GHz signals, where larger grains have larger moments of inertia and emit at lower frequencies. The presence of spinning dust may affect measurements of tensor-to-scalar ratio. Observations by *Planck* indicate a low AME polarisation. In particular, foreground forecasting by [Remazeilles et al. \(2016\)](#), shows that even when including future *B*-mode missions such as *LiteBIRD*, even a small AME polarisation could be problematic for component separation with future, ultra-sensitivity data, unless more polarisation data that samples the peak of the AME is obtained.

If spinning dust is the primary candidate, [Draine and Hensley \(2016\)](#) claim that the AME should have a negligible polarisation. This is due to quantum suppression, where the discrete energy levels of very small grains don't allow for dissipative processes from grain rotational kinetic energy to generate emission. This results in a suppression of alignment with the local magnetic field and hence a negligible polarisation. They claim that if a future AME observation finds the polarisation fraction to be greater than 0.01% for $\nu > 10$ GHz, this will suffice as evidence that the AME is not dominated by rotational emission from nano-particles and that starlight anisotropy may have generated a small degree of polarisation in the rotational emission.

If the AME is dominated by magnetic grain materials, the polarisation fraction is expected to be significantly higher. Magnetic materials, such as silicates and iron, are heavily depleted in the gas phase and are expected to be locked into grains ([Hoang and Lazarian \(2016\)](#)). It is known that amorphous silicates are a significant component of the ISM.

An investigation done by [Hensley \(2017\)](#), computed the charge distributions of different carriers of magnetic dust in different regions of ISM. They found that the

emissivity of Fe and Si nano-particles are qualitatively similar but that only pure iron grains match the emission profile that was seen in the galactic average AME at 30 GHz (Stevenson (2014)). This means that since only purely Fe grains match the galactic emission profile, the dust grains may be susceptible to polarisation effects from the Galaxy’s magnetic field.

Draine and Lazarian (1998) and Draine and Hensley (2013) stated that free-flying iron nano-particles could achieve a polarisation fraction of approximately 30%. Hoang and Lazarian (2016) calculated the alignment efficiency as a function of size, finding that grains larger than 1 nm are poorly aligned where as grains smaller than 1 nm can produce a polarised emission of between 10 – 30%.

Small grains could achieve alignment via paramagnetic relaxation however this may be suppressed at high rotation frequencies. Ultra-small grains spinning extremely rapidly may be able to align through resonance relaxation in which rotational splitting of energy levels becomes important, if aligned in this manner, spinning dust would have a polarisation of around a percent. A sufficient measurement of the polarisation of the AME between 10-30 GHz would help distinguish between spinning and magnetic dust mechanisms (Dickinson et al. (2018a)).




Figure 1.6: Taken from Dickinson et al. (2018a) showing the spinning dust emissivity as a function of frequency for different states of the interstellar medium. These were produce by an adapted SPDUST2 code, based on the work done by Draine and Lazarian (1998).

Polycyclic-Aromatic Hydrocarbons (PAHs) are potential spinning dust carriers and may be highly abundant due to their high resistance to intense environments in the ISM. Ali-Haimoud et al. (2015) analysed the spectra of the Perseus molecular

cloud at 25 GHz over a 3 GHz bandwidth. They did not make a detection, however, they set upper limits on the abundance of individual PAHs with the assumption that the AME from Perseus is entirely due to spinning PAHs. A detection of PAH spectral lines would be an indication of spinning PAHs as the abundant AME mechanism (Ali-Haimoud et al. (2015)).

These theoretical predictions of polarisation are often a maximum and many line-of-sight effects can cause de-polarisation. To get a better idea of the polarisation of the AME, a comparison between the observed AME polarisation upper limits and the observed thermal dust polarisation in the same region may give a better indication into the spinning dust polarisation fractions (Dickinson et al. (2018a)).

Currently, QUIJOTE has not reached the nominal sensitivity it claimed to need in order to observe synchrotron. However, the QUIJOTE team carried out an observation of the Perseus molecular cloud (G159.6-18.5) and detected no polarisation, they then obtained upper limits on the polarisation fraction of $\Pi < 6.3\%$ and $\Pi < 2.8\%$ at 12 GHz and 18 GHz respectively (Génova-Santos et al. (2015)).

There are four dust regions and two HII regions that have published upper limits on polarisation. These are the G159.6-18.5 Perseus cloud, the ρ Ophiuci molecular cloud, LDN1622 and the Pleiades reflection Nebula. The HII regions are LPH96 and the Helix Nebula (NGC7293) (Planck Collaboration et al. (2013b)). The first detection of AME in the Perseus region was presented with the *COSMOSOMAS* experiment and the *WMAP* (23 – 94 GHz) data by Watson et al. (2005). They found a clear emission at 11 and 16 GHz. They also reported a polarisation fraction of $\Pi = 3.4^{+1.5}_{-1.9}\%$ at the 95% confidence level. This polarisation emission favours electric dipole emission over magnetic dipole emission. Data from the Effelsberg telescope at 2.7 GHz proposed that G159.6-18.5 acts as a ‘Faraday-Screen’, which is a phenomenon where the background emission of light undergoes a rotation in polarisation angle due to the combination of the Galaxy’s magnetic field and ionized gas in the ISM (Reich et al. (2004)). The most recent upper limits on the AME polarisation are shown in Figure 1.6, taken from Figure 2. of Dickinson et al. (2018a).

Part of the science goal of the SKA is to study molecular clouds, photo-dissociation regions (PDRs), and circumstellar disks with the intention of understanding the AME alongside stellar evolution. Collapsing dust clouds have typical sizes of 1 arcmin or less which makes them well matched to SKA baselines. The SKA will also have to make below 20 GHz observations of the synchrotron and free-free emission in order to separate these from any AME signals. ALMA has the resolution to detect the rotational spectrum of PAHs. The SKA offers a single instrument that could cover

all the relevant frequencies as well as synchrotron and free-free emission (Dickinson et al. (2014)).

A major motivation for NextBASS is to distinguish between magnetic dust mechanisms and resonance relaxation mechanisms, through detecting a polarisation of either 10 % – 30 % or 1% respectively. This would benefit the understanding of the variation of AME across the sky, and also shed light on the mechanisms that produce its polarisation (if any).

1.3 CMB Experiments

1.3.1 Pre-2020

The first experiment to confirm the blackbody nature of the CMB was the *COBE* satellite in 1990 (Mather et al. (1990)). Following this, balloon and ground experiments such as BOOMERanG (de Bernardis et al. (1999)), CBI (Padin et al. (2001)) and DASI (Kovac et al. (2002)) made early measurements of the CMB intensity and E -polarisation power spectra. *WMAP* was the first instrument to measure the CMB power spectrum at a range of multipoles ($\ell = 7 - \ell = 1000$). After *WMAP*, the *Planck* satellite launched in 2009 has generated incredibly accurate measurements of the intensity and both E and B polarisation of the CMB hence generating the best estimates of cosmological parameters for the Λ CDM model to date.

At higher multipoles, $\ell > 500$, the lensing B-mode becomes dominant. The experiment POLARBEAR directly detected lensing B -modes from the CMB-polarisation signal and found their results to be consistent with Λ CDM estimates (The Polarbear Collaboration et al. (2010)). POLARBEAR2 is an update of POLARBEAR1, adding an additional frequency channel at 95 GHz and increasing the number of detectors from 1,274 to 7,588. POLARBEAR2 aims to constrain $r \approx 10^{-2}$ (Inoue et al. (2016)).

The first experiment targeted at making a measurement on the primordial B -mode signal was the BICEP1 ground based telescope. This has since been expanded to BICEP2 + Keck Array and has been expanded once more (adding detectors) to BICEP3 (Grayson et al. (2016)). In 2014, BICEP Collaboration et al. (2014) reported an abnormally large value of $r = 0.2_{-0.05}^{+0.07}$, however further analysis from *Planck* revealed that this high value arose from B-mode contributions from thermal dust (BICEP2/Keck, Planck Collaboration et al (2015)).

Other experiments that measured the CMB power spectra at smaller angular scales are ACTpol and SPTpol. ACTpol measured the intensity and polarisation

power spectra to characterise gravitational lensing, estimated a value for the neutrino mass and study the growth of structure from SZ observations of galaxy clusters (Thornton et al. (2016)). *SPT-3G* will achieve similar science goals with an angular resolution of 1arcmin (Aylor et al. (2017)). Balloon experiments that have accompanied these include EBEX launched in 2011 (The EBEX Collaboration et al. (2010)), and SPIDER which first launched in 2015 and is on-going (Gualtieri et al. (2017)). These experiments that observe $\ell > 500$ cover the higher frequency regime.

In order to separate foregrounds from the CMB, we require more data from foreground dominant regions of frequency space. This requires further low frequency observations to disentangle synchrotron, free-free and AME along with expansion into the sub-mm regime to constrain thermal dust emissions. The experiments that are currently observing the CMB at resolutions of $< 1^\circ$ have been summarised by Jones et al. (2018) and are shown in Figure 1.7. These larger scale foreground experiments, however, currently do not have sufficient sensitivity to synchrotron in order to remove it effectively.

It is not just sensitivity that is required for effective foreground cleaning, Remazeilles et al. (2016) have shown that without additional frequency channels, foregrounds cannot be effectively modelled. They also found that any errors in foreground modelling can introduce a bias in values of r . The QUIJOTE (Q-U-I-Joint Experiment) operates at 11, 13, 17 and 19 GHz contributing low frequency data for foreground removal on smaller regions of sky. QUIJOTE intends to monitor the polarisation of synchrotron and AME emissions by means of two large surveys in polarisation. The 11-19 GHz survey achieves a raw sensitivity of $15 \mu\text{Kdeg}$ (Poidevin et al. (2018)). QUIJOTE has attempted measurements limits on the AME polarisation and the additional frequency channels has been shown to improve errors on r (Génova-Santos et al. (2015), Errard et al. (2016)).

The S-band Polarisation All Sky Survey (S-PASS) has recently released a paper on their survey of polarised emissions over the southern sky at 2.3 GHz. S-PASS measures a region of the sky where synchrotron is dominant but depolarisation from Faraday rotation is also a dominating foreground effect. S-PASS can further constrain synchrotron models and investigate the spectral variations across the sky at a high resolution and a high SNR (Carretti et al. (2019)).

Survey / Telescope	Frequency [GHz]	FWHM [arcmin]	Declination Coverage	Stokes ^a	Sensitivity ^b		Status ^c	Reference(s)
					noise	offsets		
Haslam (various)	0.408	51	All-sky	<i>I</i>	1 K	3 K	3	Haslam et al. (1982)
Dwingeloo	0.82	72	-7° to $+85^\circ$	<i>I</i>	0.2 K	0.6 K	3	Berkhuijsen (1972)
CHIPASS (Parkes)	1.394	14.4	$< +25^\circ$	<i>I</i>	0.6 mK	30 mK	3	Calabretta et al. (2014)
DRAO (26-m) ^d	1.4	36	$> -29^\circ$	<i>QU</i>	12 mK	30 mK	3	Wolleben et al. (2006)
Villa Elisa ^d	1.4	35.4	$< +10^\circ$	<i>IQU</i>	9 mK	50 mK	3	Testori et al. (2008)
Stockert ^d	1.42	35	$> -30^\circ$	<i>I</i>	9 mK	50 mK	3	Reich & Reich (1986)
GMIMS-HB N	1.28–1.75	30	$> -30^\circ$	<i>IQU</i>	12 mK	unknown	1	Wolleben et al. (2010a)
STAPS (Parkes)	1.3–1.8	15	$< 0^\circ$	<i>IQU</i>	unknown	unknown	1	Haverkorn (priv. comm.)
HartRAO	2.326	20	-83° to $+13^\circ$	<i>I–Q</i>	25 mK	80 mK	3	Jonas et al. (1998)
S-PASS (Parkes)	2.3	9	$< 0^\circ$	<i>IQU</i>	0.1 mK	unknown	1	Carretti et al. (2013)
GEM	4.8–5.2	45	-52° to $+7^\circ$	<i>QU</i>	0.5 mK	unknown	0	Barbosa et al. (2006); Tello et al. (2013)
C-BASS	4.5–5.5	45	All-sky	<i>IQU</i>	0.1 mK	1 mK	0	This paper
QUIJOTE	11–19,30,40	≈ 60	$\gtrsim 0^\circ$	<i>[I]QU</i>	25 μ K	unknown	1	Génova-Santos et al. (2015a)
WMAP	22.8–94	49–15	All-sky	<i>IQU</i>	4 μ K	1 μ K	3	Bennett et al. (2013)
Planck LFI	28.4–70	32–13	All-sky	<i>IQU</i>	3 μ K	1 μ K	2	Planck Collaboration et al. (2016c)
Planck HFI	100–353	10–5	All-sky	<i>IQU</i>	0.2–0.5 μ K	1–5 μ K	2	Planck Collaboration et al. (2016c)
Planck HFI	545, 857	5	All-sky	<i>I</i>	0.4, 0.8 μ K	1 μ K	2	Planck Collaboration et al. (2016c)
CLASS	38–217	90–18	-68° to $+22^\circ$	<i>QU</i>	0.4 μ K	unknown	0	Harrington et al. (2016)

^a [I]QU denotes surveys where total intensity (Stokes I) is measured but with much larger systematic errors than for the linear polarization (Stokes Q and U). I–Q denotes a single linear polarization.

^b Approximate average total intensity sensitivity in Rayleigh-Jeans temperature after convolution to 1° FWHM resolution: “noise” is local rms; “offsets” is global systematic uncertainty.

^c Status 0: observations ongoing; 1: observations complete, reduction in progress; 2: preliminary results released; 3: Final data released.

^d An all-sky 1.4 GHz map in IQU has been assembled from the Stockert, DRAO and Villa Elisa surveys (Reich et al. 2004; Testori et al. 2008), but full details of its construction have not been published, and it is not clear if the currently-available version is the final one.

Figure 1.7: Table taken from Jones et al. (2018) indicating the most recent large-scale experiments, their frequencies, angular scales and sensitivity.

The ‘C-Band All Sky survey’ (C-BASS) is a full polarisation, all-sky survey mapping the sky at 5GHz. It was designed to provide complementary data to the *WMAP* and *Planck* missions along with future CMB experiments. At 5 GHz, the sky is dominated by synchrotron radiation but also suffers slightly from Faraday rotation depolarisation. With a sensitivity of 7.3 mKdeg, this makes the results from C-BASS the most powerful provider of information on the Galactic magnetic field. Using pixel simulation Bayesian analysis, Jones et al. (2018) demonstrate that when including a C-BASS data point with *Planck* and future satellite survey data, the CMB amplitude, synchrotron amplitude and spectral index can be much better constrained. The results of these pixel simulations are given in Figure 1.8, indicating that without C-BASS, it is extremely difficult to constrain the synchrotron amplitude and spectral index (Jones et al. (2018)).

1.3.2 Post-2020

Several space based telescopes have been proposed for missions with the main goal of improving constraints on r . The experiment that will launch soonest, is the *LiteBIRD* satellite. This satellite was proposed by the Japanese Space Agency (JAXA), in collaboration with NASA (Matsumura et al. (2014)). The ESA experiment *COrE* along




Figure 1.8: The probability distribution of the B -mode polarisation component for 3° pixels in a region of sky with low diffuse foreground emission. The dotted lines are Planck and a ‘LiteBIRD’-like satellite. The solid line includes the C-BASS data point. The vertical lines indicate the true parameter used in the model (Jones et al. (2018)).

with the American experiment *PIXIE* are no longer going ahead, however the forecasting methods put forward for these experiments have improved our understanding of the CMB parameter space (CORE Collaboration et al. (2016)).

The next generation of CMB experiments will be achieving significant improvements in sensitivity along with an expansion in frequency range. The *LiteBIRD* satellite will be looking at making a measurement of r , anticipating to constrain it down to $r \leq 10^{-2}$. The ground based experiments that will also contribute to measurements on r and constrain cosmological parameters along with measuring neutrino masses will be part of CMB-S4. CMB-S4 is an expansion of the group of experiments that make up the ground-based CMB collaboration, CMB-S3. CMB-S4 will include upgrading current experiments with additional detectors along with data obtained by the Simons’ Array (Abazajian et al. (2016)). Another balloon experiment that may accompany these include the balloon mission LSPE, (The LSPE Collaboration et al. (2012)) now LSPE-SWIPE that may potentially measure the sky at 95 GHz, 145 GHz and 245 GHz, no longer including the 43 GHz part of its design (de Gasperis et al. (2016)).

The decadal review by NASA, chooses four primary science cases and selects a mission to investigate each. The most recent decadal review has put inflationary physics as a primary candidate worthy of funding (PICO Collaboration et al. (2017)).




Figure 1.9: The polarisation sensitivity of ground-based experiments, C-BASS, QUIJOTE, CLASS and satellite missions *LiteBIRD*, *PICO* and *Planck*. The brightness temperature is in CMB temperature units, and the tensor to scalar ratio is calculated from the ratio of the the relative amplitudes of the *B*-mode to *E*-mode amplitudes. These sensitivities have been taken from Remazeilles et al. (2018) and scaled to a 3° pixel size, C-BASS (Jones et al. (2018)), QUIJOTE (Rubiño Marín et al. (2010)) where the sensitivity per pixel has been scaled to a 3° , CLASS (Errard et al. (2016)), *Planck* (Planck Collaboration et al. (2016b)), *LiteBIRD* is given by Table 2 of Remazeilles and Chluba (2018), *PICO* (PICO Collaboration et al. (2017)).

With this, the “Probe of Inflation and Cosmic Origins” (*PICO*) satellite has been proposed. This will consist of 21 frequency channels from 21 GHz up to 799 GHz. The main aim is to probe the energy scale at which inflation occurred and exclude inflationary models. They expect to set an upper limit on the tensor-to-scalar ratio with an error of $\sigma_r = 5 \times 10^{-5}$, $r < 10^{-4}$ at a 95% confidence level, 700 times lower than current constraints. These predictions include internal delensing and foreground removal. The mission will last four years achieving a noise of $0.63 \mu\text{K arcmin}$ in polarisation (Sutin et al. (2019)).

Experiment	σ_r
BICEP3+Keck+Planck	1.3×10^{-2}
EBEX10K+Planck	6.1×10^{-3}
AdvACTPol+Planck	3.7×10^{-3}
PIPER+Planck	8.3×10^{-3}
SPIDER+Planck	1.2×10^{-2}
Simons Array+Planck	4.5×10^{-3}
SPT-3G+Planck	4.5×10^{-3}
StageIV+Planck	1.8×10^{-3}
StageIV+Planck+C-BASS	7.4×10^{-4}
CLASS+Planck	3×10^{-3}
LiteBIRD	5.2×10^{-4}
LiteBIRD+C-BASS	1.2×10^{-3}
LiteBIRD+QUIJOTE	1.1×10^{-3}
Stage-IV+LiteBIRD	1.4×10^{-4}
Stage-IV+LiteBIRD+C-BASS	3×10^{-4}
*PICO	5.0×10^{-5}

Table 1.2: This Table shows combinations of experiments’ error on the tensor-to-scalar ratio taken from [Errard et al. \(2016\)](#), using Fisher matrices a polarised synchrotron and dust component along with their best performing CMB de-lensing algorithm. *PICO estimate taken from [PICO Collaboration et al. \(2017\)](#)

Due to the large difference in power between the CMB B-mode fluctuations and the foregrounds, component separation for polarisation is much more sensitive to foreground modelling uncertainties than measuring intensity. Particularly dust mis-modelling, [Remazeilles \(2018\)](#) found that mis-modelling a 2 component modified blackbody (MBB) as a single component MBB caused a bias on the re-constructed CMB B -mode at $r = 5 \times 10^{-2}$ by more than 3σ . This indicates that frequencies below 60 GHz and above 400 GHz are critical for B -mode separation. [Remazeilles \(2018\)](#) show that a PICO mission with a frequency coverage from 43 GHz to 462 GHz introduces a bias at large angular scales on the recovered B -mode spectrum because of the residual dust contamination, hence the full spectral range from 5 GHz to 800 GHz is necessary for low r detections.

It could be argued that increasing the frequency range of observations will introduce additional foregrounds. However, by introducing more frequency channels, the component separation is improved upon by the additional information available. So even if further foregrounds were present, the component separation methods would shed light on the appropriate spectral models for these and subsequently improve constraints further.

Figure 1.9 shows sensitivities for a 3 degree pixel for future and current experiments. These are superimposed on top of what we expect the polarised sky to look like. The three red lines indicate the varying degrees of synchrotron spectral indices, expected to lie between -2.8 and -3.1. The yellow line indicates the thermal dust component and the blue and orange lines are the CMB E and B mode frequency spectra in CMB temperature units respectively. ¹

We can clearly see that across the whole frequency spectra, foregrounds dominate. We can also see that at low ℓ , achieved by the full sky surveys, only PICO can obtain a sensitivity below the limit imposed by thermal dust at the lowest foreground levels around 100 GHz. It is therefore very important that the foregrounds are well sampled and modelled. The NextBASS frequencies would fill the gap uniformly in logarithmic frequency space, matching a similar frequency coverage of the lower frequencies to PICO.

1.4 X-BASS and NextBASS

1.4.1 Motivation

Even with C-BASS and QUIJOTE, the most foreground dominating region in frequency space is not sufficiently sampled at a sensitivity equivalent to the post-2020 experiments. These experiments will dedicate some observing time to foreground subtraction but cannot sacrifice observations required for their other science cases. As shown in Figure 1.4, the frequency space between 5 – 30 GHz is poorly sampled.

The Oxford Experimental Cosmology group has been working on the next generation of broad-band RF components and Digital Backend (DBE) technology, particularly for new telescopes such as SKA and MEERKAT. This technology opens a door to broadband sky mapping, made possible with faster and more bandwidth capable Field Programmable Gate Arrays (FPGA's).

The idea of this thesis is to design a foreground experiment capable of observing continuously in frequency, in the poorly sampled region of frequency space. To

¹Assuming the amplitudes of the E and B modes to be $A_{\text{CMB}}^E = 1e^{-6}$ and $A_{\text{CMB}}^B = 9e^{-9}$. The CMB brightness temperature is calculated as

$$T_{\text{CMB}} = \frac{A_{\text{CMB}} x^2 \exp x}{(\exp x - 1)^2} \quad (1.10)$$

such that $x = \frac{hf}{kT}$, where h is the *Planck* constant, f is the frequency, k is the Boltzmann constant and T is the temperature of the CMB, 2.726 K. We know from Equation 1.9 that r is the ratio of the power spectra of B and E , this means the ratio of amplitudes is $A_{\text{CMB}}^B/A_{\text{CMB}}^E \approx 1e^{-6}/9e^{-9} \approx 0.001$. So for $r = 0.001$, we expect to see a different in brightness temperature of only 10^{-2} .

achieve effective CMB cleaning, this experiment must therefore match the post-2020 experiments in sensitivity. The experiment proposed to do this has been named the “Next Band All-Sky-Survey”, a two-channel instrument with a single feed between 7 – 15 GHz and a feed array between 15 – 30 GHz. The feed array is required in order to achieve an equivalent sensitivity to synchrotron at 100 GHz of $1\mu\text{Karcmin}$ to match the *Planck* sensitivity (Planck Collaboration et al. (2016b)). Since the lower part of the band, from 7 – 15 GHz, only requires one feed horn to achieve this sensitivity, I propose a separate experiment to cover this part of the band, this is the “X-Band All-Sky-Survey”.

In this thesis, I demonstrate that “X-BASS” achieves a sensitivity to sufficient constrain the synchrotron emission and provide better measurements of the variation of synchrotron spectral index as a function of frequency. X-BASS therefore, operates as a stand-alone foreground instrument and a path-finder, or proto-type for “NextBASS”. X-BASS is also capable of disentangling the free-free and synchrotron emission in intensity observations. It will not however, be able to make a sufficient detection of the spinning dust peak, as it does not sample up the 23 GHz necessary for AME. “NextBASS”, has the capacity to both constrain synchrotron and makes measurements of the AME in polarisation to a sensitivity that has not been done before. This will both benefit the CMB community in terms of the accuracy of their foreground models, as well as increase our understanding of Galactic dust models and the mechanisms involved in the generation of polarised emission from spinning dust grains. In the next subsection, I will outline the planning and decisions made during the course of this thesis, which has resulted in the final design of NextBASS and X-BASS.

1.4.2 Objectives

As previously stated, more low-frequency data is required because current constraints on the CMB *B*-mode amplitude are not stringent enough, and that these constraints are improved with more accurate foreground data. I have also addressed why the foregrounds themselves are important to understanding the Galaxy’s magnetic field. In Table 1.3, I have listed some of the main drivers for a future CMB Foregrounds experiment, that X-BASS and NextBASS will be capable of achieving.

Experiment	Science Goal
X-BASS	To constrain synchrotron spectral curvature, which has not yet been achieved by any other foregrounds instrument.
X-BASS & NextBASS	To shed light on current polarisation foregrounds and whether they are correctly modelled.
X-BASS & NextBASS	Provide new foreground removal methods with updated models.
X-BASS & NextBASS	Achieve high fidelity, broad-band maps with low artefacts.
NextBASS	Obtain an upper limit on the AME polarisation better than what has been done to date in regions of weak spinning dust emission, and to detect a 1% level of polarisation in the brightest foreground regions.
NextBASS	To determine whether there are additional polarisation foregrounds that have not yet been detected by an increase in sensitivity between 15 – 30 GHz.

Table 1.3: Table showing the science goals and motivations behind X-BASS and NextBASS. X-BASS will be focused on the determination of the synchrotron spectral index and curvature, which is shown in Chapter 2. Both X-BASS and NextBASS will improve forecasting and modelling techniques, however it will take the sensitivity of NextBASS to determine whether the AME is polarised to 1%, and potentially detect other foregrounds in this region.

To meet these science goals, we therefore need an experiment that has excellent polarisation sensitivity and can simultaneously measure both temperature and polarisation. The simplest way of achieving this is by designing a radiometer-style experiment, similar to the architecture of C-BASS. Chapter 5 of this thesis is dedicated to radiometer design and the differences between the C-BASS, X-BASS and NextBASS receivers.

1.4.3 Planning

During the course of this PhD, a number of changes were made to the instrument design of NextBASS. X-BASS was actually later introduced as a concept after preliminary forecast modelling indicated that the synchrotron could be well constrained between 6-20 GHz. The original “X-BASS” was merely a proto-type of the NextBASS

receiver to be used with the 2m Clover dishes, [Maffei et al. \(2005\)](#), as a proto-type for the NextBASS receiver. For the purposes of this thesis, X-BASS is an independent instrument, the receiver of which will be implemented on the C-BASS South optics, once C-BASS is decommissioned.

The original plan for “NextBASS ” was to observe from 6 to 40 GHz in three channels. This would mean overlap with the satellite CMB experiments, which begin at 40 GHz and mean an excellent opportunity to probe thermal dust emissions in the regime where they become the dominant foreground. This would have separated the band into 3 channels, 6 – 12 GHz, 12 – 20 GHz and 20 – 40 GHz. In Spring 2017, this choice of frequency channels was abandoned for a number of reasons. The main reason being that the AME peak was broken into two bands, meaning it would not be possible to sample the whole AME peak at a constant sensitivity. The second reason was that the future satellite and ground experiments introduced lower frequencies, such as *LiteBIRD*’s 40 GHz channel, and *PICO*’s ambitious goal of 21 GHz. These experiments published that they will have dedicated foreground removal efforts using these channels ([PICO Collaboration et al. \(2017\)](#), [Suzuki et al. \(2018\)](#)).

The final reason is a matter of practicality, a two-band focal plane is easier to implement than a three-band, due to the focal plane size. A foreground experiment above 30 GHz would require significantly more feeds, meaning either design a larger telescope with a larger focal plane or consider a trade-off between frequency range and sensitivity. The optics of NextBASS are to be a similar resolution to C-BASS , hence fixing the diameter to 6m with a slightly under-illuminated aperture. The NextBASS frequencies were then changed to a 2 channel design from 7 – 15 GHz and 15 – 30 GHz, the AME peak being sufficiently sampled in the second channel.

NextBASS could be done as a single instrument, or as two separate instruments. Part of the preliminary research for this thesis was to consider all potential candidates for the NextBASS survey. As I previously mentioned, the X-BASS instrument could be either an individual instrument, or the lower channel of the NextBASS survey. In this case, the 15 - 30 GHz channel could be implemented on the 2m Clover dishes ([Maffei et al. \(2005\)](#)), with up to four high-frequency feeds, or a full-feed array implemented smaller, 3-4m telescope. A new 4m telescope is preferable since the Clover optics do not have a sufficiently large enough focal plane that can host the required number of feeds to obtain the low sensitivity measurements necessary for NextBASS. However, the atmospheric modelling of the C-BASS site in Klerefontein South Africa showed that due to the atmospheric water vapour content, the C-BASS (or future

X-BASS) site is an inappropriate choice for any instrument operating at frequencies above 22 GHz.

As a result, it was decided to keep X-BASS and NextBASS as two different instruments. Therefore, NextBASS is an ideal 6m telescope with a large focal plane and feed array, covering both the frequency channels discussed in a ground-based location that is less affected by atmospheric conditions. As a result, the sensitivity of the ‘band-1’ NextBASS channel, which matches the X-BASS frequencies, will have a slightly lower sensitivity. Figure 1.10, shows a flow diagram to show the development of the X-BASS and NextBASS instruments.

1.4.4 Instrument Requirements

This thesis has not quite taken on the structure of a systems engineering-style investigation. This is because I have chosen to use the simulations to demonstrate the impact of the instrument on the science, instead of using the simulations to identify the requirements of the instrument. As a result, I’ve made some assumptions on the instrument design that are widely accepted throughout the CMB community as being necessary for a polarised foreground experiment.

A few of these assumptions can be made about the type of survey, such that a 6 m instrument will achieve a $\approx 0.5^\circ$ beam resolution at 7 GHz, which is the necessary resolution required to observe CMB fluctuations. The survey will last 5 years in comparison to that of the C-BASS North instrument, as the predecessor of both X-BASS and NextBASS.

I have decided that it is necessary for the optical design to achieve a cross-polarisation of -40 dB and a beam efficiency of over 95% in a 1° beam. These are generally accepted as a necessary requirements for a CMB instrument. The cross-polarisation of -40 dB can be thought of as an indicator of the error on the polarisation such that it represents the power in the unwanted direction of linear polarisation. As a result, a -40 dB cross-polarisation means the polarisation we are trying to measure has $\approx 10^4$ times more power than that in the unwanted polarisation direction. This minimizes the leakage between linear polarisations and ensures accurate Stokes parameter maps, and subsequently accurate CMB E and B -mode measurements.

Similarly, the beam efficiency means that we want less than 5% of the total power to be present in the sidelobes of the main beam, as these sidelobes increase the system temperature of the optics and result in a lower sensitivity. Hence, both the optical and feed horn designs have been optimised to meet these requirements to ensure that, much later on in the experiments lifetime, during the map-making process, any

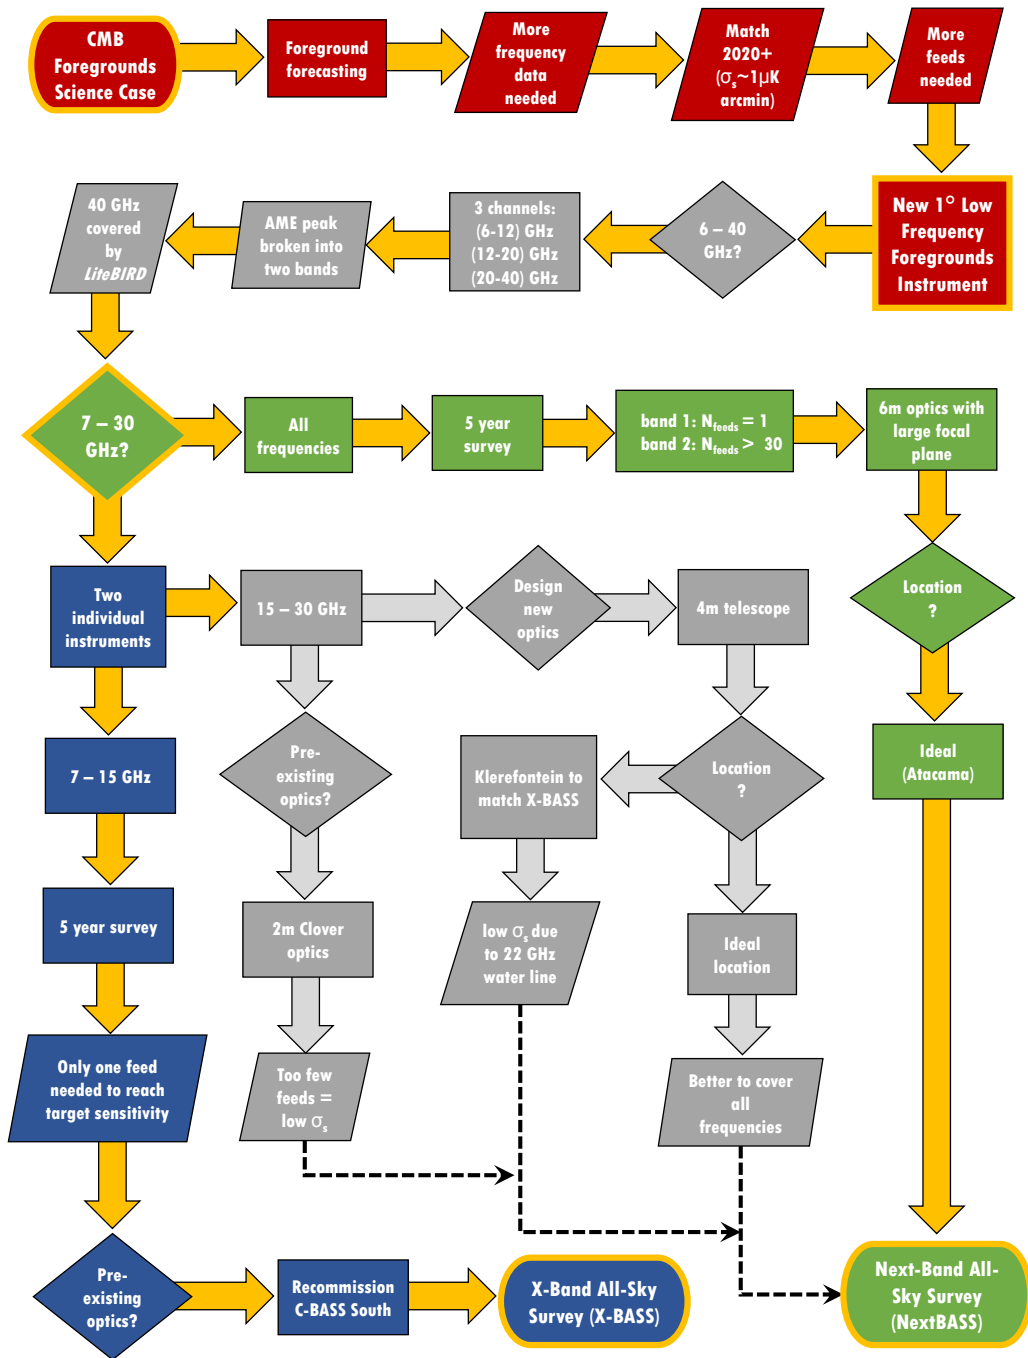


Figure 1.10: This flow diagram indicates the key decisions and alternative ideas discussed for the NextBASS experiment. The right-hand side shows that only one feed is required to reach a $1 \mu\text{K arcmin}$ which could be implemented on the C-BASS optics as a standalone experiment, ‘X-BASS’. The left-hand side is the design of an ideal instrument that is capable of reaching the target sensitivity across the whole band. The grey boxes are ideas that were not usable.

Requirement	X-BASS	NextBASS
Frequency	7 – 15 GHz	7 – 15 GHz, 15 – 30 GHz
Bandwidth	20%	20%
Number of Feeds	1	1, 33
Optics	6m Shaped Cassegrain (C-BASS South)	6m Cross-Dragone
Location	Klerefontein, SA	Plano de Chajnantor, Chile (or La Palma)
Cross-polarisation	-40 dB	-40 dB
Beam Efficiency (1°)	> 95%	> 95%
Edge illumination	-40 dB	-25 dB
System temperature	20 – 25 K	15 – 29 K
Sensitivity	12 μ Kdeg	11 μ Kdeg, 2 μ Kdeg

Table 1.4: A summary of parameters decided on for the X-BASS and NextBASS instruments. X-BASS is a recommissioning project of the C-BASS South, operating with a single feed at 7 – 15 GHz. The edge taper and cross-polarisation matches accordingly with the C-BASS South optics, optimised to minimise ground spillover. NextBASS operates in two bands, NextBASS¹ from 7 – 15 GHz with a single feed and NextBASS² from 15 – 30 GHz requiring 33 feeds to achieve a similar sensitivity to synchrotron as done by the lower channel. The optics are to use a cross-Dragone configuration, maintaining a low cross-polarisation. The system temperature is based on calculations involving the receiver and the propagation of the sky temperature based on the atmospheric conditions at each location.

instrumental polarisation signals are well characterized and reduced to have a minimal impact on the data being taken. Table 1.4, I list the requirements of the X-BASS and NextBASS instruments.

1.5 Thesis Overview

The structure of this thesis works through the design of the X-BASS and NextBASS instruments that have been used in the science simulations. Now that the key science aims have been outlined, Chapter 2 will be a presentation of foreground forecasting methods and the results of NextBASS and X-BASS in conjunction with current experiments such as C-BASS, *Planck*, *WMAP* and Haslam, along with the post-2020 *LiteBIRD* instrument. Chapter 2 will demonstrate how well X-BASS and NextBASS can constrain foreground parameters along with an investigation into polarised foreground models and whether using the incorrect model can have a noticeable affect on foreground cleaning.

In Chapter 3, I describe the optical design work carried out for both X-BASS and NextBASS, including focal plane modelling for the NextBASS optical candidates and the use of C-BASS photogrammetry data to test the surface accuracy of the dish up to 15 GHz. Chapter 4 then describes the feed horn design for both NextBASS and

X-BASS, where I have optimised a feed horn to match the C-BASS optics and used far-field Gaussian beam matching to design the NextBASS feed.

In Chapter 5, I use the literature to explain why a radiometer is the best choice of instrument for a new low frequency foregrounds experiment and I then outline the different receiver architectures that could be used for X-BASS and NextBASS. I conclude with a schematic of the X-BASS architecture and an architecture for both NextBASS band-1 (NextBASS¹) and band-2 (NextBASS²). The end of this Chapter is concerned with the design and testing of proto-type microstrip filters, an essential part of the instrument development for these experiments.

Chapter 6 is a robust description of the generated sensitivities that were used in the forecasting in Chapter 2. This includes atmospheric modelling of different potential sites for the NextBASS experiments, followed by a detailed model of a frequency dependent receiver chain. I then explain the sensitivity generation of NextBASS and how these are affected by more realistic modelling of the beam and different survey parameters.

Finally, I will summarise the work done in this thesis and address what has been achieved in each Chapter and how it relates to the motivation and justification of X-BASS and NextBASS. I will then describe the future work and development that needs to be done for X-BASS and NextBASS, and the CMB foregrounds community.

Chapter 2

The impact of X-BASS and NextBASS on CMB Foreground Component Separation

“The surest way to become Tense, Awkward and Confused is to develop a mind that tries too hard – one that thinks too much.”

Benjamin Hoff, The Tao of Pooh

This chapter investigates the complexity of diffuse Galactic foregrounds and how increasing the number of frequency data points can help resolve that complexity. The forecasting techniques used in this chapter are based on simulations outlined in Chapter 2, Section 2.2, Jew (2017). To demonstrate the impact of C-BASS on the results of CMB component separation, Jew (2017) simulated the component separation of C-BASS, using a parametric pixel-fitting process for individual pixels. C-BASS was used in conjunction with existing surveys such as Haslam, *WMAP* and *Planck* and a ‘*LiteBIRD*-like’ future survey based on the sensitivities of *LiteBIRD* from Suzuki et al. (2018). (Jew (2017)). The temperature simulations are carried out over 1° pixels and with the sensitivities and frequencies from HASLAM, *WMAP*, *Planck* and C-BASS, where as the polarisation simulations use 3° pixels to obtain a trade-off between sensitivity on the sky, and resolution. The *LiteBIRD* frequencies are necessary to constrain thermal dust emission.

The central frequency of C-BASS is low enough for the map to be dominated by synchrotron and free-free emission but high enough so that Faraday rotation, and the resulting depolarisation are small across the sky. It was shown that C-BASS enables more accurate removal of contaminating foregrounds from measurements of the CMB, particularly in polarisation where the *B*-mode is weak. The C-BASS pixel

simulations also included a free-varying synchrotron curvature term, C_s , which compensates for the spectral flattening due to self-absorption of photons (this discussed in Section 1.2.1). It was found that in polarisation, *Planck* and *LiteBIRD* were unable to constrain C_s , but the addition of the C-BASS data point helped constrain all synchrotron spectral parameters and reduce the bias in estimates of the CMB amplitude. These simulations have been used in the C-BASS project paper (Jones et al. (2018)) and are being developed for the C-BASS forecasting paper (Jew et al. (in prep.)). The conclusion of these simulations was that more low-frequency maps are required to allow more degrees of freedom when fitting for extra free-parameters, and that additional low-frequency surveys would also help reduce the degeneracy between synchrotron, free-free and spinning models. These conclusions are the motivation for the NextBASS simulations, which will investigate how further low-frequency data points in conjunction with current and future experiments can constrain more complex foreground models with additional free-parameters.

The C-BASS sky model is based on a simplified version of the *Planck* sky model in Table 4 of Planck Collaboration et al. (2016b). The true foreground amplitudes in each pixel use the component separation results by Planck Collaboration et al. (2016b) to estimate foreground intensities for different pixels. The CMB amplitude is set to $75\mu\text{K}$ in intensity and zero in polarisation using a power spectrum with the standard cosmological parameters for $r \approx 0.1$ at the angular scale of the simulated pixel. The pixel sensitivities for the surveys values have been scaled to 1° in intensity since all components, including the CMB amplitude are detected at a high signal-to-noise on this scale, whereas in polarisation, 3° pixels have been used to give maximum pixel-to-pixel fluctuations in the CMB B -mode signal.

The NextBASS simulations will utilise Bayesian parametric fitting for the sky model with the PYMC python package (Patil et al. (2010)) to explore the posterior distribution using the Metropolis Hastings algorithm (see Appendix A) in order to fit the sky model onto the simulated data. Each set of simulations has been run with a different sky model, this means a different set and/or number of free-parameters for each run. To investigate how the addition of NextBASS to C-BASS can improve foreground constraints, I use the same sky models as used by Jew (2017). This means investigating a model that has a straight synchrotron spectrum of spectral index, $\beta_S = -3.1$, and another model that includes an additional free-parameter to describe the synchrotron spectral curvature, C_s , which I set to be $C_s = 0.3$. This value for the synchrotron curvature has been arbitrarily selected to account for the positive curvature of the synchrotron spectrum at higher energies. It is the upper end of the

theoretical limit to demonstrate the difference in spectra when correct and incorrect models have been fitted to the simulated datasets (Remazeilles et al. (2018)).

I do this with an increasing number of data points. In intensity, I start with *Planck*, *WMAP* and Haslam, then I add C-BASS, I then add X-BASS and finally I perform a final run with *Planck*, *WMAP*, Haslam, C-BASS and NextBASS. In a similar fashion, the polarisation simulations begin with *Planck* and *LiteBIRD*, then add C-BASS, X-BASS and finally the NextBASS map values. I also perform an investigation into whether the different combinations of maps can determine whether a sky model has been fitted to the data correctly. I do this by taking a sky model that includes a true synchrotron curvature term of $C_s = 0.3$, but I fit for a straight spectral index. If a large bias is introduced on any of the foreground parameters for a given combination of instruments, then it becomes obvious that the data is sensitive to the incorrect model fitting.

To advance from the C-BASS simulations, I will be adding another layer of complexity in polarisation that will explore the addition of a 1% AME component. The NextBASS frequencies cover the rise and fall of the AME peak frequency, making them ideal to cover any additional complex foregrounds in polarisation. I add the spinning dust amplitude, A_{AME} as 1% of the intensity amplitude for each pixel and the peak frequency as the additional two free-parameters. I then introduce a similar modelling error to the synchrotron component where I include the AME parameters as fixed values in the model, but run a simulation that does not fit the data for any polarised AME.

This chapter will describe in detail the structure of each the NextBASS simulations and the results that go with, demonstrating the X-BASS and NextBASS can improve foreground constraints considerably, along-side C-BASS, *Planck* and *LiteBIRD* for different complexities of foreground models.

2.1 Component Separation

There are a number of different methods that are used to simulate a data set and investigate constraints on both the foreground and CMB parameters using data taken at different frequencies. These include dealing with individual pixels of the sky or modelling full sky data. In this section, I will outline the four most commonly used methods. These are template fitting, which uses a series of sky templates to fit data across regions of sky; internal linear combination (ILC), which is a blind component separation method that uses the variance to estimate foreground contributions;

independent component analysis (ICA), which uses the frequency domain of auto or cross-correlation spectra and a mixing matrix to parametrise foregrounds; and Bayesian parametric fitting, which relies on a priori knowledge of the models.

2.1.1 Template Fitting

Template fitting is also known as cross-correlation analysis. It assumes that a sky map vector \mathbf{T} at a given frequency, is the sum of template map vectors, \mathbf{t}_j , multiplied by template correlation coefficients, α_j and corrupted by a noise term, \mathbf{n} , which includes instrumental contributions alongside the CMB. This is given by:

$$\mathbf{T} = \sum_j \alpha_j \mathbf{t}_j + \mathbf{n}, \quad (2.1)$$

where \mathbf{T} is the sky map vector at a given frequency.

The cross-correlation coefficient scales the j^{th} template to the frequency of the map. The CMB map, $\hat{\mathbf{T}}_{\text{CMB}}$, is then constructed from the sky and templates, where the cross-correlation coefficient is found by minimising the variance on the CMB. The covariance matrix, \mathbf{C} contains the noise properties of both the instrument and the CMB. The χ^2 function and the resulting CMB map, $\hat{\mathbf{T}}_{\text{CMB}}$ are :

$$\chi^2 = (\mathbf{T} - \sum_j \alpha_j \mathbf{t}_j)^T \mathbf{C}^{-1} (\mathbf{T} - \sum_j \alpha_j \mathbf{t}_j), \quad (2.2)$$

and

$$\hat{\mathbf{T}}_{\text{CMB}} = \mathbf{T} - \sum_j \alpha_j \mathbf{t}_j, \quad (2.3)$$

where α_j are the coefficients that minimise χ^2 .

In this method, even when the signal-to-noise ratio of components in individual pixels is low, they can still be constrained because the template is fitted across the whole sky. The main problem of the template-fitting technique is that it has to assume that the spectral properties of each component are uniform across the whole sky. Using fixed foreground components means extracting the CMB map is difficult when there are multiple sky templates. Template fitting can only be used for large sky areas and also requires a suitable template map for each component, i.e. a map that is dominated by a single emission mechanism. None the less, template-fitting has proven to be an effective method of foreground removal, with usage on *WMAP* (Komatsu et al. (2009)) and the simulated *Planck* data (Leach et al. (2009)). More recently, the C-BASS collaboration used the template fitting method to generate a total intensity

map of the North Celestial Pole (NCP) region of the sky. They combine *WMAP* and *Planck* data using the C-BASS map as the synchrotron template to investigate contribution of diffuse foreground emissions at frequencies $\approx 20 - 40$ GHz (Dickinson et al. (2018b)).

2.1.2 Blind and Semi-blind Component Separation

Blind and semi-blind component separation methods makes the assumption that all sky components are statistically independent of one another. These methods require fewer a priori assumptions than any other separation method, requiring only an a mixing matrix. By making minimal assumptions about the foreground model, the priority can be shifted to reconstructing the CMB from its well characterised blackbody spectral distribution. The advantages of blind and semi-blind methods is their ability to deal with unknown or complex foreground contamination when reconstructing a cleaned CMB signal. This is particularly advantageous for real data sets as it means the instrumental effects and data set combinations can be dealt with immediately, leaving the complex foreground modelling and CMB reconstruction for later in the analysis.

Internal linear combination (ILC) is a blind component separation method. The main assumption is that the model can be written as a linear combination of the 1-D vectors with respect to frequency of that describe the CMB map, $\mathbf{T}_{\text{CMB}}^{(i)}$, the diffuse galactic foregrounds, $\mathbf{T}_{\text{f}}^{(i)}$, and an experimental noise term, $\mathbf{n}^{(i)}$, where i is the observing channel frequency:

$$\mathbf{T}^{(i)} = \mathbf{T}_{\text{CMB}}^{(i)} + \mathbf{T}_{\text{f}}^{(i)} + \mathbf{n}^{(i)}. \quad (2.4)$$

Each index of this vector represents a pixel on the sky at a given frequency. These temperatures are expressed in CMB temperature units¹, which means the CMB component is independent of the observing channel. The image for each channel, i , is given specific weight, w_i , to minimise the foreground and noise terms in the residual

¹The CMB temperature units are related to the Rayleigh Jeans temperature units by the relation:

$$T_{RJ} = T_{CMB} \exp\left(\frac{h\nu}{k_B T_{CMB}}\right) \times \frac{\exp\left(\frac{h\nu}{k_B T_{CMB}}\right)}{\left[\exp\left(\frac{h\nu}{k_B T_{CMB}}\right) - 1\right]^2}, \quad (2.5)$$

where h is the Planck constant, ν is the frequency, k_B is the Boltzmann constant and T_{CMB} is the temperature of the CMB (2.71 K). This conversion changes the CMB blackbody spectral shape to be flat and taper off at a frequency of ≈ 1000 GHz.

map. This means the solution type can be written as (Bennett et al. (2003)):

$$\hat{\mathbf{T}}_{\text{CMB}} = \sum_i (w_i \mathbf{T}^{(i)}), \quad (2.6)$$

$$= \sum_i (w_i \mathbf{T}_{\text{CMB}}^{(i)} + w_i \mathbf{T}_{\text{f}}^{(i)} + w_i \mathbf{n}^{(i)}). \quad (2.7)$$

If the constraint $\sum_{i=1}^{N_0} w_i = 1$ is imposed, then the summation of the CMB map over all frequencies can be described by:

$$\hat{\mathbf{T}}_{\text{CMB}} = \mathbf{T}_{\text{CMB}} + \sum_i (w_i \mathbf{T}_{\text{f}}^{(i)} + w_i \mathbf{n}^{(i)}). \quad (2.8)$$

The only variable terms are in the summation, so under the assumption that the CMB is independent from the foreground and noise terms, the weights have to minimize the variance of the CMB map. The process of minimising the noise term requires many values from which to calculate the variance meaning that it is not possible to perform ILC on a single pixel.

The ILC method can also be used in spherical harmonic space by introducing weightings for different angular scales. The sky is dominated by foreground emission on large angular scales, where as the smaller scales are usually dominated by noise. To do this version of ILC, the CMB power spectrum coefficients, $a_{\ell m}$, are estimated from measurements of $a_{\ell m}^i$ at the corresponding frequencies, i . These power spectrum coefficients are given by:

$$\langle |a_{\ell m}|^2 \rangle = \mathbf{w}_{\ell}^T \mathbf{C}_{\ell} \mathbf{w}_{\ell}, \quad (2.9)$$

where, \mathbf{w}_{ℓ} refers to a weighting system designed to minimise the power spectrum coefficients $a_{\ell m}$. The resulting weighting function is then a summation over the column vector \mathbf{e} consisting of ones when normalised, which then normalises the weight such that $\mathbf{e} \cdot \mathbf{w}_{\ell} = 1$. In this case, the weights are given by (Eriksen et al. (2004)):

$$\mathbf{w}_{\ell} = \frac{\mathbf{C}_{\ell}^{-1} \mathbf{e}}{\mathbf{e}^T \mathbf{C}_{\ell}^{-1} \mathbf{e}}. \quad (2.10)$$

It is also possible to use the ILC method in the Needlet domain. Needlets are a special type of spherical wavelet that allow for a simple reconstruction formula. Needlets are localised to a finite number of multipoles in spherical harmonic space. The localisation properties of the wavelets mean the weights of components can re-adjust depending on the local conditions of foreground contamination both over the sky and over the angular scales. This method has been used by Delabrouille et al. (2009) to produce a CMB map from the 5-year *WMAP* temperature data. The *Planck*

Collaboration have developed a pipeline ‘NILC’ to produce maps using this method (Planck Collaboration et al. (2016b)). Remazeilles (2018) has used the NILC pipeline to reconstruct the μ -anisotropies in the CMB maps whilst performing foreground removal.

One commonly used, semi-blind component separation method is Independent Component Analysis (ICA). Since there are spatial correlations between the foregrounds, ICA can be used to disentangle the CMB from the noise and a generalised foreground model on smaller scales. ICA then blindly recovers the full mixing matrix by using the independence property of different components. Similar to ILC, it assumes that a set of sky maps at a range of frequencies are each linear sums of the CMB and its foregrounds. The reason it is only semi-blind is because it assumes that the foregrounds are highly non-Gaussian in nature. The FASTICA algorithm developed by Hyvarinen (1999), is an algorithm that performs component separation by maximising the non-Gaussianity of data. Before the algorithm is performed, the data have undergone a linear transform to create a new, uncorrelated data matrix. After a number of iterations, the algorithm generates a weight vector to maximise the non-Gaussianity to extract the parameters. FASTICA can also be used in multiple component extraction.

In spherical harmonic space, ICA can be implemented as Spectral Matching ICA algorithm, SMICA, (Delabrouille et al. (2009)). By establishing a parametric model of the sky to the auto- and cross-spectra of input maps, the resulting power spectrum, C_ℓ^{ij} is then described by the sum of the auto- or cross-spectra of the CMB, $C_\ell^{\text{CMB},ij}$, the foregrounds, F_ℓ^{ij} , and the noise terms, N_ℓ^{ij} , where i and j represent the maps used in the model:

$$C_\ell^{ij} = C_\ell^{\text{CMB},ij} + F_\ell^{ij} + N_\ell^{ij}. \quad (2.11)$$

Often the spectra are binned into ℓ bins such that each bin parametrises the foregrounds by a symmetric amplitude matrix and a mixing matrix. This means that correlation between foreground parameters is accounted for. The advantage of this method is that if the spherical harmonic coefficients are independent random variables with a variance proportional to the power spectrum, by minimising the negative log-likelihood, any unknown parameters can be identified. The level of blindness in SMICA can be adjusted by adding any additional parametric component to allow for a gradually more complex design of the foreground model.

2.1.3 Bayesian Parametric Fitting

The overall idea of parametric component separation is to fit a parametric model of the sky and possibly instrumental parameters, to data at a range of frequencies. Bayes theorem relates the probability of a model given the data to the probability of the data given the model, which is easier to calculate. The probability of the model is based on prior knowledge of the model and the data. We may or may not have know the prior function for a parameter in a model. If we do not know the prior on a model, we can choose an uninformative prior and investigate whether the given data will bias the separated parameters.

In Bayesian parametric fitting, a model of the foreground components is parametrised. We model the true sky emission with a vector of parameters, θ , and we want the posterior probability distribution of those parameters given the observations \mathbf{d} . The posterior distribution is proportional to the prior on a parameter, $p(\theta)$, times the likelihood, $\mathcal{L}(\theta)$ and is given by:

$$p(\theta|\mathbf{d}) = \frac{p(\mathbf{d}|\theta)p(\theta)}{p(\mathbf{d})} \propto \mathcal{L}(\theta)p(\theta), \quad (2.12)$$

where $p(\mathbf{d}|\theta)$ is the probability of the data given the model, $p(\theta)$ is the prior, $p(\mathbf{d})$ is the probability of the data, which is difficult to calculate.

Bayesian methods allow us to marginalize over nuisance parameters. The sky model contains both fixed and free-varying parameters. The free parameters are used to describe properties of each foreground component that we are interested in. The aim of using Bayesian parametric model fitting is to demonstrate the resulting error bars on each foreground parameter for a combination of maps and regions. Bayesian methods are extremely useful as they allow us to calculate the propagation of errors from the map sensitivities to the sky spectrum. However, if the sky model is multi-dimensional and the parameters are highly correlated, then Bayesian modelling becomes computationally expensive. This method is widely used in the literature and is applied in the COMMANDER algorithm developed by [Eriksen et al. \(2008\)](#), and [Eriksen et al. \(2004\)](#), used to generate the *Planck* CMB maps [Planck Collaboration et al. \(2016b\)](#).

2.1.4 The importance of component separation

Each of these methods are well documented and have been used in different CMB analysis and forecasting scenarios. When modelling an experiment that looks at a large percentage of the sky, we can either model full-sky datasets or focus on individual

regions of sky; both have their advantages. For an experiment that is aiming to make a measurement of CMB B -modes at low multipoles, the full-sky approach to simulating data can be beneficial, as the variation of spectral indices across the sky can change the extracted CMB power spectrum. For experiments looking at smaller regions of sky, or focusing purely on foreground extraction, the single-pixel method is less computationally expensive meaning more accurate sky models can be fit to each region. For example, the *Planck* Collaboration developed a pixel-based algorithm for their parametric fitting code, COMMANDER. They have also developed this to produce all-sky CMB and foreground component maps (Planck Collaboration et al. (2016b)).

Errard et al. (2016), use a Fisher matrix analysis to forecast how well current and future CMB missions perform. When including the low-frequency survey data-point of C-BASS, the noise degradation factor, which specifies the amount that foregrounds increase the final uncertainty, is decreased by up to a factor of 3. In another paper, Remazeilles (2018) showed that ‘lever points’ at the frequency extremes are needed to constrain synchrotron and thermal dust. Using the frequency channels of the proposed NASA mission, the Probe of Inflation and Cosmic Origins, *PICO* (Sutin et al. (2019)) in forecasts, they demonstrate that without frequencies > 400 GHz, the dust temperature is poorly constrained with lower accuracy leading to a bias on the reconstructed B -mode of $r = 10^{-3}$. It is now clear that frequencies below 60 GHz and above 400 GHz are extremely important in extracting the CMB B -mode power spectrum at the level of a tensor-to-scalar ratio of $r \approx 10^{-3}$ (Remazeilles (2018)).

2.2 NextBASS Simulations

Bayesian parametric fitting is most appropriate because the science goals of NextBASS include determining the properties of the diffuse Galactic Foregrounds, not just achieving a clean extraction of the CMB map. By using uninformative priors on the foregrounds, we expect to see a truthful representation of how the increasing number of frequency data can improve constraints. This is useful for identifying whether a component has been modelled incorrectly. Since I am simply testing the performance of two new instruments regarding their ability to constrain foregrounds, single pixel simulations are sufficient to explore a variety of different foreground intensities without the computational intensity of running full-sky component separation methods.

This section covers the different spectral models, parameter values and priors used in the NextBASS simulations. The spectral models used in these simulations are based on models that are widely used in the literature, each model represents the physical emission mechanism that it is describing. Each of these models are described by a set of parameters. The Bayesian fitting method separates the parameters for each simulation to generate simulated data. However, the initial, or ‘true’ parameter value is required to start the model. The ‘true’ parameter values have been chosen from foreground component separation data from the *Planck* satellite (Planck Collaboration et al. (2016b)).

Every free-parameter in the model requires a prior that does not introduce bias into the simulated data. In this section, I also discuss the choice of priors with an over view on the Jeffreys prior. The Jeffreys prior is an uninformative prior, and has been used as the prior for the free-parameters in these simulations. I then specify the prior functions for the free-parameters in the NextBASS simulations.

To test whether a combination of experiments are capable of determining if a sky model has been fit to the data correctly, a modelling error can be introduced to the simulated data. This is done by running the simulation with a ‘true’, or correct model, but performing the spectral fitting process for an incorrect model. If the given combination of frequencies and sensitivities cannot sufficiently constrain the model parameters, then the simulated data will be unaffected by whether or not the model has been fitting correctly. However, if the combination of experiments are sensitive to the foregrounds, then the incorrect fitting process will result on a visible bias on the posterior distribution. This means that if we were to be fitting models to real data, the resulting bias could indicate that we are not using a correct sky model. Therefore, it is important to perform these simulations to demonstrate that more frequency data is essential to rule-out incorrect sky models. In the NextBASS polarisation simulations, I use modelling errors to investigate whether the additional frequency channels of NextBASS can distinguish between correct and incorrect model fitting. I perform the modelling errors on the synchrotron curvature and with a polarised AME component and present the results later in this chapter.

2.2.1 Spectral Models

The first stage in setting up the simulations is to define a set of spectral models that describe the sky. These models are based on what current observations have found. In intensity, we expect to see emission from synchrotron, free-free emission,

thermal dust and the anomalous microwave background. We know that free-free is unpolarised therefore it is excluded in the polarisation simulations. We also expect the AME to be negligibly polarised at 1% (see Section 1.2.4) so this is typically excluded in polarisation simulations, however, to investigate whether NextBASS could detect a 1% polarised AME component, I include the AME in some of the polarisation simulations.

The CMB is described by a blackbody spectrum in terms of T_{CMB} , which is the mean temperature of the CMB set to 2.7255K for these simulations. By setting $x = h\nu/k_{\text{B}}T_{\text{CMB}}$, where h is the Planck constant, k_{B} is the Boltzmann constant and ν is the frequency, the Rayleigh-Jeans brightness temperature, s_{CMB} is given in terms of the CMB amplitude A_{CMB} by (Fixsen (2009)):

$$s_{\text{CMB}} = A_{\text{CMB}} \frac{x^2 e^x}{(e^x - 1)^2}, \quad (2.13)$$

In polarisation, the CMB B -mode is set to zero. This achieves the equivalent to assuming perfect delensing of the E -mode signal such that the tensor-to-scalar ratio, r is equal to zero. This scenario means that a non-zero detection of the B -mode signal can then be interpreted as the bias introduced by a lack of frequency data, or the model itself.

Synchrotron emission has a spectrum that can be approximated to a power-law with a negative spectral index which is typically between 2.5 – 3.2. The nature of synchrotron emission means that along any line of sight, there are multiple populations of emitting electrons with potentially different spectral indices. The moment expansion treats the spectral function as an addition of spectral indices along the line of sight, with the intention of expressing the spectrum using the fewest number of parameters without losing important information on the spectral shape (Chluba et al. (2017)). The main draw-back with the moment expansion model is that there are currently not enough frequency data to constrain these multiple synchrotron parameters. The synchrotron spectral function can be written as:

$$\begin{aligned} \langle I_\nu(p) \rangle = & A_0 \left(\frac{\nu_c}{\nu_0} \right)^\beta \left(1 + \frac{1}{2} \omega_{22}^* \ln^2 \left(\frac{\nu_c}{\nu_0} \right) \dots + \right. \\ & \dots \frac{1}{6} \omega_{222}^* \ln^3 \left(\frac{\nu_c}{\nu_0} \right) \dots + \\ & \left. \dots \frac{1}{24} \omega_{2222}^* \ln^4 \left(\frac{\nu_c}{\nu_0} \right) + \frac{1}{120} \omega_{22222}^* \ln^5 \left(\frac{\nu_c}{\nu_0} \right) + \dots \right) \end{aligned} \quad (2.14)$$

where the $\omega_{2\dots 2}^*$ terms correspond to successive moments that depend on the synchrotron amplitude to define their spatial weighting. If these are uncorrelated the the

overall weighting does not alter the shape of the spectral energy distribution (SED). This simplifies considerably when considering the SED to be Gaussian, meaning the spectrum can be described to be a power-law with a curved spectral index. The power law that describes the synchrotron spectrum used in these simulations is given by :

$$s_s = A_s \left(\frac{\nu}{\nu_0} \right)^{\beta_s + \frac{1}{2} C_s \ln(\nu/\nu_0)}, \quad (2.15)$$

where A_s is the amplitude at a frequency ν_0 , β_s is the effective spectral index for synchrotron and C_s is the curvature term.

Synchrotron emission can be polarised up to 70 % in ordered magnetic fields (see Section 1.2.1). The interstellar magnetic field has a turbulent component and therefore the polarisation fraction of diffuse galactic synchrotron emission will be lower than this across the sky. The south galactic cap, and at other high galactic latitudes, synchrotron emission can be up to 40% polarised (Carretti et al. (2019)). At lower frequencies and close to the Galactic plane, the synchrotron emission will be less polarised due to Faraday depolarisation (Vidal et al. (2015)). For these simulations, we focus only on the B -mode polarisation, since we want to show the impact of NextBASS on the CMB B -mode amplitude. In future simulations, it should be noted that there is evidence that in real Galactic emission, the E -modes are a factor of ≈ 2 times brighter than the B -modes (Kamionkowski and Kovetz (2016)).

Free-free emission is generally well understood, and can be approximated by the two parameter model which depends on the effective emission measure EM and the physical electron temperature of the free-free emitting cloud, T_e . The model can be described by (Draine (2011)):

$$s_{\text{ff}} = T_e (1 - \exp^{-\tau}) \quad (2.16)$$

where

$$\begin{aligned} \tau &= 0.05468 T_e^{-3/2} \nu_9^{-2} EM g_{\text{ff}}, \\ g_{\text{ff}} &= \log \left(\exp \left[5.960 - \sqrt{3}/\pi \log \left(\nu_9 T_4^{-3/2} \right) \right] + e \right), \end{aligned} \quad (2.17)$$

such that ν_9 is the frequency in GHz and T_4 is the electron temperature divided by 10,000. In these simulations, polarised free-free emission has been neglected as it is assumed to only be weakly polarised over the whole sky (N.Macellari et al. (2011)). This is a reasonable approximation for a sky model, however on smaller scales, it is possible that free-free emission can be significantly polarised at the edges of HII regions (see Section 1.2.2).

The anomalous microwave emission, although poorly understood can be fairly well approximated by shifting a template spectrum in logarithmic frequency-brightness space. This is described by:

$$s_{\text{sd}} = A_{\text{AME}} \left(\frac{\nu_0}{\nu} \right)^2 \frac{F(\nu\nu_{\text{p0}}/\nu_{\text{p}})}{F(\nu_0\nu_{\text{p0}}/\nu_{\text{p}})}, \quad (2.18)$$

where F is the template spectrum, A_{AME} is the amplitude at frequency ν_0 , ν_{p} is the peak frequency and ν_{p0} is the peak frequency of the template. This prescription was originally taken up by the *WMAP* and also used in *Planck* models (Gold et al. (2008), Planck Collaboration et al. (2013a)).

Currently, there are only upper limits of 1% expected in AME polarisation (Draine and Lazarian (1998), Dickinson et al. (2018a)). This would be neglected usually however, given the sensitivity of NextBASS, it may be possible to obtain information on the behaviour of a polarised AME component and in particular, comment on the additional complexity in foregrounds. It has been suggested in other forecasting papers that even a weak additional polarised foreground, may be enough to affect measurements of the tensor-to-scalar ratio. For example in Remazeilles et al. (2018), using Commander, the bias on $r = 10^{-3}$ is reduced from 3.0×10^{-3} to 0.6×10^{-3} when including a 1% AME polarisation and full lensing in the component separation.

Interstellar dust grains radiate thermally as grey bodies. Their spectra can be modelled by a modified blackbody (MBB) spectrum given by:

$$s_{\text{d}} = A_{\text{d}} \left(\frac{\nu}{\nu_0} \right)^{\beta_{\text{d}}+1} \frac{\exp(\gamma\nu_0) - 1}{\exp(\gamma\nu) - 1}, \quad (2.19)$$

where $\gamma = h/(k_{\text{B}}T_{\text{d}})$, A_{d} is the amplitude of emission at reference frequency ν_0 , T_{d} is the thermal temperature of the dust grain and β_{d} is the emissivity spectral index. Due to the asymmetric dust grains aligning with the local magnetic field, polarised thermal dust fractions vary from 0 to greater than 20% (Ade et al. (2014)). The *Planck* Collaboration (Planck Collaboration et al. (2013b)) find that the ratio of B to E mode power of polarised thermal dust emission is around 0.5. In these models, the polarised power only considers the B -mode polarisation signal. There are significantly more complex thermal dust models being developed by Hensley and Bull (2018) that better describe the behaviour of thermal dust at low and high frequencies. The grey body model is suitable for NextBASS simulations because thermal dust does not dominate the frequency range we are interested in.

2.2.2 Parameter values and priors

For the Bayesian parametric model fitting, we require a prior, $p(\theta)$. If you don't have any prior knowledge of the true parameter value then you need to use an uninformative prior. An obvious choice of uninformative prior is a flat or uniform distribution. When we encode the prior, we can use a Gaussian as a Gaussian is the maximum entropy prior for a parameter with known expectation value and variance. Uniform priors can introduce bias and are not invariant under re-parametrisation. The aim of these single pixel simulations is to investigate an experiment's ability to constrain foregrounds. This means we want to choose a prior that will not bias the extracted parameters due to the model and will only introduce bias if there is not enough data available. This means we want to maximise the impact from the data alone on the parameter constraints. We must therefore avoid the use of priors that are known to introduce bias.

In 1946, Sir Harold Jeffrey introduced a prior that is invariant under re-parametrisation (Jeffreys (1946)). For a model with just one parameter, the Jeffreys Prior is also the objective reference prior, which maximises the relative entropy between the prior and the posterior with more than one parameter (Eriksen et al. (2008)). There are two commonly adapted methods of forming the Jeffreys Prior.

In the multivariate version, the Jeffreys prior is given by the square root of the determinant of the Fisher information matrix, $I(\theta)$. There are a number of problems that arise with the multivariate Jeffreys Prior. When using this prior, the maximum posterior estimates of the mean and standard deviation for a normal distribution give incorrect degrees of freedom, meaning that this prior has poor convergence properties. It has also been shown that the multivariate Jeffreys prior introduces significant biases into extracted, maximum posterior parameter estimates. In fact, Jeffreys himself advised against the use of this multivariate prior.

An alternative is the independence Jeffreys-rule prior. In this case, the prior is the product of the separate Jeffreys Rule priors for each parameter considered individually. For a parameter θ_i , the independent prior is given by equation 2.20. The full independence Jeffreys-rule prior is then proportional to the products of the prior. This prior achieves good convergence and is uninformative. The independence Jeffreys-rule prior can be written as:

$$\pi(\theta_i) \propto \sqrt{-E(\theta) \left[\left(\frac{\partial \log L}{\partial \theta_i} \right)^2 \right]}, \quad (2.20)$$

where $E(\theta)$ is the expectation value with respect to θ and L is the log-likelihood for a parameter, θ_i with a given distribution and the full set for parameters is:

$$\pi_{\text{Jeffreys}}(\theta_i) \propto \prod_i \pi(\theta_i). \quad (2.21)$$

In the NextBASS simulations, the free parameters that require Jeffreys priors are the synchrotron spectral index, β_S , the synchrotron curvature, C_s , the dust spectral index, β_d , the dust temperature, T_d , the free-free emission measure EM and the AME peak frequency, ν_p . The priors on the amplitudes of the models are kept as uniform distributions. Due to the complexity of the sky model, the free-free electron temperature, T_e , has been fixed to its true value. However, if convergence was more likely then a Jeffreys prior would also be appropriate. The limits on the priors are broad which means that any convergence is a true representation of the data. In Table 2.1 I list the priors used for the NextBASS simulations, these are the same priors used in Jew et al. (in prep.).

2.2.3 Simulated Observations

The pixels chosen for the simulations have come from a variety of different areas of the sky. They well represent a variety of foreground intensities in different combinations. These have been based off the pixels selected in (Jew (2017)). The pixels include a typical high signal foreground pixel taken from Planck Collaboration et al. (2016b) in the ‘Galactic Plane’, a low foregrounds pixel, ‘Off-Plane’, an AME pixel ‘Polaris Flare’, a dusty pixel ‘Near Orion’, and two foreground rich areas ‘Barnards Arc’ and ‘Lambda Orionis’. The ‘North Polar Spur’ (NPS) has also been included along with a zero foreground component which gives all foreground components an amplitude of zero. In Table 2.2 I list each component and the corresponding parameters.

In the C-BASS single pixel simulations performed in Jew et al. (in prep.), a synchrotron curvature term, C_S , was added to the synchrotron spectral model. It is possible to see how well the foreground parameters can be constrained by either correctly or incorrectly fitting the synchrotron component. In the C-BASS simulations, the first model used a straight spectral index with no curvature and the simulation was correctly fit for this. Then a free-varying curvature with the true value of $C_S = 0.0$ was added. A key investigation carried out by these simulations was setting the true curvature to be $C_S = 0.3$, but fitting for a straight spectral index with zero-curvature free-parameter.

θ	$\log(\pi(\theta))$	Limits
A_s	constant	$[0, 10^4]\text{K}_{RJ}$, $[-50, 50]\text{mK}_{RJ}$
β_s	$\frac{1}{2} \log \left(\sum_i \left(\frac{1}{\sigma_i} \frac{Y_{s,i}}{A_s} \log\left(\frac{\nu_i}{\nu_0}\right) \right)^2 \right)$	$[-4, -2]$
C_s	$\frac{1}{2} \log \left(\sum_i \left(\frac{1}{\sigma_i} \frac{Y_{s,i}}{A_s} \log^2\left(\frac{\nu_i}{\nu_0}\right) \right)^2 \right)$	$[-0.5, 0.5]$
A_d	constant	$[0, 10^4]\text{K}_{RJ}$, $[-100, 100]\mu\text{K}_{RJ}$
β_d	$\frac{1}{2} \log \left(\sum_i \left(\frac{1}{\sigma_i} \frac{Y_{d,i}}{A_d} \log\left(\frac{\nu_i}{\nu_0}\right) \right)^2 \right)$	$[0.8, 2.2]$
T_d	$\frac{1}{2} \log \left(\sum_i \left(\frac{1}{\sigma_i} \frac{Y_{d,i}}{A_d} \left[\frac{\nu_0}{1-\exp(-h\nu_0/kT_d)} - \frac{\nu_i}{1-\exp(-h\nu_i/kT_d)} \right] \frac{1}{T_d} \right)^2 \right)$	$[12, 45]\text{K}$
EM	$\frac{1}{2} \log \left(\sum_i \left(\frac{1}{\sigma_i} \frac{T_e \tau}{EM} \exp(-\tau) \right)^2 \right)$ where $\tau = f(T_e)EM$	$[0, 10^4]\text{cm}^{-6}\text{pc}$
T_e	constant	
A_{sd}	constant	$[0, 10^4]\text{K}_{RJ}$, $[-100, 100]\mu\text{K}_{RJ}$
ν_p	$\frac{1}{2} \log \left(\sum_i \left(\frac{1}{\sigma_i} \frac{Y_{sd,i}}{A_{sd}} \frac{\nu_{p0}}{\nu_p^2} \left[\frac{F'(\nu_0 \nu_{p0} / \nu_p)}{F(\nu_0 \nu_{p0} / \nu_p)} \nu_0 - \frac{F'(\nu_i \nu_{p0} / \nu_p)}{F(\nu_i \nu_{p0} / \nu_p)} \nu_i \right] \right)^2 \right)$	$[5, 70]\text{K}$
A_{CMB}	constant	$[-1, 1]\text{K}_{CMB}$

Table 2.1: The priors for each of the models used in the NextBASS simulations. The Jeffreys priors have been calculated by (Jew (2017)) for the C-BASS single pixel simulations. The top row of the limits gives the prior distribution limits for the intensity simulations, and the bottom row gives the same for the polarisation simulations. Note that the only difference is such that the polarisation amplitudes are much less than in intensity.

Component	Parameter	Galactic Plane	Off -Plane	NPS	Polaris Flare	Near Orion	Barnards Arc	Lambda Orionis	Zero Foregrounds
Synchrotron	$A_s^{(I)}$ [K _{RJ}]	47.5	5.88	39.5	14.4	11.0	22.7	16.6	0.00
	$A_s^{(B)}$	6.10	0.708	5.99	0.16	1.23	2.24	1.16	0.00
	β_s				- 3.1 -				
	C_s				- 0.3 -				
	$\nu_{s,0}^{(I)}$ [GHz]				- 0.0408 -				
	$\nu_{s,0}^{(B)}$ [GHz]				- 5 -				
Free-free	EM [cm ⁶ pc]	361	0.100	4.86	20.3	1.59	331	152	0.00
	T_e [K]				- 7000 -				
AME	$A_{AME}^{(I)}$ [μ K _{RJ}]	708	0.00	49.3	167	22.9	207	85.5	0.00
	$A_{AME}^{(B)}$ [μ K _{RJ}]	7.08	0.00	0.49	1.67	0.23	2.07	0.86	0.00
	ν_p [GHz]				- 25.0 -				
	$\nu_{AME,0}$ [GHz]				- 22.8 -				
CMB	$A_{CMB}^{(I)}$ [μ K _{RJ}]				- 75 -				
	$A_{CMB}^{(B)}$ [μ K _{RJ}]				- 0.00 -				
Thermal dust	$A_d^{(I)}$ [μ K _{RJ}]	2080	12.8	49.2	410	61.4	448	232	0.00
	$A_d^{(B)}$ [μ K _{RJ}]	44.8	0.335	3.72	2.70	0.614	9.98	1.61	0.00
	β_d	1.55	1.63	1.53	1.63	1.55	1.48	1.59	-
	T_d	17.5	24.9	21.8	18.1	21.5	21.2	19.0	-
	$\nu_{d,0}^{(I)}$ [GHz]				- 545 -				
	$\nu_{d,0}^{(B)}$ [GHz]				- 353 -				

Table 2.2: The ‘true’ parameter values used to calculate the prior distributions for each single pixel. These pixels were selected from the Planck sky maps obtaining a fair representation of the sky with each having different quantities of foreground components.

The NextBASS simulations that have included the synchrotron curvature term will have the true value set to $C_S = 0.3$. The choice of this curvature term accounts for the maximum theoretical rise in amplitude of the synchrotron spectrum at the higher frequencies, compared to the amplitudes of a straight spectral index of $\beta_S = -3.1$. In the polarisation simulations I will perform a set of three investigations on the synchrotron curvature. The first set of simulations use a straight synchrotron spectral index, which has been correctly fit to the data. The next set of simulations introduce a synchrotron model with a curvature free-parameter, this model is then correct fit to the data. Finally, I run a simulation with a synchrotron curvature term, however I fit a straight spectral index to the data.

In the polarisation simulations I will also include the free-parameters that describe a 1% AME polarisation in addition to the synchrotron curvature term. The polarisation amplitude is given by taking 1% of the intensity amplitude for that particular pixel. Using a similar method to the error modelling of the synchrotron curvature term, I will also investigate the effect an incorrect sky model can have when including a polarised AME component. To do this I will run the sky model to include the 1% AME component however fit for a polarised sky that has no AME. This will identify

whether the additional frequency channels of NextBASS can be used to determine if the foregrounds are more complicated than the simulated data suggests.

As I've already discussed, the information from one experiment alone will not be able to constrain sky models. As a result, we need to use a combination of experiments to get a meaningful result. For intensity, we know that pre-existing experiments have demonstrated that they have excellent Temperature sensitivity. This means that I will use *WMAP*, *Planck* and HASLAM published sensitivities to generate the simulated data. However, for the polarisation simulations, the current experiments do not have the power to measure the sky contribution due to thermal dust. This is due to a limitation in higher frequency data where we would obtain the highest dust amplitudes. To account for this, I have included a *LiteBIRD*-like experiment that will add the high frequency channels at the proposed *LiteBIRD* sensitivities in order to control the behaviour of the dust in the simulated data. This will also show how the combination of *LiteBIRD* and its predecessor, *Planck* perform in the generation of these simulated data sets.

The X-BASS and NextBASS sensitivities have been obtained through detailed modelling of the receiver temperature including atmospheric effects arising from the 23 GHz water line. This is fully covered and explained in Chapter 6. In Table 2.3, I list the sensitivities of all experiments used in these simulations. The values given for *LiteBIRD* come from a recent forecasting paper by [Remazeilles and Chluba \(2018\)](#). The choice of pixel size for the simulations has been chosen to obtain a higher sensitivity in polarisation by selecting a three degree pixel size. The intensity simulations are kept in units of $1\mu\text{Kdeg}$.

The different combinations of experiments can be seen in Table 2.4. These are based off the simulations done by the C-BASS Collaboration ([Jew \(2017\)](#), [Jew et al. \(in prep.\)](#)). For the Intensity simulations, I have included Haslam, *WMAP*, *Planck* and C-BASS as the pre-existing experiments. I then add in X-BASS and NextBASS in succession to show how adding more frequency channels improves constraints. For polarisation the only data available is from *Planck* and C-BASS, along with the modelling of a '*LiteBIRD*-like' experiment, the sensitivities of which have been taken from ([Suzuki et al. \(2018\)](#)) and then scaled to a 3° pixel.

The simulated data is then generated by using the proposed sensitivities scaled to either a 1° beam in intensity, of a 3° beam in polarisation. This represents a noise error in the model, such that the simulated data is given a sensitivity constraint based on the performance of an instrument. I have excluded CLASS and QUIJOTE in these pixel simulations as they cover different angular scales to X-BASS and NextBASS. In

Experiment	ν [GHz]	σ_I [μ K deg]	σ_P [μ K deg]
C-BASS	5.0	73.0	73.0
Haslam	0.408	2.5×10^6	-
<i>WMAP</i> K ^c	23	5.82	-
<i>WMAP</i> Ka	33	4.18	-
<i>WMAP</i> Q	41	3.52	-
<i>WMAP</i> V	61	3.79	-
<i>WMAP</i> W	95	3.92	-
<i>Planck</i> 30 ^d	28.4	2.45	3.30
<i>Planck</i> 44	44.1	2.57	3.9
<i>Planck</i> 70	70.4	3.08	4.5
<i>Planck</i> 100	100	1.00	1.53
<i>Planck</i> 143	143	0.333	0.72
<i>Planck</i> 217	217	0.261	0.60
<i>Planck</i> 353	353	0.198	0.57
<i>Planck</i> 545	545	0.0855	-
<i>Planck</i> 857	857	0.0319	-
<i>LiteBIRD</i> 40 ^e	40	-	0.613
<i>LiteBIRD</i> 50	50	-	0.393
<i>LiteBIRD</i> 60	60	-	0.325
<i>LiteBIRD</i> 68	68	-	0.265
<i>LiteBIRD</i> 78	78	-	0.222
<i>LiteBIRD</i> 89	89	-	0.192
<i>LiteBIRD</i> 100	100	-	0.150
<i>LiteBIRD</i> 119	119	-	0.125
<i>LiteBIRD</i> 140	140	-	0.0967
<i>LiteBIRD</i> 166	166	-	0.105
<i>LiteBIRD</i> 195	195	-	0.0950
<i>LiteBIRD</i> 235	235	-	0.125
<i>LiteBIRD</i> 280	280	-	0.217
<i>LiteBIRD</i> 337	337	-	0.318
<i>LiteBIRD</i> 402	402	-	0.615

Experiment	ν [GHz]	σ_I [μ K deg]	σ_P [μ K deg]
X-BASS	7.4	12.11	4.04
X-BASS	8.3	11.46	3.12
X-BASS	9.4	11.06	3.69
X-BASS	10.1	11.06	3.69
X-BASS	11.7	10.72	3.57
X-BASS	13.2	10.45	3.48
X-BASS	14.9	11.10	3.70
NextBASS ¹	7.4	11.39	3.80
NextBASS ¹	8.3	10.68	3.56
NextBASS ¹	9.4	10.24	3.41
NextBASS ¹	10.1	10.23	3.41
NextBASS ¹	11.7	9.79	3.26
NextBASS ¹	13.2	9.35	3.12
NextBASS ¹	14.9	9.84	3.28
NextBASS ²	15.9	2.21	0.74
NextBASS ²	17.9	1.93	0.64
NextBASS ²	20.1	1.71	0.65
NextBASS ²	22.5	1.91	0.64
NextBASS ²	28.4	1.81	0.60

Table 2.3: The instrument sensitivities for all experiments that are not X-BASS and NextBASS have been obtained from the sources below. These are based on Table 1 in (Jew et al. (in prep.)). The sensitivities of X-BASS and NextBASS have been calculated in Chapter 6 of this thesis. Note that the 3° pixels are to improve the signal-to-noise ratio of the polarisation simulations, without losing too much resolution in practice.

^a (Jones et al. (2018))

^b 10% of median antenna temperature and 0.6mK added in quadrature

^c Bennett et al. (2013)

^d *Planck* Collaboration (Planck Collaboration et al. (2016b))

^e Table 2 of Remazeilles and Chluba (2018)

these simulations, I have successfully demonstrated that the addition of the frequency channels improves constraints on the CMB amplitude and the other foreground parameters.

Each of these experiments have been used in different combinations for intensity, I , and polar intensity, P , such that each simulation will add in more frequency data points. In intensity, the fewest number of frequency data points will be attained from the combination of Haslam, *WMAP* and *Planck*, and in polarisation it will be *Planck* and *LiteBIRD* alone. I will perform separate simulations with the addition in succession of the C-BASS 5 GHz data point, X-BASS and then the NextBASS data. In Table 2.4, I outline the nomenclature I have used to refer to these map combinations. The table also specifies the key that indicates each map combination in the posterior density parameter plots. ²

Naming System	Map combination	Key
I^{hwp}	Haslam + <i>WMAP</i> + <i>Planck</i>	<i>purple dotted</i>
I^{hwpc}	Haslam + <i>WMAP</i> + <i>Planck</i> + C-BASS	<i>orange solid</i>
I^{hwpcx}	Haslam + <i>WMAP</i> + <i>Planck</i> + C-BASS + X-BASS	<i>blue dashed</i>
I^{hwpcn}	Haslam + <i>WMAP</i> + <i>Planck</i> + C-BASS + Next-BASS	<i>green dashed</i>
P^{pl}	<i>Planck</i> + <i>LiteBIRD</i>	<i>purple dotted</i>
P^{pcl}	<i>Planck</i> + C-BASS + <i>LiteBIRD</i>	<i>orange solid</i>
P^{pclx}	<i>Planck</i> + C-BASS + <i>LiteBIRD</i> + X-BASS	<i>blue dashed</i>
P^{pcln}	<i>Planck</i> + C-BASS + <i>LiteBIRD</i> + NextBASS	<i>green dashed</i>

Table 2.4: Map combinations and appropriate key for posterior density function. The nomenclature helps describe how the addition of more data points in succession can affect the overall constraints on the extracted foreground parameters.

The intensity simulations have two levels of complexity. The basic model, the results of which I will refer to as the I_B simulation results, has a straight synchrotron component, free-free, AME and thermal dust. Then another model with a curved synchrotron component, later referred to as the I_S simulations, AME, free-free and thermal dust. I found that the curved synchrotron model was not converging so I used tighter limits on the priors for the dust temperature, the dust spectral index,

²The X-BASS and NextBASS sensitivities have been calculated in depth in Chapter 6. They are based on the generation of the system temperature through an atmospheric model and propagated through a receiver model. The exposure times have been scaled appropriately sample the sky over a 5 year period, assuming a Gaussian solid angle for each beam.

Name	Sky Model Parameters	Data Fitting Parameters
I_B	$T_e = 7000\text{K}$ (fixed), $A_{CMB}, A_s, \beta_s, EM, A_{ame}, \nu_{peak}, A_d, \beta_d, T_d$	$A_{CMB}, A_s, \beta_s, EM, A_{ame}, \nu_{peak}, A_d, \beta_d, T_d$
I_S	$T_e = 7000\text{K}$ (fixed), $A_{CMB}, A_s, \beta_s, C_s, EM, A_{ame}, \nu_{peak}, A_d, \beta_d, T_d$	$A_{CMB}, A_s, \beta_s, C_s, EM, A_{ame}, \nu_{peak}, A_d, \beta_d, T_d$
P_B	$A_{CMB}, A_s, \beta_s, A_d, \beta_d, T_d$	$A_{CMB}, A_s, \beta_s, A_d, \beta_d, T_d$
P_S	$A_{CMB}, A_s, \beta_s, C_s, A_d, \beta_d, T_d$	$A_{CMB}, A_s, \beta_s, C_s, A_d, \beta_d, T_d$
P_C	$A_{CMB}, A_s, \beta_s, C_s, A_{ame}, \nu_{peak}, A_d, \beta_d, T_d$	$A_{CMB}, A_s, \beta_s, C_s, A_{ame}, \nu_{peak}, A_d, \beta_d, T_d$
P_S^{err}	$C_S = 0.3$ (fixed) $A_{CMB}, A_s, \beta_s, A_d, \beta_d, T_d$	$A_{CMB}, A_s, \beta_s, A_d, \beta_d, T_d$
P_C^{err}	$A_{ame} = 1\%$ (fixed), $\nu_{peak} = 25\text{ GHz}$ (fixed), $A_{CMB}, A_s, \beta_s, C_s, A_d, \beta_d, T_d$	$A_{CMB}, A_s, \beta_s, C_s, A_d, \beta_d, T_d$

Table 2.5: This table shows the names given to the NextBASS simulations and the parameters that were used to generate the sky model and the parameters used in the data fitting process. The fixed parameters have been specified, all other parameters were free varying using the priors listed in Table 2.1.

the synchrotron spectral index and the synchrotron curvature. [$(T_d: 12\text{--}45[\text{K}])$, $(\beta_d: 1.3\text{--}1.7)$, $(\beta_s: -2.9\text{--}-3.3)$, $(C_s: 0 - 0.5)$].

The lack of convergence is due to the increased number of parameters with respect to the number of data points. In the synchrotron curvature simulations in intensity, there are 11 free-parameters. If we consider HASLAM, *WMAP*, *Planck* and C-BASS, then there are only 16 data points. Since the number of data points is similar to the number of parameters we are trying to fit, then the model can find it difficult to converge. Even with the extra 12 channels provided by NextBASS, the model still has difficulty converging due to the small error bars of NextBASS compared to the other experiments. This could be solved by using a different Monte Carlo technique, which I discuss at the end of this chapter.

The polarisation simulations have three levels of complexity. Starting with a basic model, referred to as the P_B simulations, using a straight synchrotron curvature and thermal dust as the only foreground components. In the P_S simulations, a curvature term is added and finally in a complex polarised sky model, I call the P_C simulations, I add a 1% polarised AME emission. In the P_C simulations, when fitting the same model as used to generate the sky model back to the data, I found that the polarised AME is only well constrained in the Galactic Plane pixel.

I have included two modelling error scenarios, both in polarisation. The first I run the P_S model with the curved synchrotron however fit for a straight spectral index (to be consistent with the previous nomenclature, let us refer to this set of simulations as P_S^{err}). The second set of modelling errors I have investigated use the same sky model as outlined for the P_C simulations, however the data is now fit with a model that excludes any polarised AME component. This set of simulations, (let us call them the P_C^{err} simulations) show an increasing bias on the CMB amplitude and other foreground components as more frequency data are added. In Table 2.5, I have listed

the fixed and free-parameters that have been used for each simulation in both the sky model and the data fitting process.

2.3 Intensity Simulation Results

In this section I present the results of the intensity simulations carried out for the X-BASS and NextBASS simulations. These include results from a set of simulations that used a straight synchrotron spectral index, and then a set of simulations that introduced a synchrotron curvature term. I found the introduction of the additional curvature free parameter resulted in a poor convergence of the model. This is due to the limited frequency data at higher frequencies and the limitations of the MCMC method. I discuss this further in this section.

2.3.1 Straight Synchrotron

The basic model in intensity, the I_B simulations, represent a sky model with the fewest number of free parameters. There are 9 free parameters for this basic temperature model, including the synchrotron spectral index β_S , the synchrotron amplitude A_S , the free-free emission measure EM, the AME peak frequency ν_p , the AME amplitude A_{AME} , and the thermal dust spectral index, amplitude and temperature β_d , A_d and T_d respectively. There is no synchrotron curvature term included in this set of simulations; the synchrotron spectral index is straight. We expect to see that adding more frequency channels then improves the constraints on foreground parameters. For the most basic model, with the fewest number of components, we can see a simple trend that introducing more data point will improve parameter constraints even in pixels with weak foregrounds.

This can be seen in Figure 2.1, which shows the posterior outcome for each parameter in the foreground-poor ‘Off-Plane’ pixel. The blue and green dashed lines represent the addition of X-BASS and NextBASS respectively to the *Planck*, *Lite-BIRD* and C-BASS data. It is clear that in this pixel, both X-BASS and NextBASS make it possible to constrain the free-free emission measure and tighten constraints on the synchrotron spectral parameters, A_S and β_S . Furthermore, the successive addition of more data points gradually removes the bias on the synchrotron foreground parameters. The AME component is poorly constrained since in this region, the amplitude is close to zero, so the introduction of X-BASS and NextBASS shows a worse constraint on the AME peak frequency, however an improved detection of a zero amplitude AME component.

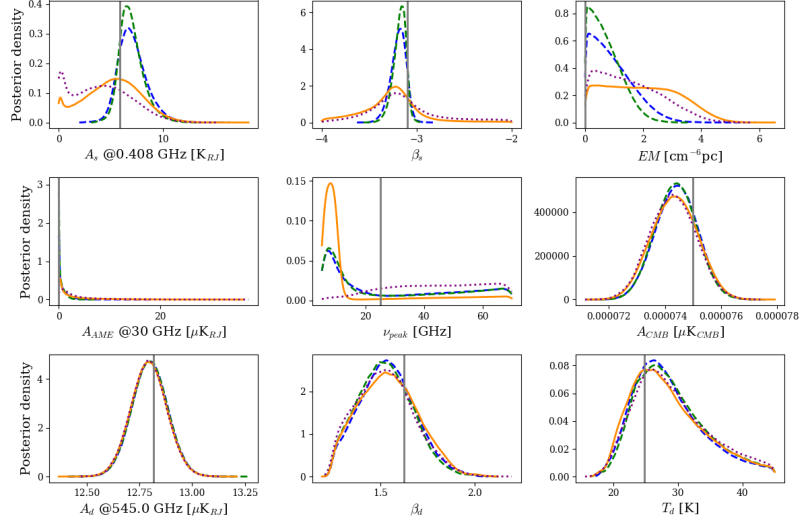
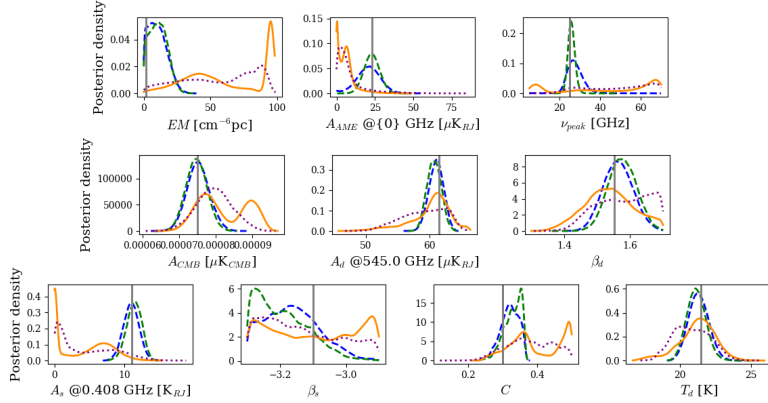


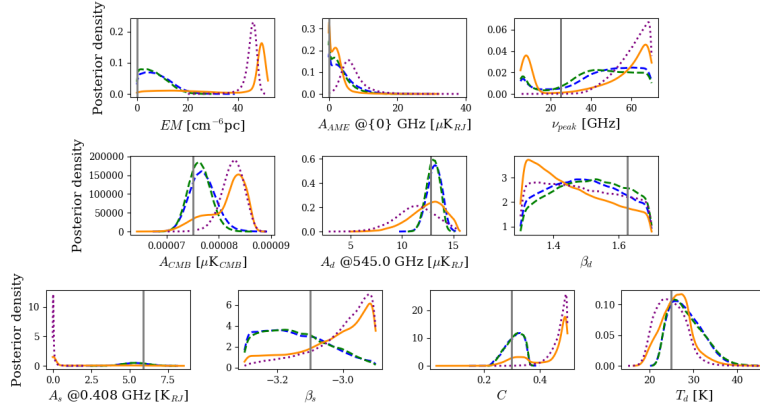
Figure 2.1: Probability density of ‘Off-Plane’ pixel for the simple intensity foreground model showing that even in foreground poor pixels, the addition of the X-BASS and NextBASS frequencies can improve constraints on the foreground parameters. (*purple dotted*: Haslam + *WMAP* + *Planck*, *orange solid*: Haslam + *WMAP* + *Planck* + C-BASS, *blue dashed*: Haslam + *WMAP* + *Planck* + C-BASS + X-BASS and *green dashed*: Haslam + *WMAP* + *Planck* + C-BASS + NextBASS.)

2.3.2 Curved Synchrotron

In the I_S simulations, I include the synchrotron curvature term as an additional free-parameter. In these simulations the curvature term, $C_s = 0.3$, I also fit for this curvature. Due to the complexity of the sky model, there are 10 free parameters. The number of data available is still limiting the convergence of the MCMC. As a result the synchrotron curvature becomes difficult to constrain. By tightening the upper and lower limits of the priors, I was able to improve the convergence on the free-free and synchrotron amplitude but the curvature and spectral index of synchrotron has still not converged. A longer MCMC chain may be required to do so. Or the addition of weakly information priors which taper at the extreme edges of the prior. This would have a similar affect to adding more data points however it can bias the parameter estimates.



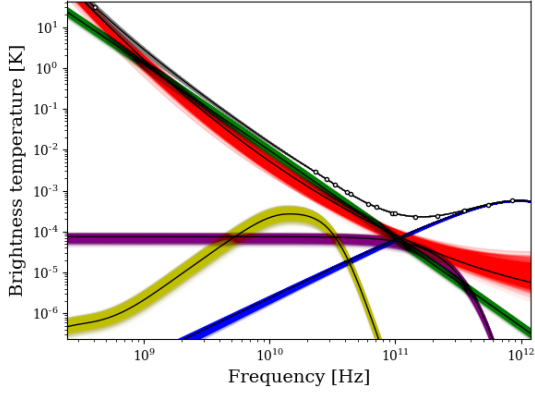
(a) Near Orion pixel: A dust rich region with weak synchrotron emission.




(b) Off-Plane pixel: A foreground poor pixel.

Figure 2.2: The results of the NextBASS, I_S simulations including a synchrotron curvature free parameter. Even with tighter limits on the priors, the chain was unable to reach convergence, however there is still sufficient improvement in constraints on the synchrotron amplitude and the CMB amplitude with the addition of X-BASS and NextBASS. (*purple dotted*: Haslam + *WMAP* + *Planck*, orange solid: Haslam + *WMAP* + *Planck* + C-BASS, blue dashed: Haslam + *WMAP* + *Planck* + C-BASS + X-BASS and green dashed: Haslam + *WMAP* + *Planck* + C-BASS + NextBASS.)


In Figure 2.2 the posterior distributions indicate that the synchrotron has not converged, however the addition of X-BASS and NextBASS has resulted in better constraints on the AME amplitude. The Off-Plane pixel has converged better but is unable to constrain the synchrotron amplitude. Finally, I demonstrate the reduction of the error bars on the foreground parameters and the CMB amplitude with the addition of increased data points in the spectrum plot 2.3.



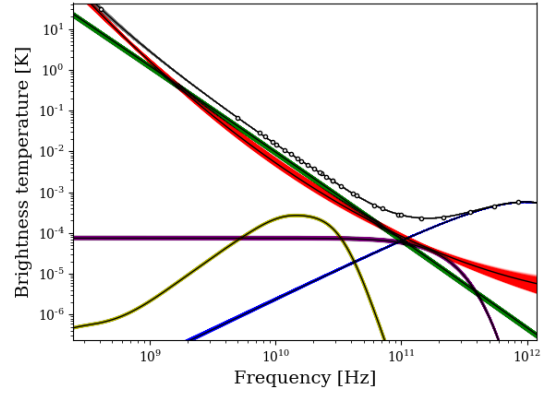
(a) HASLAM+*WMAP*+*Planck*



(b) HASLAM+*WMAP*+*Planck*+ C-BASS



(c) HASLAM+*WMAP*+*Planck*+ C-BASS
+ X-BASS



(d) HASLAM+*WMAP*+*Planck*+ C-BASS
+ NextBASS

Figure 2.3: The extracted spectrum for the Lambda Orionis pixel. The black lines indicate the true spectrum and the coloured error bars indicate the foregrounds as generated by the simulations. The circular markers indicate the frequency channels included in the fit. As we fill in the gap between Haslam and the lowest *Planck* frequency, the spectra error bars become smaller. The CMB amplitude error is also reduced. *Red=synchrotron*, *green=free-free*, *yellow=spinning dust*, *blue=thermal dust*, *purple=CMB*, *grey=total signal*.

2.4 Polarisation

In polarisation, the CMB has fewer polarised foregrounds, therefore the number of parameters in the model that describe the polarised sky is fewer than for intensity. As a result, introducing additional model complexity becomes easier, as the chains will still converge. This allows us to investigate more complicated sky models.

2.4.1 Basic Polarisation Foregrounds

The most simplistic model in polarisation only uses 6 free parameters as outlined in Table 2.5. In this model I use a straight synchrotron spectral function, and fit for a straight synchrotron spectral index, I also exclude any contributions from the anomalous microwave emission. No modelling errors are introduced in this set of simulations. This creates a baseline to compare the effects of more complicated foregrounds when adding more components to the model.


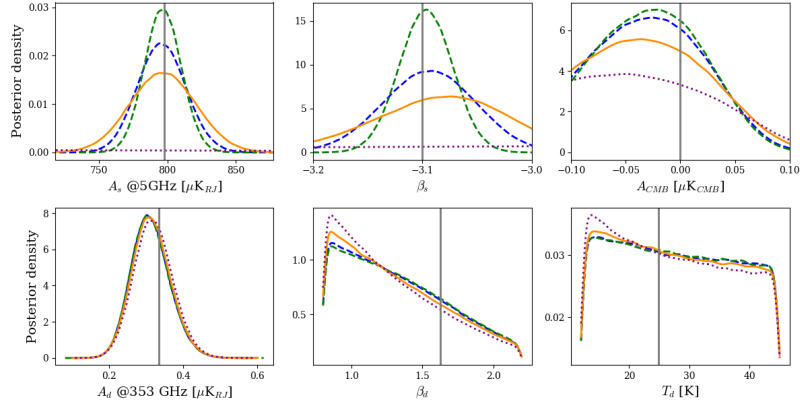
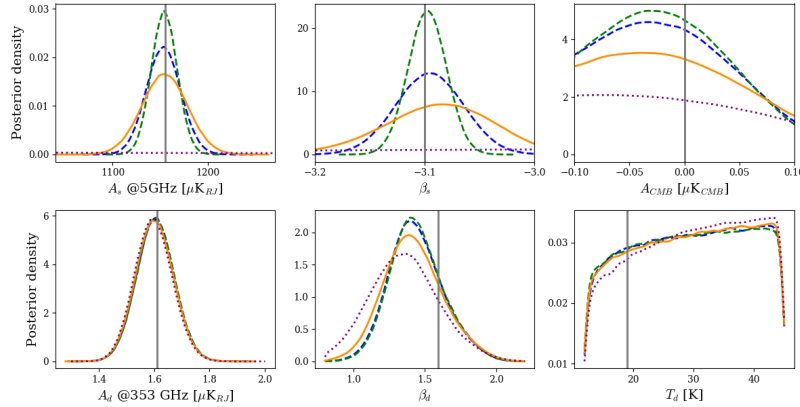


Figure 2.4: Improvement of constraints on the basic model in polarisation (P_B) for the Polaris Flare pixel. Each line in the spectrum shows a trace of the MCMC. The circular data points indicate the frequency data run for that simulation and the black lines indicate the true spectrum. The CMB has been set to zero amplitude so the extracted B -mode spectrum is a result of the detection limits of the combination of maps. The black error bars indicate the standard deviation of the noise sensitivity for each experiment. In this weak foreground pixel, *Planck* is unable to constrain any foreground signal weaker than that frequency channels absolute sensitivity. *Red=synchrotron*, *blue=thermal dust*, *purple=CMB*, *grey=total signal*.



(a) Off-Plane pixel.



(b) Barnards Loop pixel.

Figure 2.5: Posterior density distributions for the free-parameters used in the NextBASS P_B simulations. The addition of X-BASS removes the bias on the synchrotron spectral index, and the addition of NextBASS further improves constraints on the synchrotron parameters. (purple dotted: *Planck* + *LiteBIRD*, orange solid: *Planck* + *LiteBIRD* + C-BASS, blue dashed: *Planck* + *LiteBIRD* + C-BASS + X-BASS and green dashed: *Planck* + *LiteBIRD* + C-BASS + NextBASS).

The spectrum plots for the ‘Polaris Flare’ pixel is given in Figure 2.4.³ The addition of X-BASS and NextBASS reduces the error bars on the spectral components. It should also be noted that the large error bars are from *Planck* as it does not have sufficient sensitivity to detect these components in this faint foreground pixel. Figures 2.5a and 2.5b give the posterior density functions for the ‘Off-Plane’ and ‘Bernard’s Loop’ pixels respectively. For this model, the addition of more data points results in tighter error bars across all pixels. There are minimal problems with bias for the

³The MCMC process adds a trace line for each component modelled in the sky, these can be seen in the spectral plots. The larger the distribution of these trace lines, the greater the error bar on that parameter

simulations with fewer data because the model contains so few parameters. These Figures show that in a simple model of polarisation, C-BASS is necessary to constrain the synchrotron, the addition of X-BASS and NextBASS simply improve upon this.

2.4.2 Polarisation Foreground with Synchrotron Curvature

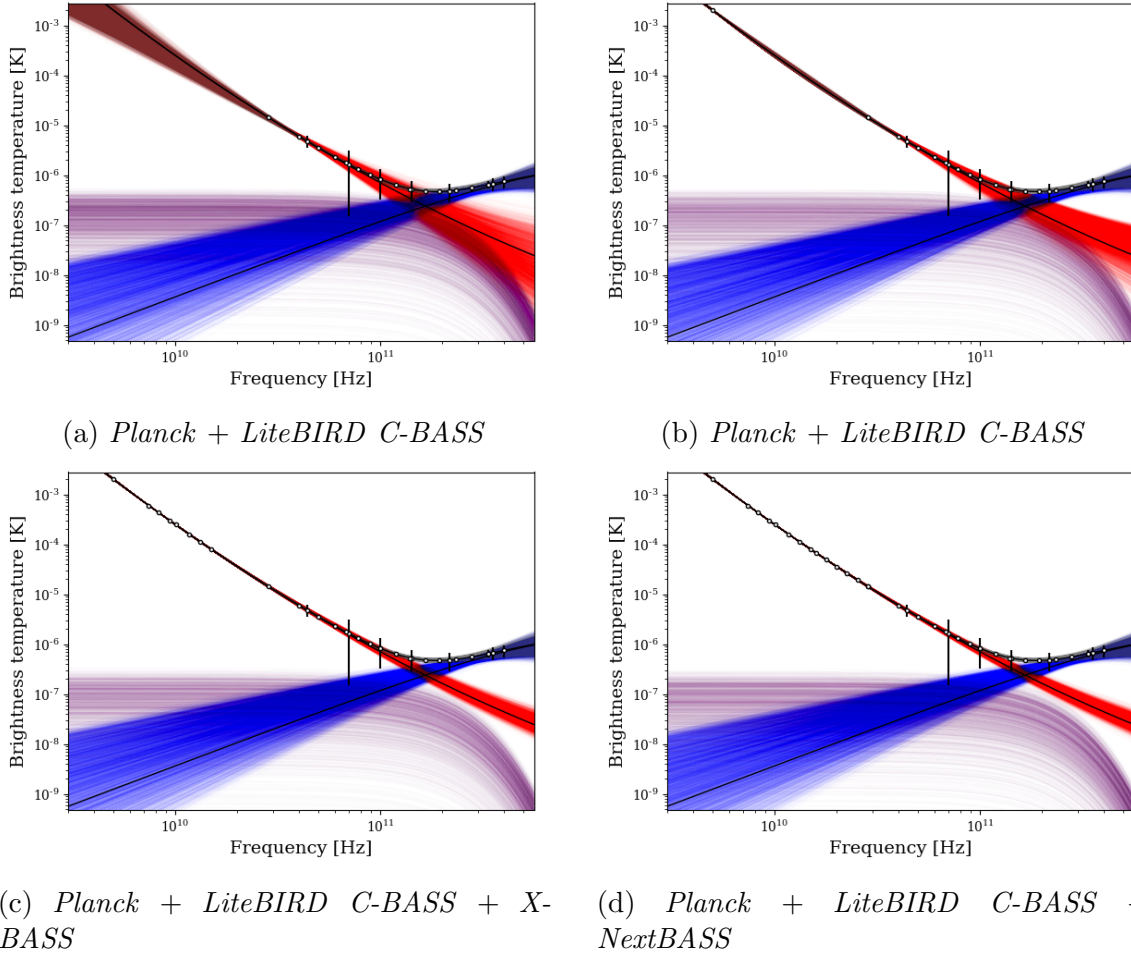
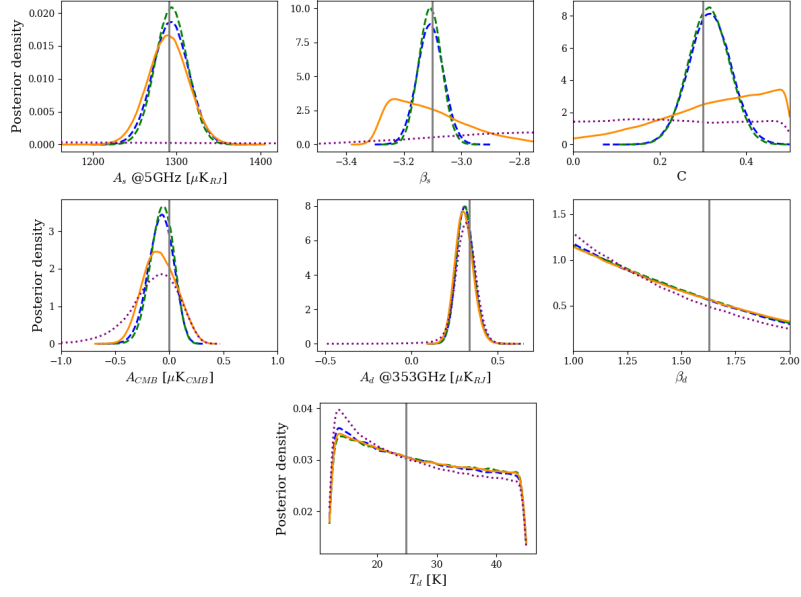
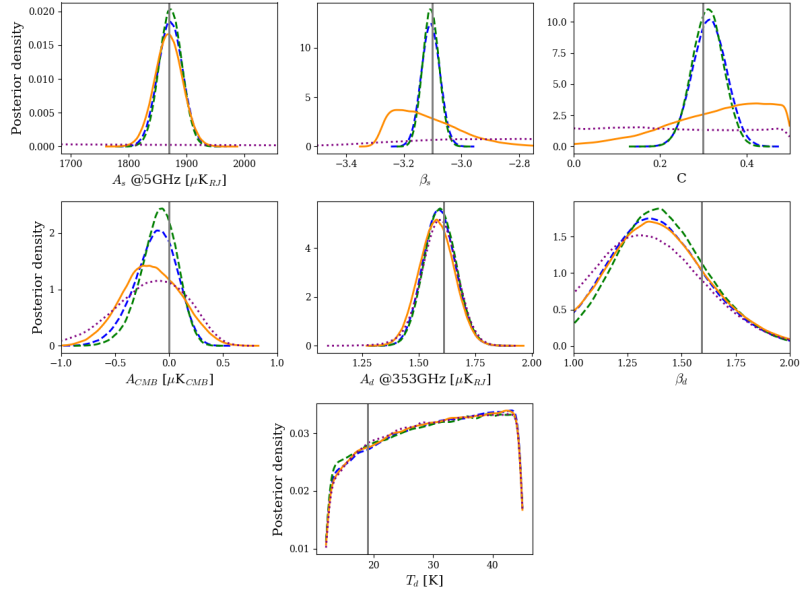


Figure 2.6: Spectra for the different combinations of maps in the Near Orion pixel, a dust rich pixel. The spectra show that with the increasing number of data points as introduced by C-BASS, X-BASS and NextBASS in succession, the error bar on the synchrotron spectral index, β_s , is reduced. The dust parameters are not well constrained by any experiment, as a result there is a bias on the dust spectral index, β_d . Red=synchrotron, blue=thermal dust, purple=CMB. (Note: the difference in x -axis for the *Planck+LiteBIRD* spectrum is due to the scaling from a lack of low-frequency data points, therefore the fitting has begun from the lowest *Planck* frequency).



(a) Off-Plane pixel



(b) Barnards Loop

Figure 2.7: Posterior density distributions for the P_S simulations that include a synchrotron curvature term of $C_S = 0.3$. X-BASS and NextBASS improved constraints on A_s, β_s and A_{CMB} , however offer little improvement on constraints on thermal dust. When a parameter cannot be constrained, the posterior tends to a straight line. (purple dotted: *Planck* + *LiteBIRD*, orange solid: *Planck* + *LiteBIRD* + C-BASS, blue dashed: *Planck* + *LiteBIRD* + C-BASS + X-BASS and green dashed: *Planck* + *LiteBIRD* + C-BASS + NextBASS).

In polarisation, I now include the free-varying synchrotron curvature component. I demonstrate that the addition of X-BASS and Next-BASS constrains C_S and reduces the bias on the CMB B -mode amplitude. We can see that the addition of NextBASS removes some of the biasing on the CMB amplitude. This is demonstrated in Figures 2.7a and 2.7b for the ‘Off-Plane’ and ‘Barnard’s Loop’ pixels respectively.

The spectrum plots for the ‘Near Orion Pixel’ (Figure 2.6) show that the addition of X-BASS reduces error bars on the synchrotron significantly from C-BASS however the addition of NextBASS does not improve on this much. It should also be noted that the low frequency foreground data do not have a significant effect on the dust constraints, suggesting that it is still poorly constrained even with the addition of *LiteBIRD*.

2.4.3 Curved Synchrotron Modelling Error

In the P_S^{err} simulations, I have used a sky model that includes a fixed synchrotron curvature term of $C_S = 0.3$ but the synchrotron curvature term is not used in the data fitting (see Table 2.5). The free parameters used in the sky model are β_S , A_S , A_d , β_d , T_d and A_{CMB} . In this modelling of errors, the differences between the foreground amplitudes of different pixels can change the bias on the extracted foreground parameters. The bias can be seen as a shift of the local maximum in the posterior density distribution plots, or as a translation of the trace frequency spectrum, compared to the true spectrum. In the pixels with stronger foregrounds, the fitting error is more obvious as the addition of X-BASS and NextBASS means the bias on becomes greater, but the error bars are reduced. For the pixels with weak foregrounds, the error bars are more poorly constrained and the bias is less obvious. As a result, the addition of X-BASS and NextBASS is essential in order to establish whether a pixel has the correct sky model fit to the data.

For a synchrotron rich pixel such as the ‘North Polar Spur’, the synchrotron is already well constrained with fewer data points and the resulting bias on the synchrotron spectral index is obvious. When the true sky model has a positive synchrotron curvature, if this term is excluded from the fitted data, then the fit compensates for this by making the spectral index more positive. In Figure 2.8, even without the lower lever point of C-BASS, there is an obvious discrepancy between the true synchrotron spectral index and what the simulation has generation. The addition of C-BASS then reduces the synchrotron error bars. The bias is obvious as the spectral index has been shifted.

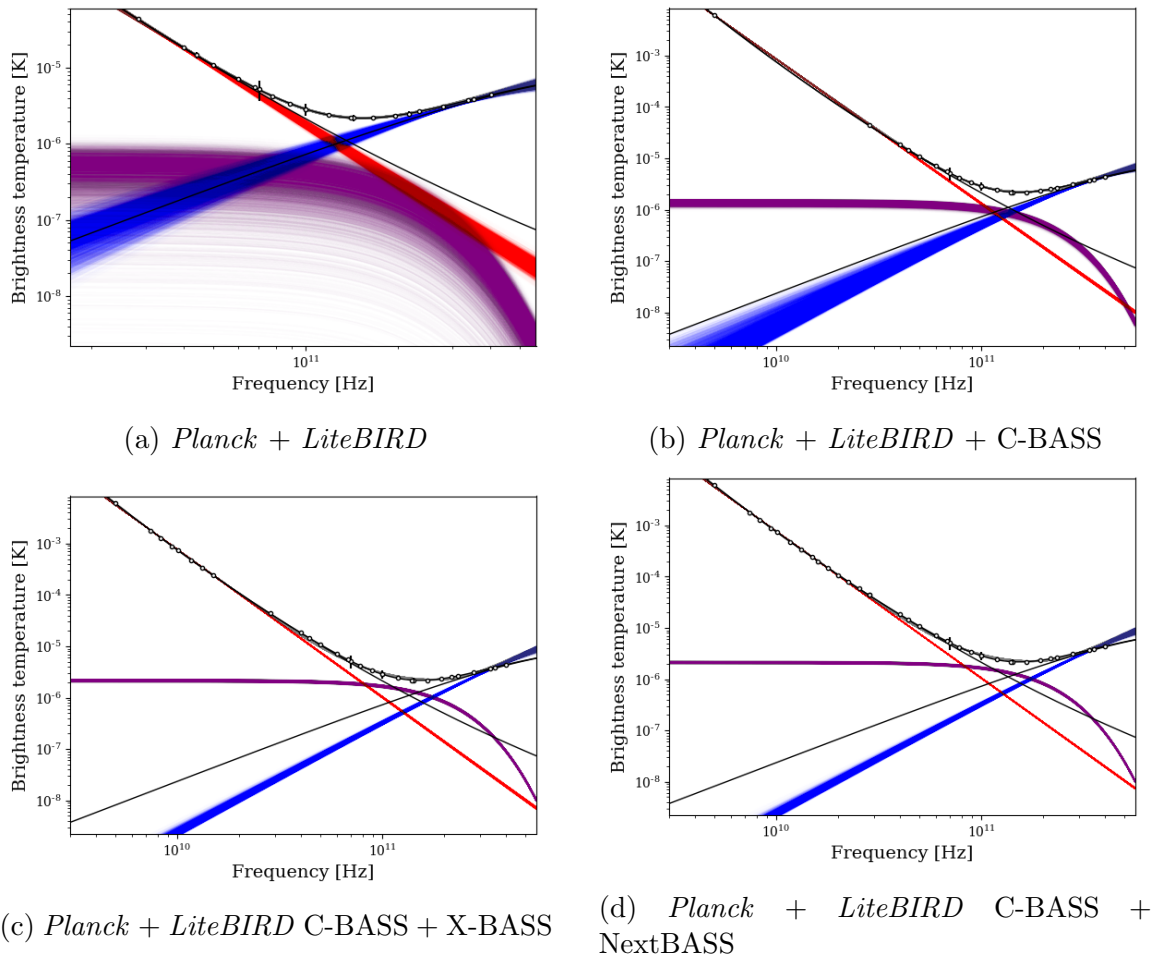
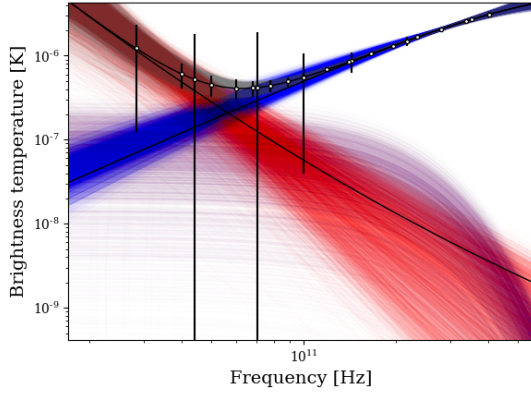
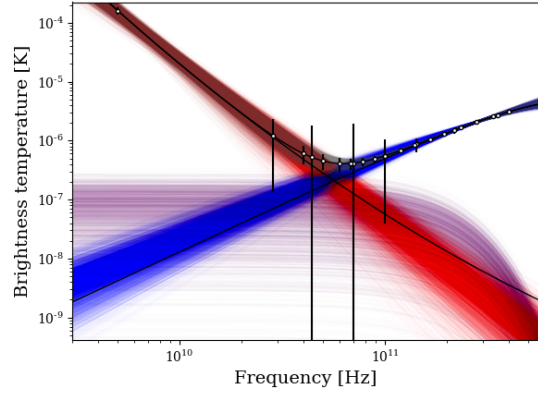



Figure 2.8: Spectra for the different combinations of maps in the North Polar Spur pixel, a pixel with strong synchrotron emission. The spectra show that by excluding the true curvature term in the fitting process, a large shift in the synchrotron spectral index can be seen. This shift is detectable with the *Planck* and *LiteBIRD* frequencies. With the addition of X-BASS and NextBASS, there is also a bias introduced on the dust spectral index. X-BASS and NextBASS make it possible to determine whether the fitting process has been carried out correctly. *Red=synchrotron, blue=thermal dust, purple=CMB, grey=total signal.* (Note: the difference in x -axis for the *Planck+LiteBIRD* spectrum is due to the scaling from a lack of low-frequency data points, therefore the fitting has begun from the lowest *Planck* frequency).




(a) *Planck + LiteBIRD*



(b) *Planck + LiteBIRD C-BASS*



(c) *Planck + LiteBIRD C-BASS + X-BASS*



(d) *Planck + LiteBIRD C-BASS + NextBASS*

Figure 2.9: Spectra for the different combinations of maps in the Polaris Flare pixel, a synchrotron poor pixel. These simulations include a fixed synchrotron curvature in the sky model of $C_s = 0.3$, however exclude curvature in the data fitting. Even with the addition of NextBASS, it is not obvious that the data has been fitted with the incorrect model. *Red=synchrotron, blue=thermal dust, purple=CMB, grey=total signal.*

In comparison to a synchrotron poor area of sky such as the ‘Polaris Flare’ pixel, the error bars in the frequency spectra are poorly constrained. In this scenario, NextBASS is better at measuring synchrotron than X-BASS. In the spectra in Figure 2.9, it is difficult to determine if there is a bias on β_s because at the NextBASS frequencies, the fitted spectrum is within the range of acceptable values for β_s .

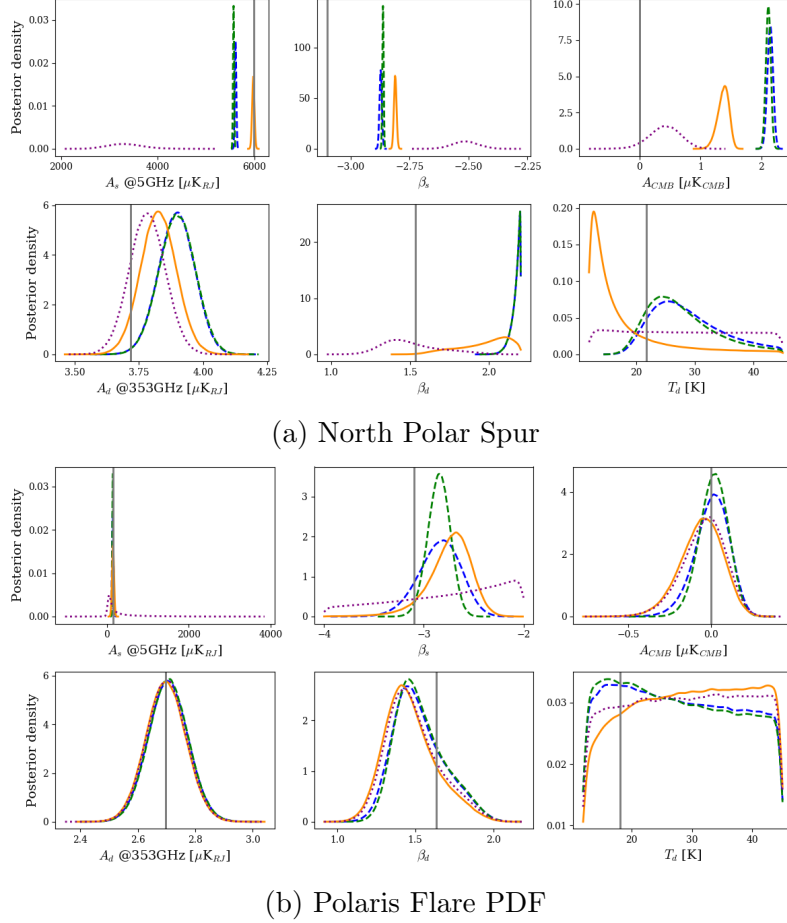


Figure 2.10: The posterior density distributions for the North Polar Spur Region and the Polaris Flare pixels respectively. The synchrotron curvature has been fixed to $C = 0.3$ in the sky model but has been excluded from the fit. For both pixels, there is a shift in β_s towards more positive values. The strong synchrotron signal in the North Polar Spur pixel means with the addition of X-BASS and NextBASS, the fitted data deviates further from realistic models. This is more difficult to recognise in weaker foreground pixels (purple dotted: *Planck* + *LiteBIRD*, orange solid: *Planck* + *LiteBIRD* + C-BASS, blue dashed: *Planck* + *LiteBIRD* + C-BASS + X-BASS and green dashed: *Planck* + *LiteBIRD* + C-BASS + NextBASS).

In Figure 2.10, I have plotted the posterior density for both of these pixels. The bias on the synchrotron parameters and the CMB amplitude is obvious in the pixel

with higher foregrounds, the North Polar Spur pixel. However, this does not hold for low foreground pixels. Identifying this in foreground rich pixels is important as an incorrect fitting procedure cannot be detected in foreground poor pixels. This is why CMB measurements are taken in foreground poor pixels, to prevent uncertainties arising from incorrect foreground modelling. However, a good CMB B -mode detection is necessary across many pixels in order to obtain an accurate measurement of the low ℓ B -mode power spectrum. Therefore, one can argue that correct foreground modelling is imperative to remove the bias on the CMB B -mode amplitude for full-sky simulations.

2.4.4 1% Polarised AME Component

In polarisation, I now include 10 free parameters as I introduce a 1% polarised AME component. The P_C simulations include polarised AME in the model and the fit. Each pixel has an AME amplitude in intensity, A_{AME}^I , I take 1% of this and use that as the true value for A_{AME}^B . I also include the AME peak frequency ν_p and the synchrotron curvature term C_S . In the case of the more complicated sky model, I explore how the addition of X-BASS and NextBASS can constrain foreground components. I also investigate whether there are considerable errors arising on the CMB B -mode polarisation due to the AME component and whether X-BASS and NextBASS can improve upon this.

Since the 1% polarised AME component is for most pixels is just below the sensitivity capability of NextBASS, the ability to constrain a model containing this component becomes reliant on whether the other foregrounds can be well constrained. In a pixel with strong foregrounds, such as the Galactic Plane, introducing the NextBASS data points significantly removes the bias on the synchrotron curvature and CMB amplitude. In this scenario, NextBASS would be able to make a detection of a 1% polarised AME in a bright region of sky.

In Figure 2.13, I give the spectrum for each of the different map combinations. When C-BASS is added to the simulation, due to the high SNR of C-BASS, there is a false detection of the AME signal. This is due to the bi-modal nature of the AME prior, which I will discuss in more detail in the summary of this section. Figure 2.12 compares relationships between posteriors in the parameter space for the map combination including NextBASS in the foreground weak, yet AME rich Polaris Flare pixel. The shape of the posterior for ν_{peak} is actually the shape of the prior, combined with the detected signal of the AME. There are three distinct distributions in the ν_{peak} vs. ν_{peak} plot of Figure 2.12, the lowest peak and the steady rise at higher

frequencies are the bi-modal prior. The peak in the middle is a weak detection of the AME by NextBASS. Depending on the constraints of the other parameters in the model specifying the accuracy of the sky model, this weak signal may be enough to classify this as a 1% polarised AME detection.

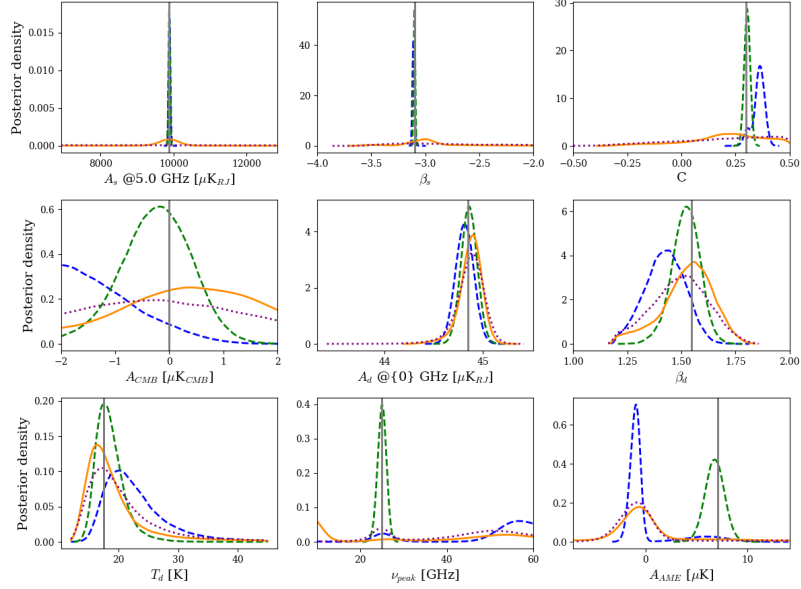


Figure 2.11: Posterior density distributions for the Galactic Plane pixel for the complex polarisation model, P_C , which includes a 1% polarised AME component. In this bright foreground pixel, NextBASS is needed to make a definitive detection of the AME polarisation and hence remove the bias on the other foreground parameters. This plot demonstrates that NextBASS is required to make a 1% polarised AME detection (purple dotted: *Planck* + *LiteBIRD*, orange solid: *Planck* + *LiteBIRD* + C-BASS, blue dashed: *Planck* + *LiteBIRD* + C-BASS + X-BASS and green dashed: *Planck* + *LiteBIRD* + C-BASS + NextBASS).

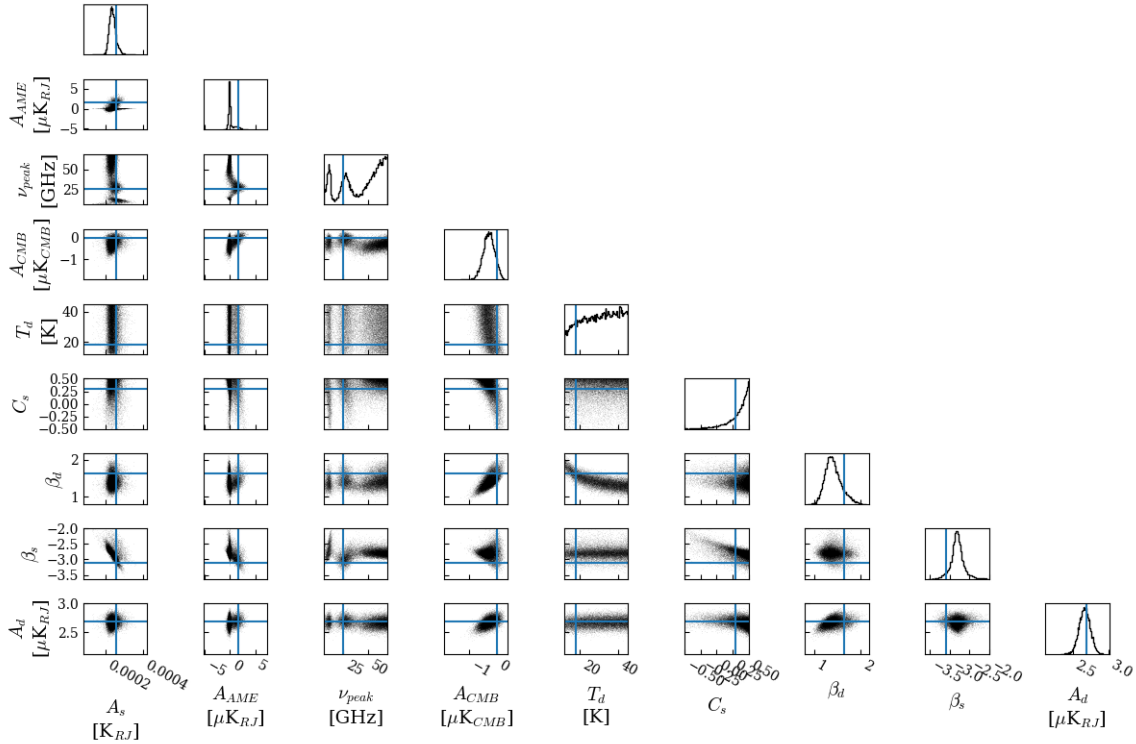

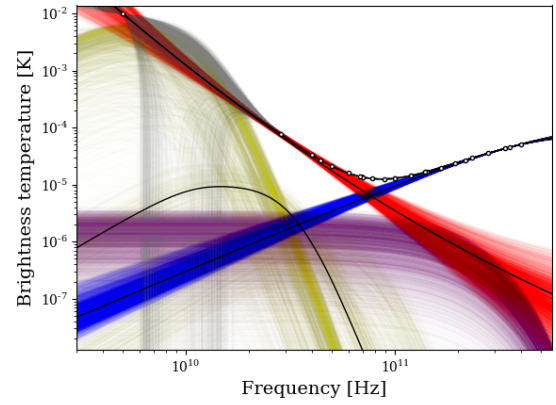


Figure 2.12: Polaris Flare pixel, a pixel with weak foregrounds but comparatively strong spinning dust emission. This simulation was for a 1% polarised AME sky model including NextBASS. This plot shows the relationships in parameter space. There is an three peak distribution of the AME peak frequency, ν_{peak} , corresponding to the bi-modal prior and a weak detection of AME made by NextBASS.


Figure 2.14 shows the resulting spectra with the inclusion of NextBASS for different pixels of differing levels of AME polarisation. These correspond to the Lambda Orionis pixel, a relatively bright foreground pixel. The Polaris Flare pixel with weak foregrounds but higher spinning dust emission, the North Polar Spur pixel with very weak spinning dust but strong synchrotron emission and the Off-Plane pixel with very weak foregrounds and negligible AME.




(a) Planck + LiteBIRD



(b) Planck + C-BASS + LiteBIRD



(c) Planck + C-BASS + LiteBIRD + X-BASS



(d) Planck + C-BASS + LiteBIRD + NextBASS

Figure 2.13: Galactic Plane pixel for the complicated polarisation model showing the improvement in error bars in the spectrum. The false detection at 5 GHz is a result of the bi-modal nature of the AME peak frequency prior. The NextBASS data points are needed to constrain the AME. Red=synchrotron, blue=thermal dust, purple=CMB, grey=total signal.

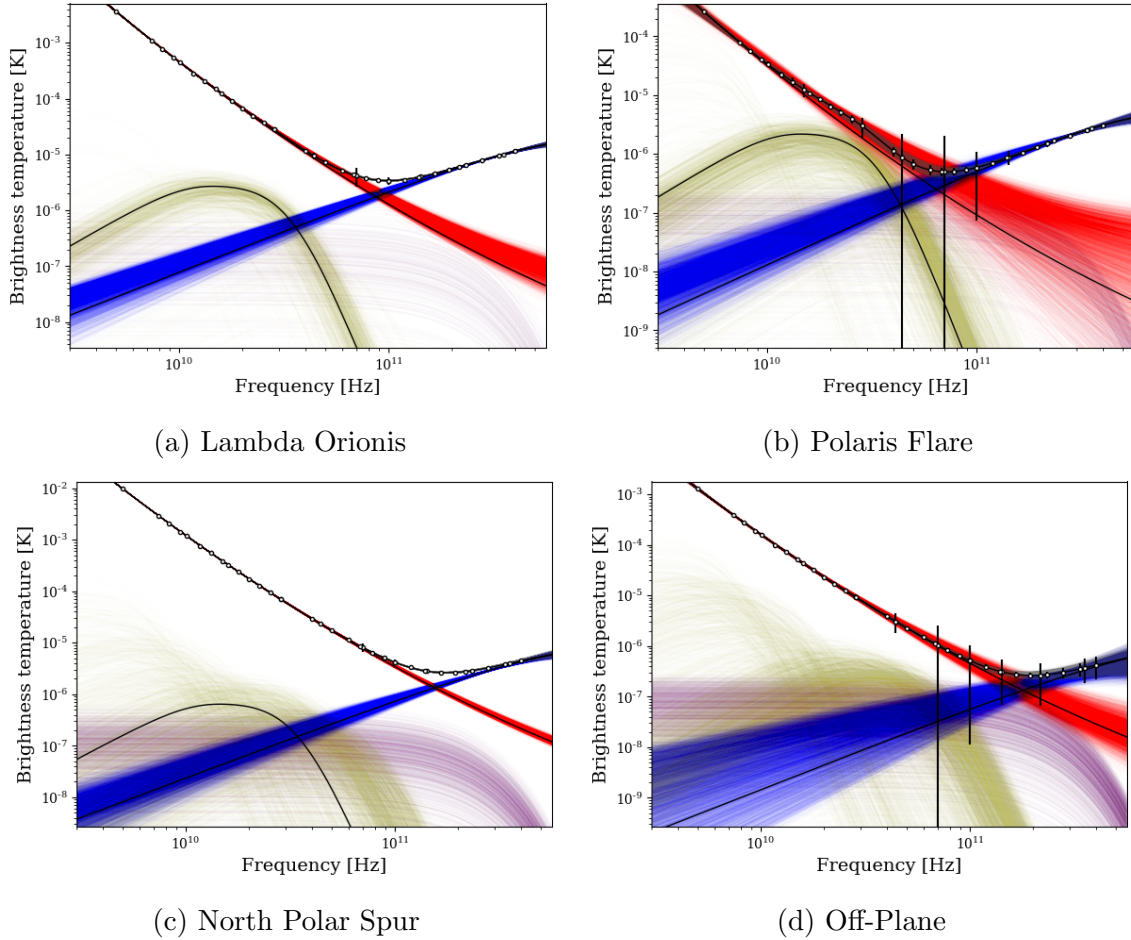



Figure 2.14: A comparison on the spectra for pixels of differing foreground intensities for the map combinations including NextBASS. These are for a polarised sky that contains a 1% AME component that has been correctly fit for. The resulting spectra show that in most pixels, NextBASS can make a very faint detection unless the AME signal is negligible, then a spurious signal is introduced. *Red=synchrotron, blue=thermal dust, purple=CMB, yellow=AME, grey=total signal.*

2.4.5 Excluding 1% AME from simulation

The final investigation I carried out is the removal of the 1% AME component in the fitting when it has been included in the sky model. The list of fixed and free-parameters used for the sky model and the data fitting for the P_C^{err} simulations can be found in Table 2.5. In the brighter pixels, we see that there is a bias introduced with the addition of more data points. Comparing the NPS and Galactic Plane pixels, we can see that due to the bright synchrotron in both, the foregrounds are well constrained. However the North Polar Spur pixel does not have significant AME to change the spectrum when an incorrect sky model is fit to the data. (See Figure

2.15). In the case of the Galactic Plane pixel, the CMB amplitude is shifted so far to a negative amplitude that it is not present in the spectrum.



(a) North Polar Spur pixel: weak AME, strong synchrotron emission (b) Galactic Plane pixel: strong foregrounds

Figure 2.15: The North Polar Spur and Galactic Plane spectra generated by the P_C^{err} NextBASS simulations that include a 1% polarised AME component in the sky model, however exclude spinning dust in the data fitting. This shows that in both pixels, due to the strong foreground signals, the synchrotron and dust spectra are well constrained. Since the AME signal in the Galactic Plane pixels is detectable by NextBASS, when no spinning dust is included in the fit, there is a noticeable shift in the synchrotron and thermal dust spectra from the true spectra. This modelling error removes the B -mode amplitude as it has become negative. This modelling error is much more difficult to detect in the North Polar Spur pixels, as the spinning dust emission is very low. *Red=synchrotron, blue=thermal dust, purple=CMB.*

In the P_C simulations, I have shown that NextBASS is required to make a detection of a polarised AME component. The P_C^{err} set of simulations investigate if excluding this component from the fitting can cause a noticeable bias on the CMB B -mode amplitude. Figure 2.16 shows the posterior density distribution for the CMB B -mode amplitude for all pixels included in these simulations. The plot shows that there is a noticeable bias in the extracted B -mode amplitude towards negative values for most pixels when NextBASS is included in the maps. This shift is considerably greater in the Galactic Plane, Polaris Flare and the Lambda Orionis pixels, which have a stronger spinning dust emission. This bias is more pronounced for NextBASS than X-BASS due to the higher channel of NextBASS which samples the majority of the spinning dust spectrum.

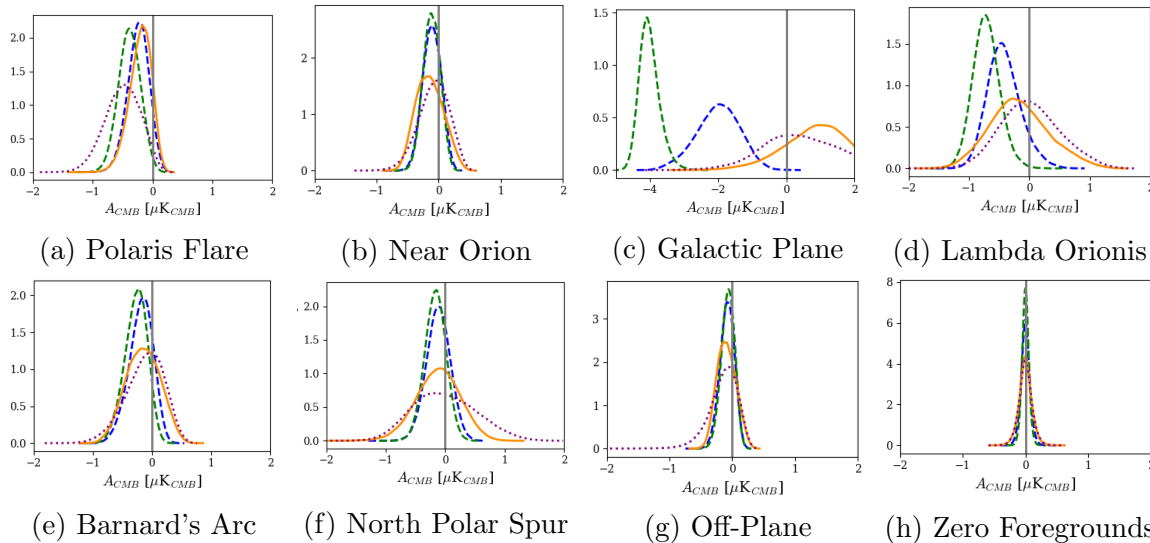


Figure 2.16: Posterior density distributions of the CMB B -mode amplitude for all sky pixels. These simulations include a 1% polarised AME component in the sky model however exclude the AME polarisation from the data fitting. As a result, there is a bias introduced on the extracted B -mode amplitude for pixels with stronger spinning dust emission. In practise, this bias can help determine whether the correct sky model has been fit to the data. Only NextBASS is capable of making this measurement on the B -mode bias, which would hence determine a 1% polarised AME detection and further improve on future CMB foreground cleaning. (purple dotted: *Planck* + *LiteBIRD*, orange solid: *Planck* + *LiteBIRD* + C-BASS, blue dashed: *Planck* + *LiteBIRD* + C-BASS + X-BASS and green dashed: *Planck* + *LiteBIRD* + C-BASS + NextBASS)

2.5 Future Work

These simulations could be improved upon in a number of ways. The two main issues that these simulations have suffered from is the inability for the intensity simulations with the synchrotron curvature term to converge, and the bi-modal properties of the AME prior. These are due to us applying non-realistic spectral properties into the parameter space. First of all, when the synchrotron intensity model includes steep curvatures and more positive spectral indices, there is a tendency for the synchrotron spectrum to curve up again at higher frequencies. We know this ‘U-like’ spectral shape is not physically possible for synchrotron emission. By repositioning the minimum of the curvature term to lie way out of the band of observable frequencies, say at 1000 GHz, the convergence of the synchrotron spectral index will be improved. This may also solve the convergence issues in the intensity simulations.

In the case of the bi-modal AME signal, we know that there is no strong emission

at 5 GHz from the C-BASS maps (Dickinson et al. (2018b)), therefore we can change the lower limit of the prior to cut off this region of frequency space. Currently the prior is set to 10 GHz and this bi-modal feature is still visible. By moving the cut-off frequency to 15 GHz, the lower peak in the bi-model prior will be removed entirely, making the AME model used in the simulations more true to the physical mechanisms that determine the spectral shape of spinning dust.

In polarisation, the amplitude of all sky components is much less. It then becomes much more difficult to constrain all of the parameters especially when there are limited data sets. This means the choice of prior becomes very important. An alternative prior to the Jeffreys prior is the reference prior. This maximises the relative entropy between the posterior distribution and the prior. This means that the data can have a maximised influence on the posterior. The main disadvantage is that reference priors for highly dimensional models are non-trivial to calculate and must be estimated numerically. This is computationally expensive, especially with minimal datasets.

Due to their computational complexity, the full reference prior has not been used in these simulations. The reference prior satisfies similar invariance principles to the Jeffreys prior and is equal to the Jeffreys' prior for single parameter models. Therefore, for the preliminary pixel component separation methods used in the NextBASS simulations, the full reference prior has not been necessary. The extra data points provided by X-BASS and NextBASS improve parameter constraints enough such that the prior distribution will not be detectable underneath the extracted parameter posterior. However, for simulations with fewer data points and more free-parameters, the reference prior would further remove any correlation between parameters and improve convergence. The reference prior method is being developed for the C-BASS simulations by Jew et al. (in prep.).

When running models with lots of parameters, the Metropolis-Hastings algorithm can take a long time to converge. An alternative already discussed could be to use the Hamiltonian Monte Carlo (HMC) methods (Girolami and Calderhead (2011)). This requires knowing the derivative of the prior for each parameter so the sampler can quickly move around the parameter space. This HMC method has been shown to work and implemented as the No-U-Turn Sampler, developed by Hoffman and Gelman (2011), can be used in PyMC3. This method is being developed for similar forecasting and foreground removal simulations by Grumitt et al. (in prep.).

Chapter 3

The Antenna Design for NextBASS

*“Rabbit’s Clever”, said Pooh thoughtfully.
“Yes”, said Piglet, “Rabbit’s clever”.
“And he has Brain”,
“Yes” said Piglet, “Rabbit has Brain.”
There was a long silence.
“I suppose”, said Pooh, “that that’s why he
never understands anything.” ”*

Benjamin Hoff, *The Tao of Pooh*

This chapter describes the optical design of the X-BASS and NextBASS antennas, to match the stringent requirements of a low-noise, polarised CMB foregrounds experiment. We require an optical system, which gives a circular beam with low cross-polarisation, and in the case of NextBASS, we also require a large focal plane area. Contributions from the far-out sidelobes must also be minimised to prevent system temperature increases, and to reduce scan-synchronous pick-up of ground spillover which would mask the sky signal.

The NextBASS optics have been purpose-designed with an unblocked aperture and a focal plane that would suit a large feed array. The two prime candidates for the optical configuration are an off-axis Gregorian (OAG) and a cross-Dragone (CXD) system. Using studies done by myself and others in the literature, I will explain my choices followed by an analytical investigation of both configurations using the advanced electromagnetic antenna modelling software, GRASP ([TICRA \(2015b\)](#)). I have also analysed the focal plane of candidate designs to assess their compatibility with a feed array. Furthermore, I discuss the design of a ground-shield for NextBASS, and how effectively the spillover can be reduced.

The X-BASS optics will use the pre-existing C-BASS South antenna, an adapted Cassegrain reflector with shaped primary and secondary mirrors. This is poorly suited to a feed array system but has excellent beam circularity and low cross-polarisation, suiting a single-feed system well. I use photogrammetric surface accuracy data of the C-BASS South dish to analyse the performance up to frequencies of 15 GHz. I perform an optical analysis of X-BASS with and without the modelled surface errors. The feed horns used in conjunction with the systems modelled in this chapter come from the final horn designs, these will be detailed in Chapter 4.

3.1 Dual Reflector Antenna Design Principles

This section is an overview of reflector antennas and the optical design principles required to develop any reflector antenna design. I used these principles to derive the different styles of reflector antennae and explore their advantages and disadvantages.

3.1.1 Telescope Design Fundamentals

A telescope takes parallel light rays from distant objects and focuses them into a single point. This can be done using a combination of focusing elements. Each element is shaped to a specific conic section. For an effective antenna design, we will consider three combinations of conic sections. The first, or only, focusing element in a system can be either a spheroid or paraboloid reflector, a second focusing element can also be added which may be either an ellipsoid or a hyperboloid reflector. Therefore, the three scenarios we can consider are a spheroid or paraboloid on its own, a spheroid or paraboloid plus an ellipsoid and thirdly, a spheroid or paraboloid plus a hyperboloid.

A conic section is a particular curve generated by taking a plane cut through a cone. Each mirror in a classical optical system is generated by a 3-D rotation of a conic section. A conic section can be described by an eccentricity, e and parameters to set the physical scale of the mirror. The parameters of the various conic sections are given in Table 3.1. Both the hyperbola and ellipse have two foci, where the parabola and circle have only one. Note that the circle is a special case of the ellipse, with the two foci coincident.

Conic Section	Typical Value	Eccentricity	Equation
circle	$e = 0$	0	$x^2 + y^2 = r^2$
ellipse	$0 < e < 1$	$\sqrt{1 - \frac{b^2}{a^2}}$	$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$
parabola	$e = 1$	1	$y = ax^2$
hyperbola	$e > 1$	$\sqrt{1 + \frac{b^2}{a^2}}$	$\frac{x^2}{a^2} - \frac{y^2}{b^2} = 1$

Table 3.1: Table of conic sections and their corresponding equations. Where a corresponds to the semi-major axis and b is the semi-minor axes.

The most simple reflector, known as a prime focus or front-fed antenna, is a focusing element which takes the rays and brings them to a single focal point where the feed is placed. A spheroid can be used, however these reflectors can be susceptible to spherical aberration at radio wavelengths (Qiu (1998)). These aberrations would be visible in the point-spread function (PSF) of the beam, as a loss of circularity. The PSF is an image of the aperture in Fourier space, and indicates the level of diffraction that the reflector at a given wavelength is susceptible to. Paraboloid reflectors do not suffer spherical aberrations, making them a better choice for any reflector (Hecht (1997)). To reduce the loss of rays at large angles, a secondary reflector can be added into the system. This is usually an ellipsoid or hyperboloid. Each of these has two foci. In a dual reflector system one focus is aligned with the prime focus of the paraboloid, rays then converge at the focal point of the secondary reflector, which determines the feed position.

On-axis designs align the primary and secondary axes, in off-axis systems this alignment is offset by some angle. Figure 3.1 shows the different combinations of conic sections to create some typical on- and off-axis reflector systems along with the ray diagrams. The angle Ψ gives the angle subtended between the primary and secondary axes for the off-axis systems.

As previously demonstrated, the most basic antenna is the axial, single-feed system. The wide angle feed required to fully illuminate the primary aperture will cause large signal pick-up from the ground. This can result in a higher system temperature. These systems work well for purposes that only require a single feed with a simple antenna architecture and are commonly used in satellite downlink. The Lovell telescope is a famous example of a prime-focus antenna and is used for observing pulsars and detecting radio transients (Michilli et al. (2018), Bhattacharyya et al. (2018)). However, for the applications we are considering, we need a very low noise instrument so a dual-reflector antenna will better meet these requirements.

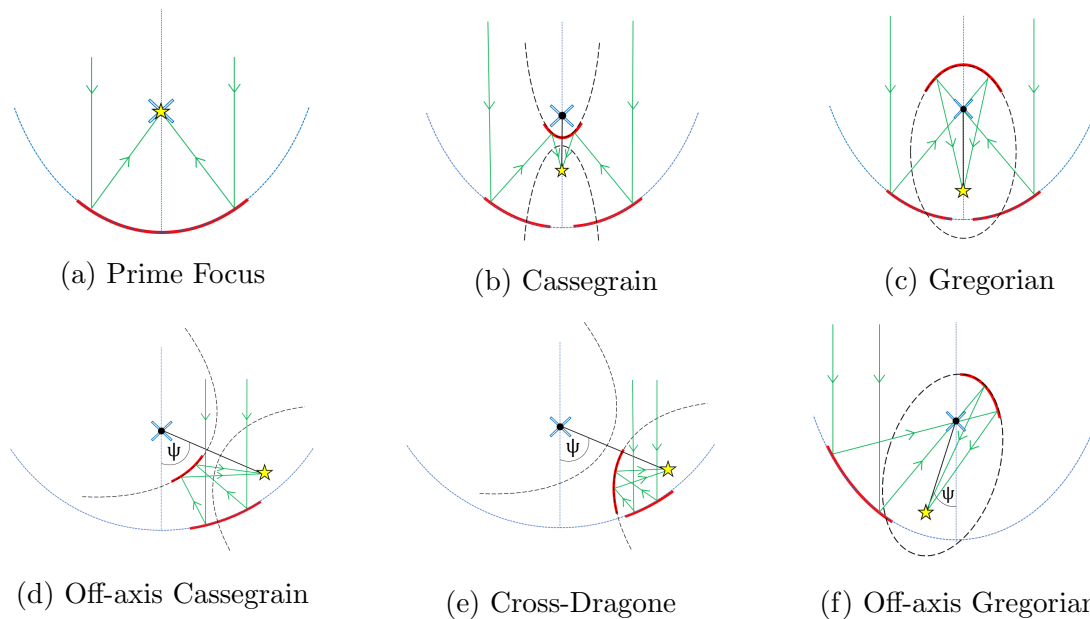


Figure 3.1: Typical reflector antennae designs based on the conic sections. Paraboloids are represented by the blue curve where the blue cross indicates the primary reflector focus and the dotted line is the vertex and primary axis of the system. The black dotted curve represents the secondary reflector and the black line joining the two foci is the secondary axis. The focal point of the dual reflector system is shown by the star. The red curves show the useful segments of the curves that make the reflectors.

3.1.2 On-axis Systems

The most common dual-reflector system is the Cassegrain. This has a paraboloid primary and a convex hyperboloid secondary. This second reflector allows the feed to be placed facing the direction of the antenna main beam. Also the backend can be located in a convenient receiver cabin beneath the primary. A Cassegrain system can often be shaped to reduce the effects of ground-spillover (as we will see with the C-BASS South optical configuration). A Cassegrain system has a much larger focal length than the parent primary, meaning a higher f -number and is thus fed by a higher-gain feed horn.

Another common on-axis system is the Gregorian reflector antenna, which also uses a parabolic mirror as the primary but has a concave ellipsoid as the secondary. In a Gregorian system, the feed can be located either below or above the primary. Using a concave mirror means greater aperture efficiencies can be achieved than with a classical Cassegrain design (Balanis (2009)).

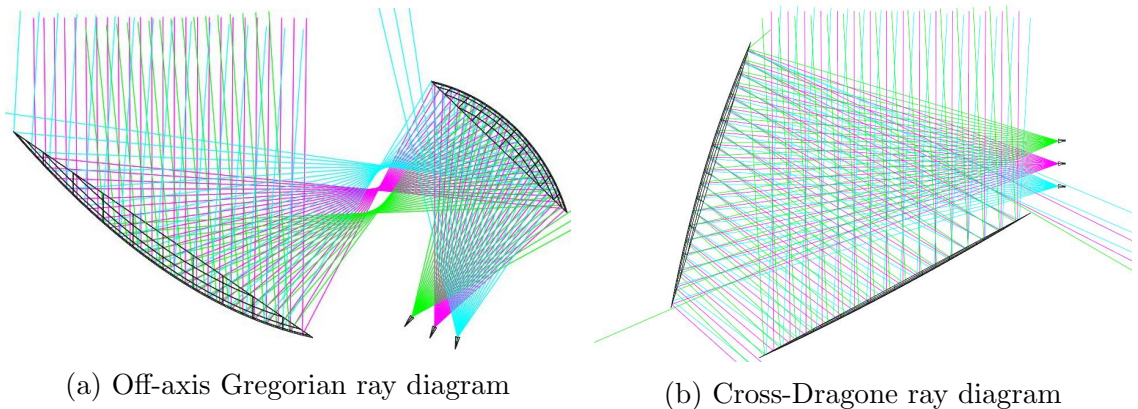


Figure 3.2: Ray diagrams of the two reflector antennae tested for the NextBASS optics. Note that there is some curvature in the position of the focal points in the OAG configuration, whereas the CXD has a flat focal plane.

3.1.3 Off-axis Systems

Having the feed system at the primary mirror centre is not always ideal as the blockages and struts required can cause a loss in aperture efficiency, a rise in cross-polarisation and introduce unwanted features in the point spread function of the beam. By making the mirrors off-axis, it is possible to avoid these blockages.

For an off-axis configuration of a classical Cassegrain, the secondary mirror can still be hyperbolic and convex but both mirrors are now used off-axis. The cross-Dragone (CXD) optical configuration is a special case of an off-axis Cassegrain, such that the secondary mirror is now concave instead of convex. This is achieved by using a segment of the negative side of the hyperboloid. Another example is the off-axis Gregorian (OAG), which uses both paraboloid and ellipsoid mirrors with the secondary still concave as in the on-axial design. For these systems, we can consider the coordinate system to be at the centre of the parabolic vertex of the primary. We define the system by some angle, Ψ , which is the angle between the primary and secondary axes (see Figure 3.1). In the OAG system, the angle, Ψ between the primary and secondary axes is much smaller than in the CXD system. These off-axis systems open up space in the focal plane making them good choices for systems with a focal-plane array. Sometimes a tertiary mirror can be added to obtain better positioning of the feed system.

Both the OAG and CXD are popular choices for new CMB experiments. Often, for the space-based experiments, the OAG is favoured due to the feed lying in the same plane as the mirror, making the receiver system support part of the primary structure. For example, both the *Planck* and *WMAP* optics were based on off-axis

Gregorians. *WMAP* has two back-to-back shaped Gregorian reflectors, with the secondary and primary axes separated by 141° to allow differential measurements over large angles (Page (2004)). The *Planck* optics obey the Mizuguchi condition (I discuss this condition further in Section 3.1.6) to preserve polarisation purity and the shaped primary mirror is an off-axis ellipsoid of eccentricity $e = -0.87$ (Tauber et al. (2010)).

However, as experiments are increasing in sensitivity, they require more feeds and hence the use of CXD's is more favoured due to its larger and more uniform focal plane. The *LiteBIRD* satellite will use a Mizuguchi-cross-Dragone design (Matsumura et al. (2014)), enabling up to 2022 bolometers in the focal-plane. Other experiments using similar designs are ACTPol, a 6m OAG with ≈ 3000 transition edge sensor bolometers (Thornton et al. (2016)), CLASS which also uses an OAG with a tertiary mirror to change the axial position of the receiver (Abazajian et al. (2016)). EBEX uses a 1.5m Gregorian-Mizuguchi Dragone with lenses and dichroic filters that separate and couple the beam to both the 150 GHz and 220-450 GHz bolometer focal-planes (The EBEX Collaboration et al. (2010)). Finally, similar systems to that could help with the NextBASS optical design are the QUIET optics, the Simons Array and QUIJOTE. QUIET used a classical side-fed Dragone-Cassegrain antenna feeding into a cryostat containing an array of 90 GHz feed horns (Bischoff (2013)). The Simons Array, a 6m CXD with a focal-plane of 22,000 polarisation-sensitive detectors (Abazajian et al. (2016)). QUIJOTE also uses a 2.25m cross-Dragone antenna, the focal plane consists of 8 horns, 2 for each frequency at 11, 13, 17 and 19 GHz (Watson et al. (2010)).

3.1.4 Primary Mirror

We will assume that all primary reflectors will be parabolic. The equation for a parabola is given in Table 3.1. The focal length of a parabola is related to the constant a and is given by $f_p = 1/4a$. Since the primary mirror is the first element in the optical system, it defines the amount of collecting area available, and hence gives an indication of the total collecting power. The ratio of focal length to the primary mirror diameter is known as the ‘ f -number’ given by equation 3.1. This gives an indication of the telescope speed, slower systems (greater than $f/6$) will suffer less from aberrations however will be larger in size.

$$f = \frac{F}{D} = \frac{f_p}{D_p}. \quad (3.1)$$

3.1.5 Secondary Mirror




Figure 3.3: Diagrams of the defining parameters of an ellipse and a hyperbola and the corresponding lengths between the focal points and some point on the curve.

As shown in the section outlining types of reflectors, the secondary mirror is usually either an ellipsoid or a hyperboloid, the equations of which we have seen in Table 3.1. It is clear from Figure 3.3, that the focal length is given by $f_s = 2c$, or the distance between two foci. For off-axis systems, we must consider the distances from some point on the curve, P to both foci. This in practice, tells us where to position the feed relative to the sub-reflector. In an ellipse, we can define two lines from a point on the surface, to either of the two foci. This also applies to a hyperbola. The two resulting vectors are R_1 and R_2 . We can assume the vector that describes the line between the centre of the secondary reflector and the position of the feed, is given by R_2 , which can be defined in polar coordinates from the origin by equation as (TICRA (2015b)):

$$R_2 = \sqrt{\frac{a^2 b^2}{b^2 \cos^2 \phi - a^2 \sin^2 \phi}} \quad (3.2)$$

The eccentricity can be used to calculate the magnification M_0 , a beam will undergo when reflected by a conic section. Since the primary mirror is a parabola, the magnification is caused by the elliptical properties of the secondary. This magnification is given by:

$$M_0 = \frac{e + 1}{e - 1}. \quad (3.3)$$

3.1.6 Mizuguchi Condition

The design of off-axis reflectors for NextBASS has been modelled in the physical optics analysis package GRASP, making use of the dual reflector wizard. This allows

the user to define the f -number, focal length of the primary and subtended angles between the two mirror axes, and the angular positioning of the feed. It is also possible to specify whether the Mizuguchi condition is enforced. The following descriptions are given in the GRASP Technical Description in Chapter 6 ([TICRA \(2015d\)](#)).

Off-set parabolic systems normally exhibit poor cross-polarisation compared to on-axis systems. They can also suffer from less circularly symmetric PSFs. To compensate for this in dual-reflector systems, the Mizuguchi condition ([TICRA \(2015d\)](#)) can be enforced. This is such that Ψ_f , the angle between the primary and secondary reflector axes is given by:

$$\tan\left(\frac{\Psi_f}{2}\right) = M_0 \tan\left(\frac{\alpha}{2}\right). \quad (3.4)$$

This is the Mizuguchi condition such that the angle α , which describes the angle between the feed axis and the secondary reflector axis obeys $\alpha + \Psi_f = 90^\circ$. When designing an offset reflector, one can ensure each feed in the system obeys this condition by pointing it to point at the centre of the secondary mirror.

3.1.7 Effective Focal Length

If the Mizuguchi condition holds, then it is possible to derive the effective focal length f_e , of the dual reflector optical system in terms of M_0 and f_p , the magnification of the secondary mirror, and the focal length of the primary mirror respectively. Using the angles between the reflector and feed axes, the total effective focal length of the optical system is given by ([TICRA \(2015d\)](#)):

$$f_e = f_p M_0 \frac{1 + \tan^2\left(\frac{\alpha}{2}\right)}{1 + M_0^2 \tan^2\left(\frac{\alpha}{2}\right)}. \quad (3.5)$$

3.2 Optical Analysis

The design principles addressed in the previous section are derived from ray-optics. In this section, I explain the methods used in the optical analysis of reflector antennas generally, and also those used specifically in the characterisation of the optics of X-BASS and NextBASS. This starts with the type of analysis used to investigate the electric current induced on the surface of a reflector, by a propagating electromagnetic (EM) wave. This is a result of the behaviour between the electrons on the conducting reflector surface and long wavelength EM waves, as described by Maxwell's equations. The currents are used in the generation of the telescope radiation, or beam, pattern,

which is a complete description of the diffraction pattern of the optics. This pattern is the resulting gain in all angular directions, produced by the currents induced on the main reflector surface by a propagating EM wave.

3.2.1 PO analysis

To compare the performance of different optical designs, each arrangement was modelled in GRASP and a Physical Optics (PO) analysis was carried out (TICRA (2015c)). The resulting beam patterns form the basis of analysing the optical systems in this chapter. PO analytical methods perform two principal calculations. It first finds the currents induced on the reflector surface by either an illuminating field (eg. from a feed horn), or directly from the currents in another reflecting surface) and then calculates the fields radiated away by those currents. The total field at any point is a sum of the incident and scattered fields is given by (Yassin et al. (2005)):

$$\mathbf{E} = \mathbf{E}^i + \mathbf{E}^s. \quad (3.6)$$

Using geometric ray optics, we can establish the physical limits of illumination on the dish. If we assume that the bright regions we are interested in lie within these limits, then we can obtain the induced PO current between these points. For a perfect conductor, the tangential component of the electric field vanishes on the reflector surface, meaning the induced current at any point can be given by (Yassin et al. (2005)):

$$\mathbf{K}_{\mathbf{E}}^{PO} = \begin{cases} \mathbf{n} \times \mathbf{H} = 2\mathbf{n} \times \mathbf{H}^i, & \text{illuminated surface.} \\ 0, & \text{surface in shadow.} \end{cases} \quad (3.7)$$

and

$$\mathbf{K}_{\mathbf{H}}^{PO} = 0 \quad \text{Everywhere.} \quad (3.8)$$

$\mathbf{K}_{\mathbf{E}}^{PO}$ and $\mathbf{K}_{\mathbf{H}}^{PO}$ are the electric and magnetic current densities respectively, \mathbf{H}^i is the incident magnetic field on the reflector and \mathbf{n} is a unit vector normal to the surface. The \mathbf{E} and \mathbf{H} fields at any point can be found using Maxwell's field equations and the equivalence method. This is in terms of \mathbf{A} , the electric vector potential. We define $R = |r - r'|$ describing the vector between the observation point, r and the integration point r' running over the surface. Hence, \mathbf{A} can be written as (Yassin et al. (2005)):

$$\mathbf{A} = \frac{\mu}{4\pi} \int \int \mathbf{K}^{PO} \frac{e^{ikR}}{R} dS. \quad (3.9)$$

The resulting field equations describing \mathbf{E} and \mathbf{H} can be written as (Yassin et al. (2005)):

$$\mathbf{E} = \mathbf{E} + i\omega\mathbf{A} - \frac{\nabla(\nabla \cdot \mathbf{A})}{i\omega\epsilon\mu}, \quad (3.10)$$

$$\mathbf{H} = \mathbf{H}^i + \frac{1}{\mu}\nabla \times \mathbf{A}. \quad (3.11)$$

In GRASP, the resulting radiated fields are calculated exactly. However, in order to find the induced currents, two approximations must be made (TICRA (2015c)):

1. the geometrical shadow region does not radiate
2. we can neglect non-uniform current near the edges

It should be noted that the first assumption is in fact un-true, but this can be overcome by adding more reflections to better model such shadows. The second assumption holds true for apertures that are more than several wavelengths across. If neither of these hold, then PTD (Physical Theory of Diffraction) analysis may be used as well. PTD separately calculates the currents flowing in the reflector edges and radiates those into the shadow region. The PTD method means that the previous limits on the illuminated regions of the dishes can be replaced by a more accurate set of limits that include shadow regions formed by diffraction edge effects. The GRASP simulations in this thesis have implemented both the PO analysis for accurate surface radiation measurements and the PTD method to accurately estimate the radiation contribution from edge currents. A more detailed description of both PO and PTD methods is given in Chapter 3 of the GRASP Technical Description, (TICRA (2015c)).

3.2.2 Modelling Spillover

The PO and PTD process generates fields induced by current distributions on the dishes. This comes from initial currents formed by incident radiation on the main-reflector, which then induces currents on the sub-reflector and creates the radiation pattern seen by the feed horn. We can model this process in the software GRASP. To do so, GRASP reverses this process and treats the system as a transmitter. The first PO step is to calculate the current distribution induced on the secondary from a feed horn pattern, the next step finds the current induced on the main-reflector from both the feed and sub-reflector. The third step generates the total field on the primary and finally total fields propagate to the far-field. I have outlined these steps below.

1. **Get Currents** > Generates induced currents on a target reflector from a source

2. **Get Field** > Generates the total field on the main-reflector and sends it to the far-field
3. **Add Field** > Adds the additional fields from the sub-reflector and feed horn and sends it to the far-field

When using the PO method to model the reflectors, the commands must be ordered correctly to obtain an accurate estimate of the beam behaviour. The simplest set of commands would be for a single, axial reflector. This would require a command to generate a current from the feed to the main reflector, and then a command to get that field from the main reflector to the sky. Adding a sub-reflector changes the first command so that the feed now illuminates the sub-reflector. We then introduce a second **Get-Currents** command to generate a current on the main reflector induced by the currents in the sub-reflector. All the current objects are then used to calculate the far field on a ‘cut’ file.

These models do not take into consideration double reflections, which can cause additional spillover lobes due to the radiating shadow regions. To model this double reflection, an additional PO surface has to be identified. This can be considered as a second ‘current’ induced on the sub-reflector by currents generated on the main reflector. In Figure 3.4, I list the commands that model the spillover due to the double reflection at the secondary. These are the commands used to perform the PO analysis in GRASP for the OAG and CXD optical arrangements.




Figure 3.4: GRASP commands to correctly account for the secondary spillover lobe. The diagram illustrates the steps in the analysis process to generate the the secondary spillover lobe for the cross-Dragone optics, this is modelled by adding a second reflection (command 3) at the sub-reflector surface, induced by the main-reflector.

3.2.3 Modelling Blockages

The off-axis optics do not suffer from blockages, therefore the PO commands including the second reflection accurately model both the forward and backward beams of the optics. This model is too simplistic for optical configurations that suffer from blockages. Similar commands for on-axis reflectors are described in the GRASP Users Manual. (TICRA (2015b)).

Blockages occur when the primary mirror is obscured by either the feed or the sub-reflector, this occurs in on-axis antennas. The second reflection PO is the correct method of modelling the forward beam of such a system, however the backwards beam that includes the outermost spillover lobes is poorly represented. In this case, we must consider a second surface on the main-reflector, which samples the currents induced by the second surface of the sub-reflector. An alternative method is to include a hole the same size as the sub-reflector rim in the primary mirror. This gives a good estimate of the forward beam but the backward beam is no longer well represented.

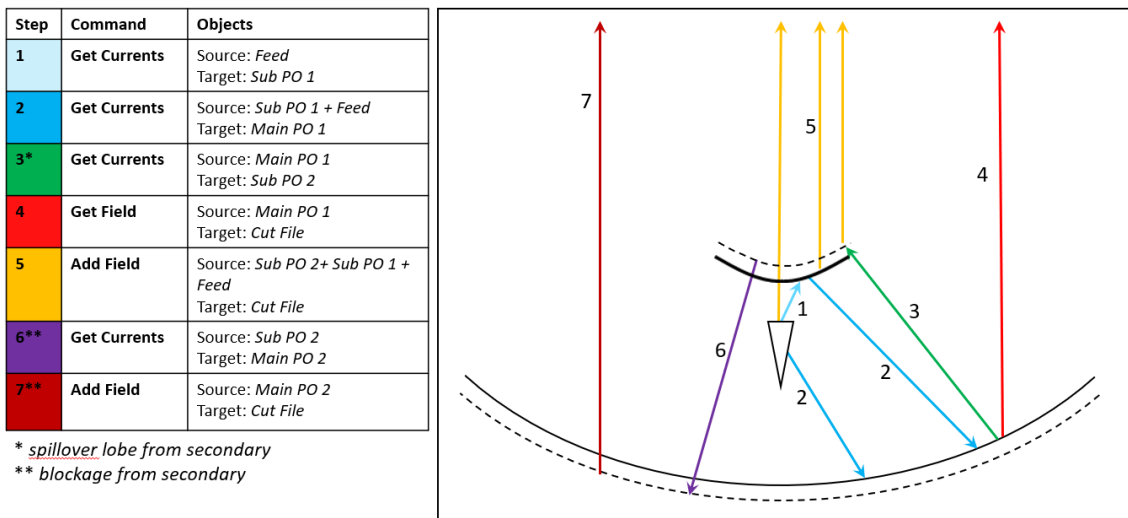


Figure 3.5: The commands required to model the blockages in an on-axis system. It is not possible to correctly model both the forward and backward beams simultaneously, therefore commands 1-5 model the forward beam, and commands 1-7 model the backward beam. By splicing these two together between 70° and 90° degrees, an accurate estimate of the real beam is obtained.

There are no struts in the C-BASS design, keeping it perfectly circularly symmetric. However, there is still a blockage of the main-reflector from the sub-reflector. In Figure 3.5, the commands 1 through 5 model the forward beam, commands 6 and 7 then model the backward beam. The most accurate model of the beam is therefore a

splice of the forward beam model and the backward beam model. I splice these two together at an angle of 70° , as the beams match well at this angle. In Figure 3.6, I show the PO and PTD generated beam pattern for X-BASS at 7.0 GHz. The red beam has been generated by commands 1-5 and the black beam has come from the full backwards set of commands. The shoulder effect around the main beam in the backwards calculation are the additional edge currents from the second sub-reflector surface, these cancel in the forward direction as the second sub-reflector field is sent to the cut. We expect the steep drop off around 110° because the main reflector does not see the outer most spillover lobes as these angles are missed by the limited size of sub-reflector.

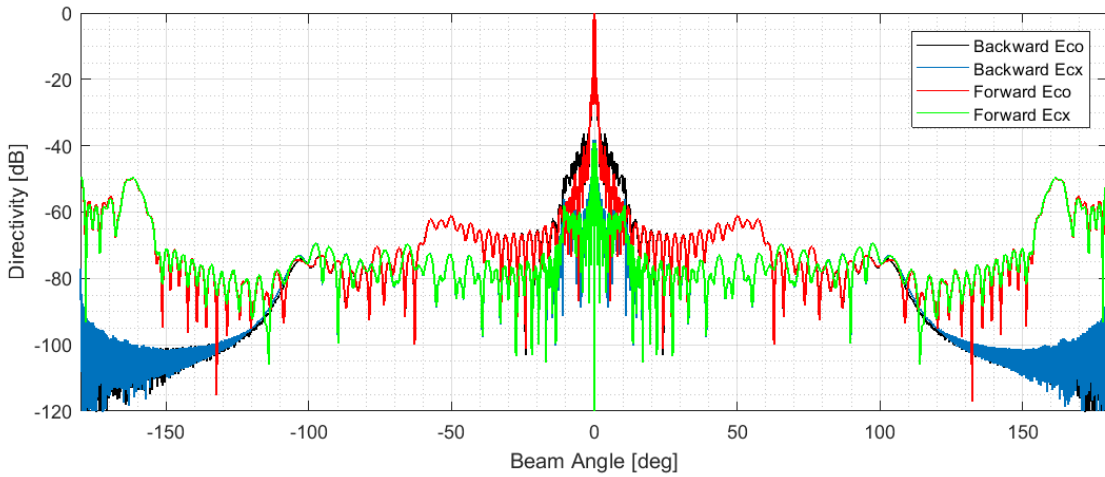


Figure 3.6: The forward and backwards beams of X-BASS, which are spliced together at 70° . The E_{co} and E_{cx} give the co-polar and cross-polar beams respectively. The co-polar beam is taken at a ϕ cut of 0.0° and the cross-polar at $\phi = 45.0^\circ$. The forward beams show a much higher level of spillover at large angles but a narrower main beam. The backward beams correctly model the outer spillover lobes however introduce a shoulder feature in the main beam due to the secondary reflector.

3.2.4 Overlap Integral

The beam pattern gives us the amount of power produced by an antenna as a function of angle in spherical polar coordinates. The overlap integral is the integral that finds the total power in a beam pattern. This depends on the telescope polar pattern. This describes the complete diffraction pattern of the optics. Integrating this beam over solid angle gives us the total power on the sky. An antenna will observe in two polarisations due to the phase differences induced across the aperture. We describe these two phase-states as co-polarisation and cross-polarisation. The total power of a

telescope is therefore the integral of the sum of the co-polar and cross-polar radiation patterns.

3.2.5 Co-polar Beam Pattern

As previously explained, using GRASP we can estimate an accurate model of the telescope beam on the sky. This gives a detailed description of the telescope gain over all angles. In polar coordinates, the beam pattern varies over both the θ and ϕ directions. We have to sample each an appropriate number of both θ and ϕ steps to get a good estimate of all sidelobes that could cause unwanted additional gain in the beam as covered in sub-section 3.2.2.

3.2.6 Cross-polarisation

The polarisation sources that need to be measured by NextBASS are the CMB and synchrotron. Polarised signals from the CMB are formed by scattering processes which produce linearly polarised signals. Synchrotron radiation is also linearly polarised since it is polarised via the alignment of the galactic magnetic field direction. Analysing the optics with linear polarisation produces a cross-polar component. This is the part of the beam pattern that is not aligned with the co-polar beam. If the cross-polarisation peaks are high with respect to the main beam, then it will be difficult to identify the polarisation direction of a source as there will be significant leakage from the optics to produce spurious polarisation signals. We therefore want this cross-polarisation signal to be as low as possible. For an instrument sensitive enough to measure polarised foregrounds, we want this level to be between -35 dB and -40 dB.

3.2.7 Calculating Total Power

The output of the beam pattern in GRASP is the imaginary and real parts of the currents in both the co and cross-polar directions. The total co-polar and cross-polar power is given by the sum of the squares of the imaginary and real components (TICRA (2015b)).

$$\begin{aligned} E_{co} &= (Re_{co})^2 + (Im_{co})^2 \\ E_{cx} &= (Re_{cx})^2 + (Im_{cx})^2 \end{aligned} \tag{3.12}$$

The total power in the entire beam is hence $I = E_{co} + E_{cx}$. Consider the telescope beam on the sky, most of the power is in the main-beam, pointing at an astronomical

source. Any sidelobes outside this beam generate a loss in total power as they are pointing at the noisy sky and or the warm ground. To investigate the telescope’s ability to capture an astronomical source effectively, we need to calculate the overlap integral between the telescope beam and the external temperature pattern across the whole sphere. To sample the beam solid angle, we can use polar coordinates to find $I(\theta, \phi)$. The total power in the beam can be expressed in terms of the overlap integral. This integral calculates the power from each cut in the ϕ direction summed over all θ , see Equation 3.13. Figure 3.7 shows a visualisation of the integral space, broken into θ and ϕ regions as seen from above the optics.




Figure 3.7: Visualisation of integral used to calculate the power across the whole telescope configuration. The resulting field is measured across a number of cuts in the ϕ direction and some $d\theta$. The finer the grid, the more accurate the integral.

$$\int_0^{2\pi} \int_0^{\pi/2} 2I(\theta, \phi) \sin\theta d\theta d\phi. \quad (3.13)$$

The factor of two accounts for Nyquist sampling of the beam when performing the integral. GRASP normalises the total power over all cuts to 4π (TICRA (2015b)). The area underneath a single cut gives the power for that value of $d\phi$. The finer the grid of ϕ cuts, the better the accuracy of the interpolation. However, using too fine a grid of ϕ points results in long computation times. I used 13 cuts which samples every 15 degrees between 0 and 180 in ϕ for NextBASS, this was well normalised. However I found that X-BASS required a finer grid of interval $d\phi = 7.5^\circ$, or 25 cuts. This is because the C-BASS South optics are shaped dishes and it requires more ϕ steps to accurately model the sidelobes. The number of theta points was chosen to be 18001, as this gave a smooth, well sampled beam.

Another improvement to the integral is to interpolate between the ϕ cuts. Since I found the beams to be fairly well normalised with my choice of cuts, I chose a linear interpolator to create more ϕ steps in the integral. I used the MATLAB `interp1` function to re-sample the data at a higher rate in order to create a finer grid than the original. I used this interpolated data to calculate the overlap integral. I found that the additional interpolation increased the accuracy of the integral as the normalisation constant was closer to 4π . However, this interpolation will only improve accuracy if the ϕ -space is sufficiently sampled in the first place.

3.2.8 Directivity

The directivity is the maximum forward gain of the telescope. In terms of the beam pattern, this is given by the maximum co-polar component of the beam pattern. This will be located at zero degrees and can be used to normalise the beam so all other features can be described by a reduction in power.

3.2.9 Peak-to-peak Cross-polarisation

Beneath the main beam there will be some contribution from cross-polarisation. In an on-axis system, there are two planes of symmetry for linear polarisation where the cross-polarisation is zero, giving a typical four-lobe structure to the cross-polarised beam. For off-axis systems, there is only one plane of symmetry, giving a characteristic two-lobe pattern. The maximum cross-polarisation level is given from the difference between the peak cross-polar power and the directivity of the main beam. The maximum cross-polarisation usually occurs at a ϕ cut of 45.0° . The target cross-polarisation level would be between -35 dB to -40 dB at $\phi = 45.0^\circ$. During observations, a CMB foreground telescope requires this level of cross-polarisation to minimise leakage between the two linear polarisation channels.

3.2.10 Sidelobe level

The sidelobes are the local maxima of the far-field radiation pattern. They are identified as ‘lobes’ and are separated by ‘nulls’. A good beam pattern will have sharp nulls meaning that there is a uniform phase pattern across the aperture. For a circular aperture that is uniformly, and fully illuminated, the first sidelobe level is given by the Airy pattern and corresponds to a peak of ≈ -17.6 dB (Hecht (1997)). A uniformly illuminated aperture distribution means a maximum possible directivity, however also

maximum sidelobe levels. The sidelobe levels can be reduced by changing uniformity, this could be achieved by designing a feed horn that produces a non-uniform illumination pattern, and to under-illuminate the dish.

A good indicator of the optical performance of the antenna would be to ensure that each configuration is capable of achieving approximately this theoretical limit. Ideally, the combined optimised feed and optics should be able to obtain between -15 dB and -20 dB for the first sidelobe level across the band.

3.2.11 Beam Width

A good indication of the beam width comes from the 3dB level of the beam. This is the full width of the at half power of the maximum gain, or directivity. This can be used as an indicator for whether the optics are diffraction limited. The diffraction limit of a telescope is defined by Fraunhofer diffraction. A circular aperture will create a point spread function known as an Airy spot. The angular size of this Airy spot indicates the maximum resolution of the beam of the sky. If the half power full beam width is approximately the Airy spot, then the optics are diffraction limited and can achieve full resolution (Hecht (1997)). For an Airy pattern, the separation, α , at which two beams have distinct peaks, is given by:

$$\alpha = \frac{\lambda}{D}. \quad (3.14)$$

3.2.12 Beam Efficiency

The beam efficiency, η_{beam} defines the percentage of power in the main-beam. It is often taken to be between the first two nulls between the first sidelobes of the main-beam. It can also be specified over any angular region that will be looking at the sky, and is defined by

$$\eta_{beam} = \frac{P(\theta_{beam})}{P(\theta_{total})} = \frac{\int_{beam} f^2(n)dn}{\int_{4\pi} f^2(n)dn}. \quad (3.15)$$

The remaining power outside of this is lost due to sidelobes and as a result can increase the system temperature. This is particularly an issue for the farthest sidelobes, which point at the warm ground. The beam efficiency therefore tells us how much power is actually on the cold sky and how much power is pointing at the warm ground (Burke and Graham-Smith (2010)).

3.2.13 Aperture Efficiency

The aperture efficiency η_{ap} is defined by the ratio of effective collecting area to the physical area of a telescope. This is given by a number between 0 and 1, where 1 means that the effective collecting area is equivalent to the physical aperture size. The effective area A_{eff} , is related to the forward gain of the telescope, G , by $A_{eff} = G \frac{\lambda^2}{4\pi}$. The aperture efficiency is then the ratio of A_{eff} to the physical dimensions of the optics, given by $\pi(D/2)^2$, where D is the diameter of the primary reflector (Burke and Graham-Smith (2010)).

$$\eta_{ap} = \frac{A_{eff}}{A_{real}} = \frac{G\lambda^2}{\pi^2 D^2} \quad (3.16)$$

There is a trade-off between beam efficiency and aperture efficiency as a fully-illuminated primary will lead to stronger sidelobes and greater spillover. Since we want to design stable systems with as low noise as possible, we benefit more from lower sidelobes than a high gain beam. For CMB foreground experiments, we are not interested in detecting distant point sources, instead we are measuring faint diffuse emission. If the forward gain of the telescope is too high, then observations of the diffuse medium will be undetectable beneath the bright point sources. We can hence say that the beam efficiency is more important than the aperture efficiency, since we want most of the power to be in the main beam but with a low enough gain to avoid saturation of the galactic emissions in the presence of point sources. However, we do still require point sources to calibrate the telescope, therefore we can say that the aperture efficiency should be less than 60 % in order to perform an accurate calibration whilst avoiding saturation of the diffuse emission.

3.2.14 Beam Ellipticity

Along with cross-polarisation and the power in the first sidelobes, a good beam must also be circular. The more circular the beam, the more uniform its point-spread-function, and the easier it can be cleaned in the post-processing stage. For cases involving beams that are off-axis or suffering diffraction effects due to surface irregularities, the beam may lose some of its circularity. A good assessment of beam asymmetry in the θ direction can be found by the difference between negative and positive values of θ at a certain dB level (typically 3 dB) with respect to the mean value. When this value is equal to zero the beam is symmetrical, anything greater than this means the beam has become asymmetrical.

$$\text{beam ellipticity} = \frac{\theta_- - \theta_+}{2(\theta_- + \theta_+)} \quad (3.17)$$

We can also investigate the circularity of the beam in the ϕ direction. This can be seen from the far-field beam patterns as the difference between the 0.0 and 45.0 degree cuts. If these two ϕ cuts are similar in gain then the beam has good circularity. If there is a considerable difference between them (greater than a few dB), then the beam is not well circularly symmetric. This results in a higher cross-polarisation across all angles.

3.3 X-BASS Optics

The X-BASS experiment is a plan to replace the existing 4.5 – 5.5 GHz receiver on C-BASS South, with a receiver covering 7 – 15 GHz. The C-BASS South were designed to achieve low cross-polarisation with a high circularly symmetric optical configuration for a 1 GHz bandwidth centred at 5.0 GHz. This means if we want the optical performance of X-BASS to match C-BASS, we must perform some stringent modelling to ensure X-BASS will meet these requirements. As a result, a feed design that will give a good illumination of the optics over the very wide frequency range 7 – 15 GHz must be implemented as well checking that the C-BASS dish surfaces will not introduce diffraction effects at these frequencies. To perform this check, information on the surface profile from photogrammetric images must be used to ensure that the surface errors will not cause unacceptable diffraction effects at the higher frequencies. I have tested this with the photogrammetric data obtained at C-BASS in 2016 (see page 219 ([Jew \(2017\)](#))). By analysing the beam patterns, we can then determine if the performance of X-BASS. The X-BASS optics must hence meet the following requirements:

1. Achieve a peak-to-peak cross-polarisation level of -35 to -40 dB consistently across the band
2. First sidelobes no greater than -15 dB
3. Far-out sidelobe levels below -40 dB
4. A beam ellipticity close to 0 to ensure a circular beam
5. A minimum of 90 % beam efficiency for a one-degree beam as this ensures that less than 10 % of the power is lost to the warm ground, maintaining a low system temperature and hence a high sensitivity

Since C-BASS was designed to be very under-illuminated, this makes the optics ideal for X-BASS as they should already achieve low spillover. I aim to show that the C-BASS South optics are able to obtain a similar performance as C-BASS but up to frequencies of 15 GHz, over the band 7 – 15 GHz.

3.3.1 C-BASS South

Both the C-BASS North and South dishes were recommissioned from old telecommunications antenna (Jones et al. (2018)). C-BASS North is a 6.1m telescope and C-BASS South is 7.6m in diameter (Holler, 2007 (priv. comm.)). The main-reflector has been kept the same apart from re-surfacing and some panel alignment to improve surface accuracy. However, for both C-BASS South and North, the secondary reflector has been redesigned (Holler et al. (2011)).

The main reflector is a ‘shaped system’, i.e. it deviates from a classical paraboloid. This was done to increase the aperture efficiency for the original telecommunications application, because of this, the secondary mirror design has also been re-shaped to compensate for the gain in far-out sidelobes. The shape has a slight peak in the centre to reduce standing waves between the secondary and the feed. The secondary is also slightly oversized to reduce the outer-most spillover rays by changing their direction to not diffract over the edge of the primary. To keep symmetry and avoid the use of struts, the secondary mirror is supported on a piece of radio transparent foam. This maintains the telescopes’ excellent circular symmetry (Holler, 2009 (priv. comm.)). The material of the radio transparent foam is Plastazote LD45 (King et al. (2010)), made of polyethylene that is transparent at frequencies up to infrared (Zotefoam (2010)).




Figure 3.8: A schematic describing the optical configuration for C-BASS. Taken from the optical design memo. Both the primary and secondary reflectors have been shaped and differ slightly from a classical Cassegrain (Holler, 2009 (priv. comm.)).

Since C-BASS is close to a classical Cassegrain, it is possible to recreate an equivalent optical configuration in GRASP. Using the Wizard in GRASP I generated a primary reflector with the same focal length as the main reflector of C-BASS, as identified by the position of the feed horn in Figure 3.8. I then changed the eccentricity of the sub-reflector required to give an angle of illumination subtended by the feed and the sub-reflector to match the 46.3deg specified by the C-BASS memo. I also ensured that the resulting sub-reflector radius was the same as that of C-BASS. Figure 3.9 shows the GRASP wizard that describes this Cassegrain.

Parameter	Equivalent Value
Primary focal length	$f_p = 1.9735\text{m}$
Secondary eccentricity	$e = 2.6$
Primary diameter	$D_p = 7.6\text{m}$
Secondary diameter	$D_s = 0.486\text{m}$
Half angle	$\theta_f = 46.33^\circ$
Equivalent Focal length	$f_e = 4.440\text{m}$
f - number	$\frac{F}{D} = 0.584$

Table 3.2: The design parameters of an equivalent Cassegrain to match the C-BASS optics.

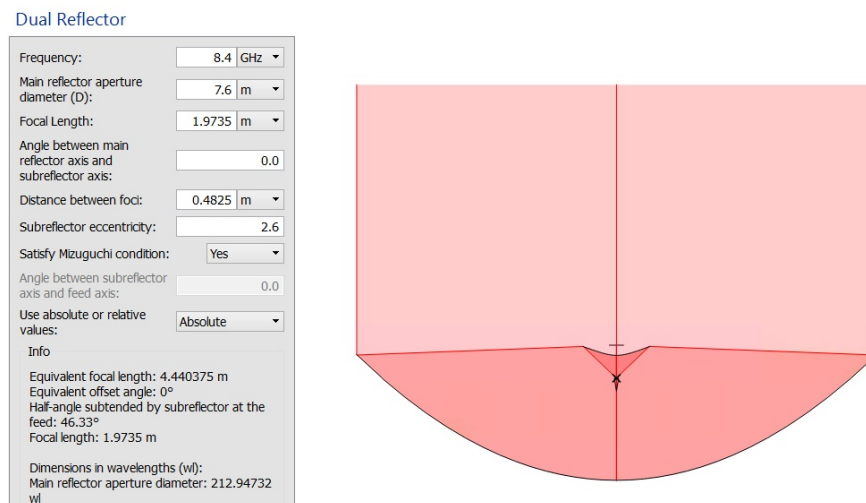


Figure 3.9: Snap-shot of the wizard used to obtain a ‘C-BASS-like’ Cassegrain reflector and hence the effective focal length and equivalent mirror eccentricity of the sub-reflector. The parameters of this equivalent Cassegrain are well-matched to the X-BASS optics.

The effective focal length of such a system will be close to that of the real C-BASS dish. In Table 3.2 I show the specifications of the equivalent Cassegrain. From this




Figure 3.10: Co-polar plot of the equivalent Cassegrain, the ideal and real optics, modelled with the photogrammetry data. Each are well matched. We can see that the shaped secondary reflector has reduced the outer sidelobes, improving performance from the equivalent Cassegrain. It should also be noted that there is only a small increase in sidelobes when including the effects from surface modelling.

we can see that the X-BASS system will have a focal length of $\sim 4.44\text{m}$ and an f -number of ~ 0.58 . Figure 3.10 shows that the classical Cassegrain and C-BASS optics have almost identical beam patterns. The real C-BASS optics achieve lower spillover overall, however suffer from larger first sidelobes. Since the main beam is identical in width, this Cassegrain system is a good equivalent to C-BASS.

The surfaces of C-BASS have been mapped by photogrammetric images. Using this data, we can map the surface to a best fit polynomial. This polynomial can be used to describe the primary mirror surface in the GRASP simulation. The fifth order polynomial that describes the primary mirror has been interpolated from photogrammetry data and can currently be described by equation 3.18.

$$z = ax^2 + by^2 + cx + dy + e \quad (3.18)$$

where the best fit 5th order polynomial values are $a = 0.933$, $b = 0.0933$, $c = -2.414 \times 10^{-4}$, $d = 1.1808 \times 10^{-4}$ and $e = -0.0143$.

The C-BASS South secondary mirror was designed by [Holler, 2009](#) (priv. comm.). To keep blockage to a minimum, the secondary was kept at a minimum size of 0.5m in radius. The mirror has been profiled to include a peak at the centre. The central tip reduces standing waves which cause large angle spillover from the centre of the mirror. The surface data was produced by [Holler, 2009](#) (priv. comm.) and has been

used in the GRASP simulation. In Figure 3.11, I show the ray diagram of the C-BASS secondary mirror, demonstrating the effect of the central peak preventing rays from being reflected back to the feed.




Figure 3.11: C-BASS Secondary mirror and feed ray diagram showing that the central peak shaping of the sub-reflector changes the direction of the central rays to reduce the overall spillover at large angles, and increase the power on the main-reflector.

3.3.2 Photogrammetry

Photogrammetry is used to measure surfaces from photographs. It enables the recovery of surface profiles. It works by taking many photographs at different angles of markers placed on the reflector surface, to obtain a measurement of the distance between the points and the depth of the image. Photogrammetry uses the redundancy in many photos of the same markers from different directions to solve for the 3-D positions of the markers. It is useful for radio astronomy as the dish profiles can be measured accurately and it can be investigated whether the surface accuracy is enough to meet the specifications of the instrument.



Figure 3.12: C-BASS South in operation Klerefontein, South Africa. The photogrammetry targets are visible on the twelve panels that make up the primary reflector and the foam cone supports the secondary mirror on top of the receiver (Taylor (2018)).

The photogrammetry results used in this experiment were obtained with the most recent measurements taken by Luke Jew and Angela Taylor in 2016. The photogrammetry data set consists of 261 surface data points. Once the points have been identified in the photographs using the targets on the dish, the relative positions of the points can be recovered using the software PHOTOMODELER. The absolute distances are then scaled from the relative positions by the distance between three known points.

Using modern cameras, it is possible to perform photogrammetry to a precision of $\frac{\text{largest object dimension}}{5 \times 10^6}$ (Luhmann (2010)). However, the camera used for the C-BASS photogrammetry data obtains a target accuracy of $\approx 0.8\text{mm}$ (Copley (2014)). After obtaining the *rms* surface distortion using Equation 3.19, we can propagate this uncertainty to give the surface accuracy to a precision of 0.87mm , hence the overall surface accuracy is $2.63\text{mm} \pm 0.44$.

Using their results I have calculated the *rms* distortions on the surface of the dish. This allows us to decide whether the surface tolerances are suitable for a feed up to 15 GHz. For C-BASS, markers were positioned on the surface at certain height, the dish is then marked with additional targets that each have a known x - y position. These are positioned at equal intervals across the dish. By taking photographs at different angles, it is possible to measure the depth of each target. Figure 3.12 shows the C-BASS South optics with the photogrammetry targets positioned on the surface.

To fit the polynomial to the surface to find the *rms* distortions on the surface, the data has to have the targets extracted. These are identified by an initial fit of the polynomial to obtain the surface depth. Any values larger than a specified value indicate the position of the target. These data points are then removed and the data is re-analysed. The error in the z direction is calculated as being the difference between the polynomial at that (x, y) coordinate and the measured value of z at the same (x, y) . The *rms* surface accuracy is calculated by:

$$z_{rms} = \sqrt{\frac{1}{n} \sum (z_p - z_m)^2} \quad (3.19)$$

where n is the total number of surface points, z_p is the z -displacement predicted by the polynomial and z_m is the z -displacement of the real dish at the same (x, y) coordinate.

Using the values for z_m for the corresponding (x, y) coordinates, I defined a new irregular spline surface in GRASP. This surface takes the photogrammetry data and interpolates a surface between the points using a cubic spline algorithm. This surface can then be used to replace the original primary reflector surface profile. Figure

3.14 shows the ray diagram with the 261 surface spline points of the GRASP X-BASS model.




Figure 3.13: The surface tolerances of the C-BASS dish taken from photogrammetry data obtained in 2016. The numbers indicate the positions of the 12 panels. The Root-Mean-Square surface distortions was found to be $\approx 2.6\text{mm}$ which is about one-tenth of 2cm wavelength at 15.0 GHz . As a result, it is expected that there will not be much distortion of the beam at 15.0 GHz .




Figure 3.14: GRASP Simulation of X-BASS with the photogrammetry data points. Each point on the surface is a photogrammetric data point. The resulting surface is a cubic spline interpolation between these points.

3.3.3 X-BASS Performance

I will now look at the analysis of the X-BASS optics using GRASP. Adding the photogrammetric data to the GRASP simulation introduces asymmetries in the beam

profile. These should be most visible at higher frequencies as the surface errors will be closer to the wavelength. Figure 3.16 shows the full beam at 15.0 GHz for both the ideal dish and the dish with the surface errors. There is little difference in the co-polar beam and since both the E_{co} $\phi = 0.0^\circ$ and E_{cx} $\phi = 45.0^\circ$ degree beams are similar, the real dish has not lost circular symmetry at large scales. The cross-polarisation however has increased as the E_{cx} 90.0° and E_{cx} 0.0° patterns are showing features that are not present in the ideal beam.

Figure 3.17 plots the beam directivity, cross-polarisation, first sidelobe level, beam efficiency within a one-degree beam, beam circularity and the aperture efficiency of the X-BASS optics with and without the surface data. These beam properties remain mostly the same with the introduction of the surface data, the major difference is that the ideal C-BASS optics are nearly perfectly circular where as the surface data shows increasing asymmetry with increasing frequency. Figure 3.18 shows the main beam patterns for X-BASS with the surface data included. Each beam pattern shows the co- and cross-polar plots for the 0.0° , 45.0° and 90.0° degree cuts. By plotting these cuts we can see that the surface errors introduce some asymmetries in the main beam however, we can see that the beam ellipticity is still good between the different cut angles.






Figure 3.15: Cumulative Power of X-BASS with the photogrammetric surface data at 1 GHz bandwidths. The higher frequencies suffer from greater spillover as the cumulative power does not converge to 1 until larger angles. This was calculated with the optimised feed horn developed for X-BASS. The cumulative power converges closer to one at narrower angles for higher frequencies as the physical diffraction effects become less severe at higher frequencies, resulting in lower spillover between the reflectors.

The cumulative power is a way of showing how much power is present in the beam

at each angle. The cumulative power plots correspond to the sidelobes seen in the beam patterns in Figure 3.16 and how the power in the beam increases as a function of angle. By combining both halves of the integral into one side, it is possible to plot the integral against beam angle. In Figure 3.15, we can see that more than 95% of power in the main beam is present in the first 2° of the beam. These simulations used the final optimised X-BASS horn design, which is described in Chapter 4 and the final specifications are given in Figure 4.6c. The channels with the large spillover lobes at 9.0 and 10.0 GHz as these have a step at 50.0°. This is the spillover from the feed and is a result of it over-illuminating the secondary. This improves considerably at higher frequencies.



(a) Ideal Dish - no surface errors



(b) Real Dish - with surface errors

Figure 3.16: Comparison of the full beam cuts at 15.0 GHz for the C-BASS optics modelled with and without the photogrammetric data. The photogrammetric data has increased the power in the cross-polarisation of the 0.0° along with creating slightly asymmetric first sidelobes. Otherwise the real beam still meets the requirements of the X-BASS optics.

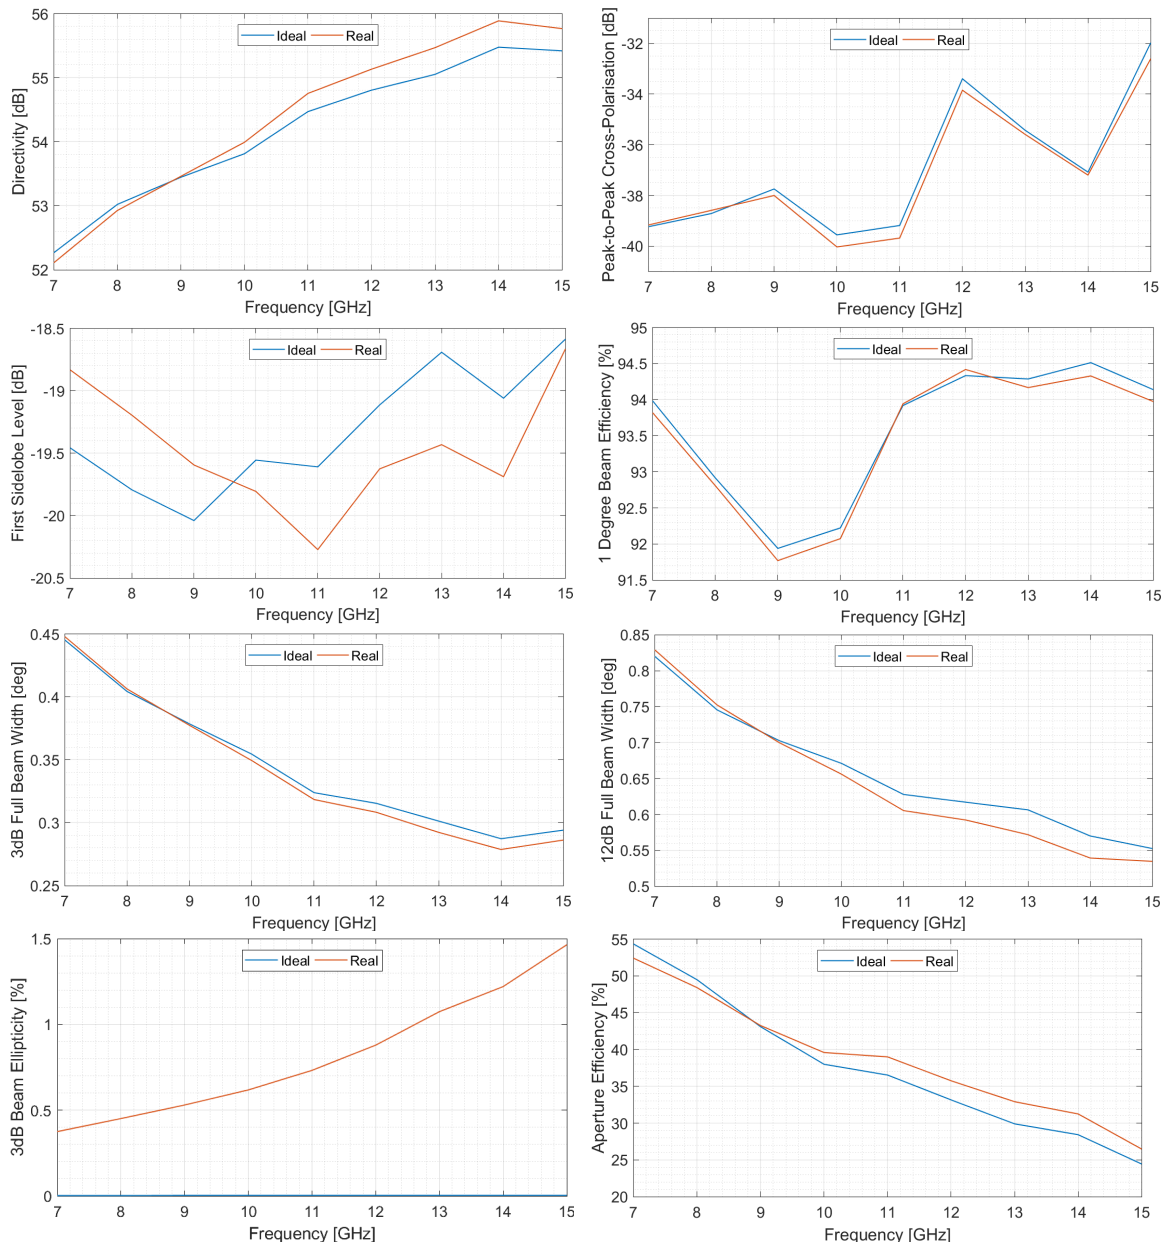


Figure 3.17: Beam Parameters for X-BASS for both the ideal dish and the dish modelled with the photogrammetry data. This shows that there the main difference the surface tolerance introduces is an increase in the first sidelobe level at low frequencies as well as a loss of circularity in the beam, as indicated by the increasing ellipticity of the beam.

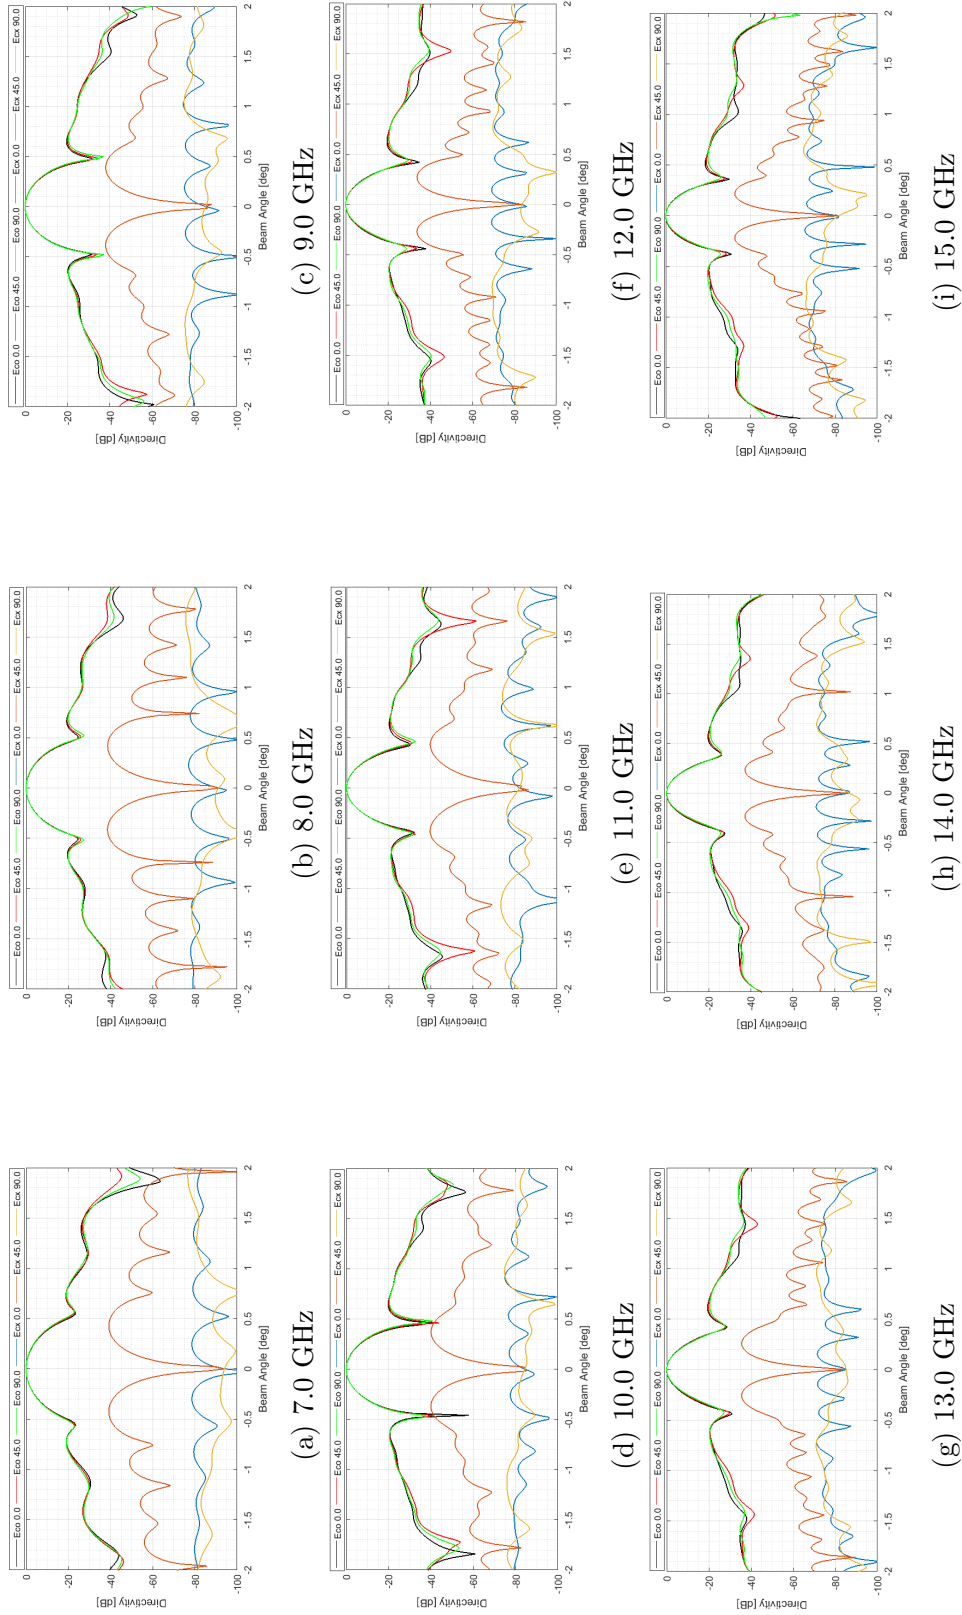


Figure 3.18: The mainbeam across the X-BASS frequency channels modelled with the surface data obtained for the C-BASS primary mirror. The plots show the co- and cross-polar components of the beam for $\phi = 0.0^\circ, 45.0^\circ$ and 90.0° . There is some asymmetry in the sidelobes, which is worse at the 15.0 GHz, however the main beam is still symmetric. Each of the ϕ cuts are well aligned, meaning the beam is highly circular.

3.4 NextBASS Optics

NextBASS is planned to be the ‘ideal’ telescope for CMB foregrounds in the range 7 - 30 GHz, which needs an optimal optical system designed for it. Both the OAG and CXD designs are un-blocked apertures and have often been used for high sensitivity instruments. In order to investigate each design, I have created both systems in GRASP and performed a full analysis of both. In order to obtain the required sensitivity needed for foreground removal at the higher frequencies, a feed array is required. Therefore, the NextBASS optics must be a system that can perform well across a large focal area, with very low sidelobes and good beam circularity and cross-polarisation.

3.4.1 Performance Targets

For NextBASS to maintain a high sensitivity, low system temperature and measure polarisation to a high fidelity, there are requirements on the optical configuration that must be made. These are similar to those set out for X-BASS however, since NextBASS is a specialised design, we need to target higher performance goals. For example, the main systematic in the C-BASS data is ground spillover, due to its design it is not possible to create a ground-shield for X-BASS. However, using an off-axis design for NextBASS means that it can be designed with a ground-shield and baffle that maintains all but the main beam at a constant temperature. This reduces the systematics at a potential increase in system temperature.

1. Achieve a peak-to-peak cross-polarisation level of -35 to -40 dB consistently across both bands to ensure minimal polarisation leakage between the two linear polarisations
2. First sidelobes no greater than -15 dB as this can contribute to contamination of the main beam
3. High beam symmetry to ensure that cross-polarisation is minimised and that the beam PDF is circular
4. Large, flat focal plane suitable for a feed array
5. Little cross-polarisation deterioration across the focal plane
6. A viable method to control ground-spillover to reduce ground systematics

3.4.2 Detailed Geometry

In Figures 3.19a and 3.19b I have shown the schematics of the two optical systems. These are to-scale, based on their 3D visualisations generated in GRASP. These schematics include the parameters that describe the full optical system, such as the angles between the main reflector, sub-reflector and feed axes. I based the CXD configuration on the Clover optics. The Clover optics were CXD telescopes of diameter 2m to be used with a bolometer array of 90, 150 and 220 GHz detectors. The angle between the primary and secondary axes, $\Phi = 65.0^\circ$, was optimised to maximise the usable area of the focal plane (this is discussed in the Clover CDR Documentation [Johnson et al., 2007](#) (priv. comm.)). The angle $\Phi = 65.0^\circ$ maintains a compact design, without the secondary being too close to the primary, as this can cause unwanted edge currents, and hence large secondary spillover lobes. I used this as a starting point to obtain the optical configuration for the 6m CXD telescope. This corresponds to a design angle of $\Phi = 65^\circ$ between the main and sub-reflector axes, and an angle of 25° between the secondary and feed axes.

The OAG has a much smaller sub-reflector diameter and hence a lower spillover. Using a paper by Shaul Hanany, ([Hanany and Marrone \(2002\)](#)), I chose the axial tilt of the OAG to be 15° . When designing the OAG optics, I change the feed position to be on the same plane as the primary mirror, this is ideal for a large focal plane that would require a heavy receiver. If the feed is not in the same plane as the primary mirror then the design would require struts that can block the primary mirror. I outline the parameters that define both the CXD and OAG systems in Table 3.3, and the corresponding schematics for each optical configuration is given in Figure 3.19.

	Parameter	CXD	OAG	Description
<i>System</i>	Ψ	65°	15°	Angle between primary and secondary axes
	α	25°	23.96°	Angle between feed and secondary axes
	$\frac{F}{D}$	3.528	0.533	F-number of system
	f_e	9.870m	7.812m	Effective focal length of equivalent paraboloid
	d_{ps}	22.151m	7.733m	Distance between primary and secondary mirror centres
<i>Primary</i>	D_p	6.000m	6.000m	Diameter of primary mirror
	f_p	21.167m	3.199m	Focal length of primary
	f'_p	27.196m	3.309m	Usable focal length of off-axis primary mirror
	θ_p	56.18°	34.14°	Angle between parabolic vertex and off-axis mirror centre
<i>Secondary</i>	D_s	6.067m	1.681m	Diameter of secondary mirror
	f_s	29.194m	3.199m	Focal length of secondary ($2c$)
	f'_s	7.576m	2.453m	Usable focal length of off-axis primary mirror
	θ_s	0°	61.083°	Rotation angle of secondary axis
	e	-2.067	0.46	Eccentricity of secondary mirror
	$2a$	-14.121m	7.466m	Semi-major axis of secondary mirror
	M_0	0.347	-2.704	Magnification from secondary
	R_1	7.981m	2.904m	Length subtended between off-axis point and first foci
	R_2	8.0627m	4.581m	Length subtended between off-axis point and second foci


Table 3.3: Parameters for the 6.0m NextBASS telescope. This includes the geometrical and optical focal lengths of the primary and secondary reflectors including the eccentricity of the hyperbolic mirror, the distance between the primary and secondary mirror (d_{ps}), the distance between the feed and the secondary reflector (R_2), the diameters of the primary and secondary reflectors (D_p, D_s) and the effective focal length of the entire telescope. e is the eccentricity of the secondary mirror.

3.4.3 Beam Properties


Both optical systems were simulated in GRASP using the methods described in section 3.2.2. The beam patterns of each optical arrangement have unique spillover features. At angles greater than the main beam, each off-axis configuration has a distinctive pattern of sidelobes due to the asymmetries of both designs. These spillover lobes

are greatest at the lowest frequencies and at the $\phi = 0.0$ cut due to more prominent diffraction effects. As seen in Table 3.3, the sub-reflector of the CXD is, by a small percentage, larger than the main-reflector. A larger secondary will reduce the primary spillover at $\approx -125^\circ$, however the double-reflection sidelobe at $\approx 50^\circ$, will increase. This is because the secondary mirror is so large, the shadow regions will be more prevalent, so we expect to see additional sidelobes when including the double-reflection. In Figure 3.20a we see the difference between single and double reflections. There is a peak at around 50° that is caused by the double reflection sidelobe. In contrast, the small sub-reflector size in the OAG telescope results in lower sidelobes. We can compare the beam patterns of the CXD and OAG in Figures 3.20a and 3.20b respectively. It is apparent that the OAG spillover is consistently below -50dB at all angles. The CXD design suffers from a sub-reflector spillover of up to -40dB at an angle of $\approx 125^\circ$. As a result, we anticipate the beam efficiency of the CXD to be lower than the OAG.

At higher frequencies, the beam efficiency will increase. This is because at shorter wavelengths, Fraunhofer diffraction across the main aperture is less prevalent. We also see a decrease in the secondary spillover because the gap between both reflectors is now even larger than the wavelength. I compare the spillover lobes in Figure 3.21 showing that an increase in frequency decreases the overall spillover. Diffraction effects from a circular aperture will be in phase when in the on-axis direction of the reflector. Within the CXD spillover lobes, there is a peak, which arises because of an effect known as the Arago spot (also known as Poissons spot) (Hecht (1997)). This is because the sub-reflector acts as a circular aperture that blocks the waves coming from a point source at infinity. The intensity of this peak is equal to that of the incoming wavefront, this is why the height of this feature does not change considerably across the band. The width of the Arago spot region depends on the diameter of the circular object and the wavelength. For a reduction in wavelength with a constant diameter circular aperture, the width of the spot region decreases slightly. We can see this in Figures 3.21a and 3.21e as a reduction of width with increasing frequency. The curvature becomes steeper because other diffraction effects across the aperture are less prevalent.



(a) Spillover lobes of the 6m Cross-Dragnone telescope. The red lobe at -50 degrees indicates the spillover from the secondary. The lobes to the right hand-side are caused by the primary mirror and the peak at the centre is a typical feature of circular obstacles known as the Arago spot, a product of Fresnel diffraction (Hecht (1997)).



(b) Spillover Features of the 6m off-axis Gregorian optics for different ϕ cuts, which show differing sidelobe positions due to the asymmetry of the OAG optical configuration.

Figure 3.20: Full beams of co- and cross-polar cuts for the CXD and OAG telescopes being investigated for the NextBASS optics. The CXD figure shows the effects of the double-reflection sidelobe, which is not visible in the OAG optics because of the significantly smaller sub-reflector.

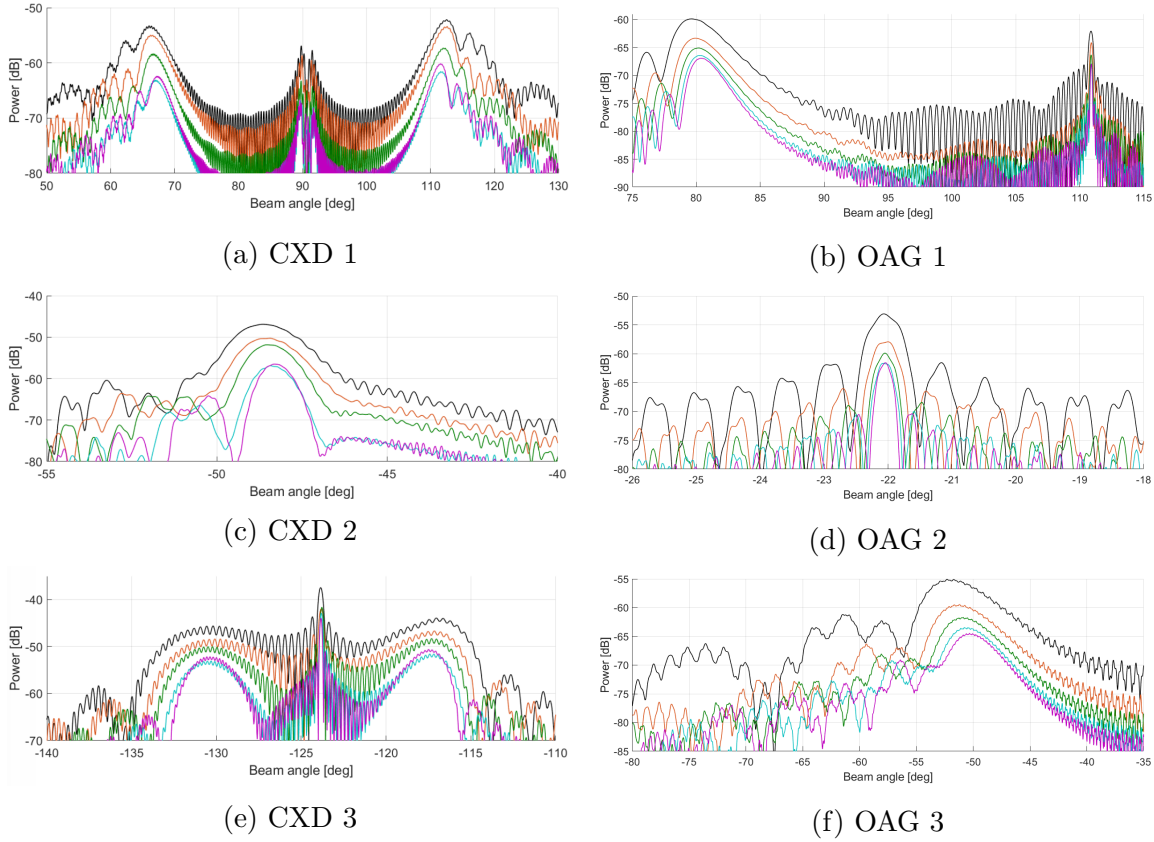


Figure 3.21: The outer most spillover lobes at different frequencies for both optical configurations. Black - 7.0GHz, red - 9.0 GHz, green - 11.0 GHz, cyan - 13 GHz, purple - 15.0 GHz. Note the axes have been scaled separately to appropriately show the change in power with frequency for spillover lobes at different angles for the CXD and OAG optics. The CXD (left-hand side) has more power in these sidelobes due to the larger sub-reflector. The OAG (right-hand side) has much lower power spillover lobes with less distinctive shapes due to its small secondary reflector. The Arago spot feature is visible in the main-reflector and sub-reflector spillover (Figures 3.21a and 3.21e respectively).

I have extracted the beam parameters from the optical simulations carried out in GRASP. These include the beam widths across the band, beam efficiency and the maximum cross-polarisation levels. Figures 3.22 and 3.23, plot these beam parameters as a function of frequency for ‘band-1’ (7 – 15 GHz) and ‘band-2’ (15 – 30 GHz) respectively. It must be noted that the CXD simulations have been done with a hyperbolic profile feed, the design of which is outlined in the following chapter. The OAG simulations used a hybrid mode corrugated feed, generated in GRASP to match the same edge taper as the hyperbolic feed for the CXD. This should only affect cross-polarisation performance and first sidelobe level.

The differences in the types of feeds modelled with the optics should only have a small effect the beam full-width half maximum and the power in the first sidelobe level. As expected, we can see that the OAG achieves slightly higher beam efficiency for band-2. As discussed earlier, this is because the OAG design has a smaller sub-reflector and therefore an increase in frequency means a decrease in spillover. I carried out some system optimisation in GRASP to qualitatively assess the effect of decreasing the size of the secondary reflector in the CXD optics. I found that reducing the size of the secondary reflector does minimise the spillover lobe from the secondary reflection. However, the overall edge taper on the primary is increased as a result, so there is actually an overall net gain in spillover. Further optimisation of the off-axis Gregorian design has not been carried out, however due to the asymmetry of the optics, a large secondary reflector would result in a far more asymmetric spillover profile. Both perform well in cross-polarisation, achieving levels well below -30 dB. Since the CXD simulation uses a realistic feed model, the OAG cross-polarisation is a slight overestimate as a result of the ideal hybrid mode feed. A similar statement can be made about the first sidelobes however for band-1, the CXD design achieves very low first sidelobe levels towards the upper middle of the band.

Looking at the 3 dB full beam width, the CXD design varies from approximately 0.54° to below 0.48° for band-1 and between 0.26° and 0.12° for band-2. We expect a diffraction limited system to achieve $\approx 0.5^\circ$ at 7 GHz, $\approx 0.23^\circ$ at 15 GHz and $\approx 0.12^\circ$ at 30 GHz. This means that the CXD configuration is diffraction limited through band-2 and the lower half of band-1. This is due to the beam width consistency the hyperbolic profile feed designed for NextBASS can achieve (I discuss this more in the following chapter). If we consider a diffraction limited system, the beam size should correspond to λ/D . At 7.0 GHz, we expect the beam size to be 0.4° , we want to maintain this resolution across the band to make the map making process easier, even with the loss in resolution. We can see that the 3 dB beams are slightly higher than this limit for both the OAG and CXD. This means there is a slight reduction in resolution compared to the theoretical value.

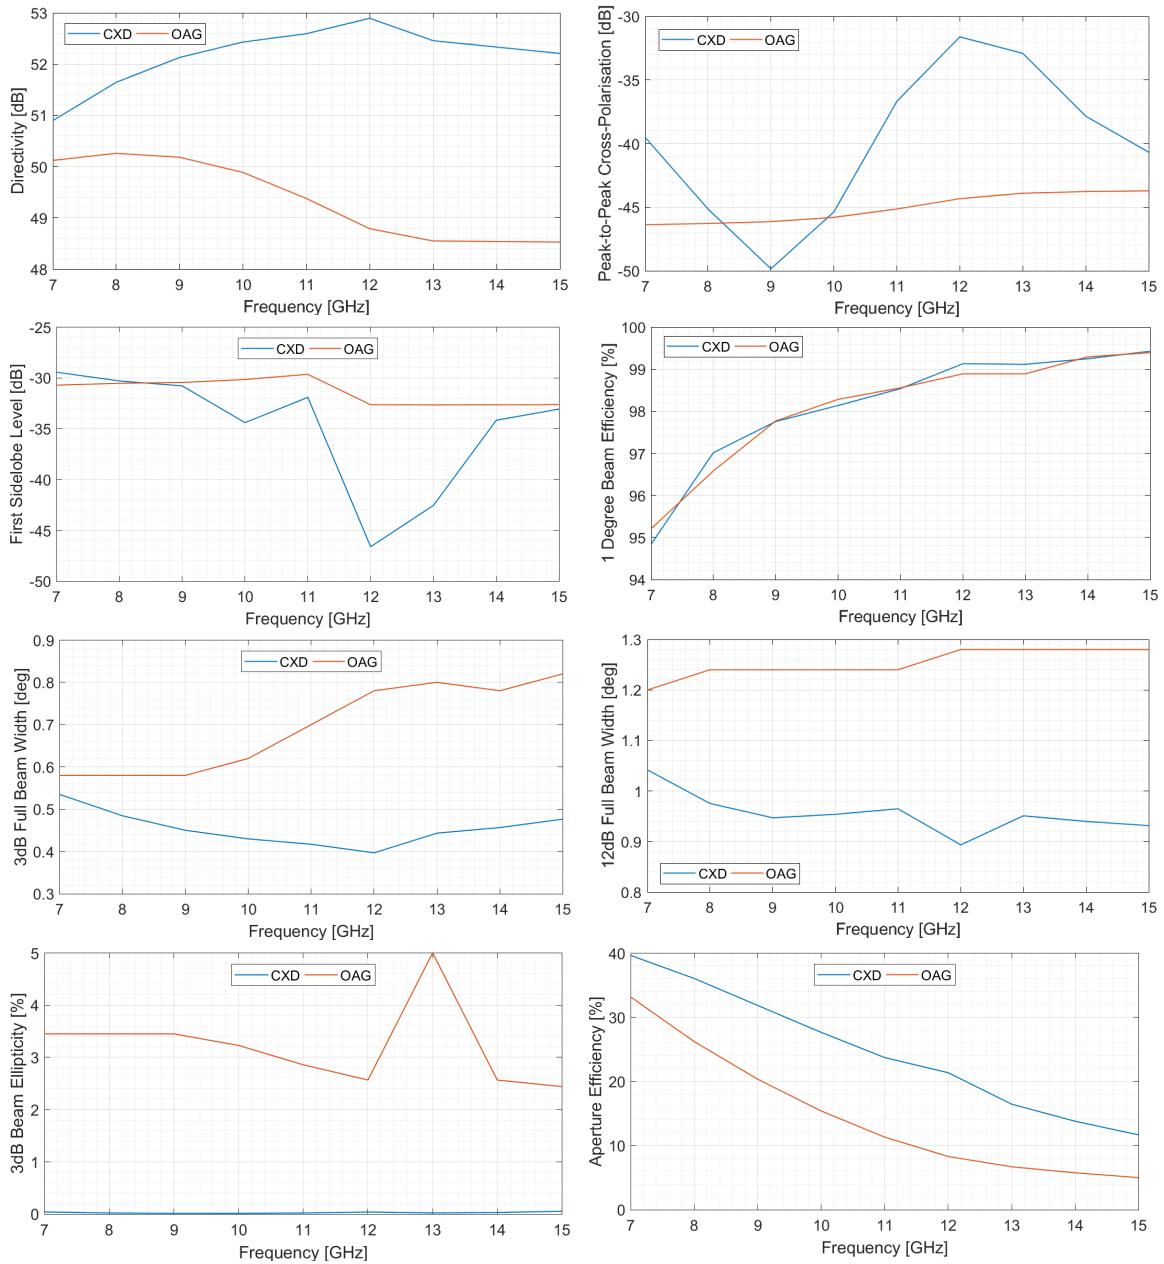


Figure 3.22: The band 1 (NextBASS¹) beam parameters for both the CXD antenna and the OAG antennas. The CXD achieves a higher directivity however does not perform quite as well in cross-polarisation. Both the CXD and OAG designs are similar in beam efficiency. The beam ellipticity is given as a percentage, such that 0% is a circular beam, and 100% is a fully elliptical beam. The CXD has a highly circular beam, however the OAG has an ellipticity of 5%.

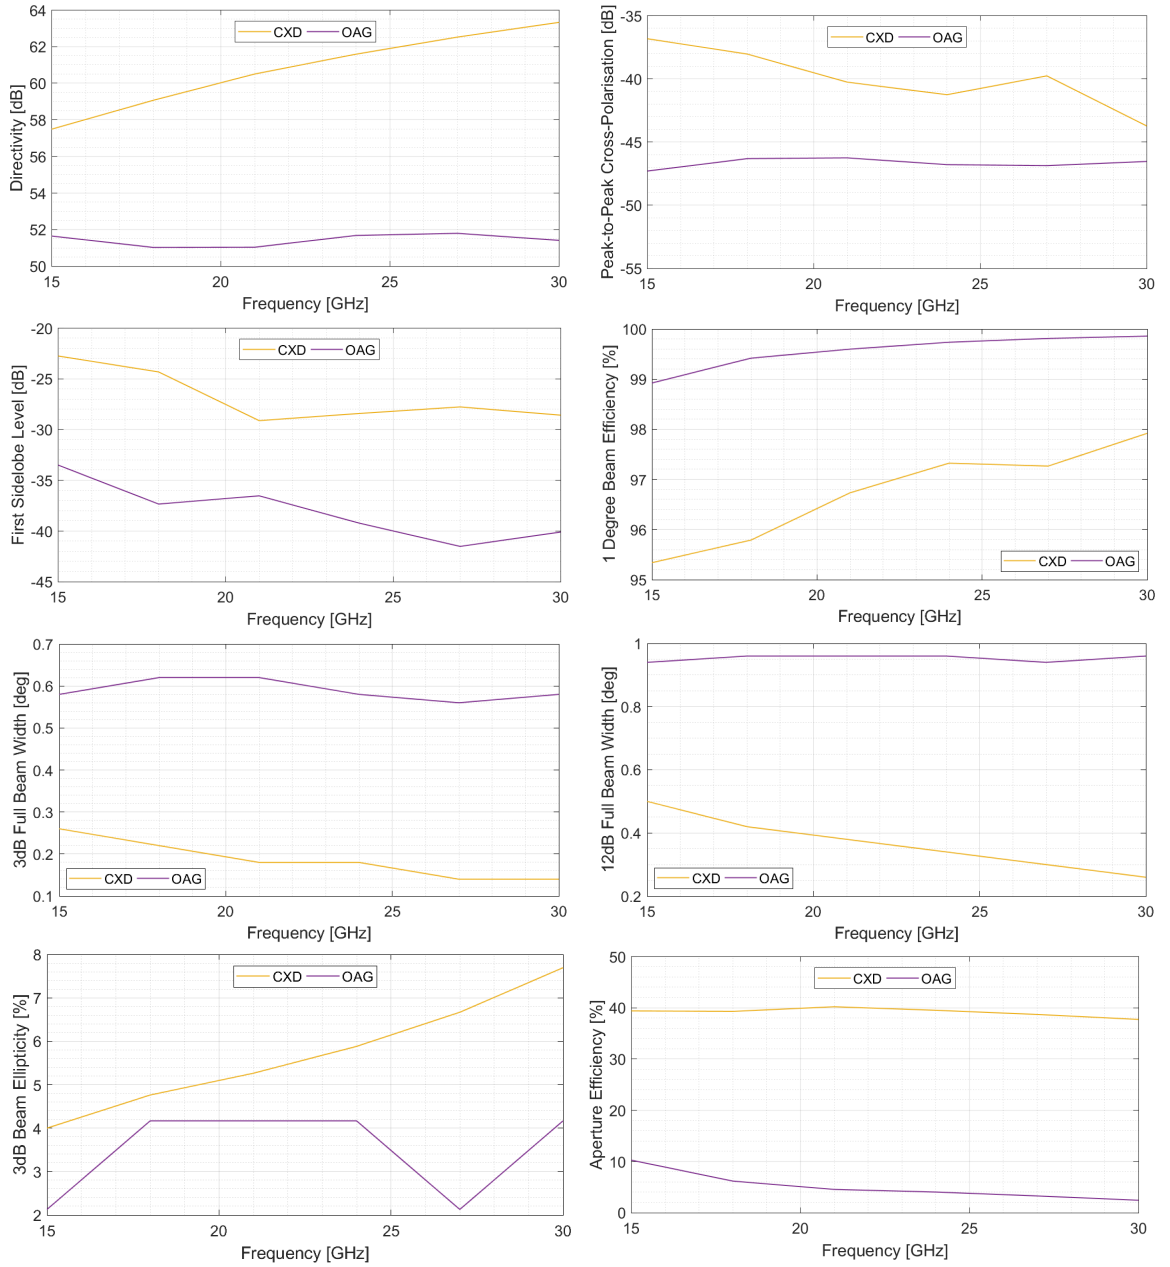


Figure 3.23: The Beam Parameters for NextBASS² (band 2) at the frequencies 15 – 30 GHz for both the CXD dish and the OAG dishes. The OAG obtains a lower cross-polarisation and first sidelobe level. The OAG beam is more elliptical than the CXD.

3.4.4 Focal Plane

At this stage, both optics have similar performances and would both be suitable for a sensitive CMB-style experiment. The next important comparison to make is of the focal plane. To investigate the shape of the focal plane, I created a grid of Gaussian

feed horns, which matched the illumination angle of the secondary mirror with a 25 dB edge taper. I spaced the grid in 50 cm intervals, generating a $1.5\text{m} \times 1.5\text{m}$ regular grid centred at the focal point of the optics. To analyse the generated fields, I created a far-field grid in the u - v plane. The position of each of the co-polar field distributions should indicate whether there is any curvature on the focal plane. The PSF of each individual beam will indicate the severity of diffraction effects at the edge of the focal plane. I investigated two feed positioning scenarios, one with all feeds perpendicular to the secondary mirror and for the other, I used the distance between the feed and the secondary mirror to derive a tilt angle, in order to point each feed towards the secondary reflector centre. Adding the tilt means that all beams are now illuminating the secondary, this should decrease distortions in the beam for the outer-edge pixels along with increasing the size of usable focal plane.

The focal point of the OAG lies on the same axis as the primary mirror. This means there is more flexibility in the positioning of a feed-array as the angular separation of the central feed and the sub-reflector is increased. When positioning the feeds that without the tilt towards the sub-reflector, the outer most beams will miss the mirror entirely. This is why it is necessary to retain the Mizuguchi condition by tilting the feeds towards the centre of the sub-reflector. A screen-grab of the OAG optics with the feeds tilted towards the centre is shown in Figure 3.24.




Figure 3.24: Screen grab of the GRASP optical configuration of the OAG telescope with each feed in the grid pointing towards the centre of the secondary reflector.

Figures 3.25 shows the co-polar and cross-polar u - v plane of the feed grid before and after tilting for the OAG. The plots clearly show that the tilt recovers some of the usable focal plane, as more rays are now hitting the sub-reflector. We also see that adding a tilt causes the focal plane to become more square and there is a slight reduction in cross-polarisation levels of the central feeds. For both the flat and tilted

feed arrangements, the OAG optical configuration displays a visible warping of the positions of the beams with noticeable asymmetries in the u - v plane. This curvature indicates that the focal plane is not flat with respect to the z -axis. For a large focal array, we can conclude that the OAG is not a suitable design as it would be difficult to position the feeds in a regular grid but with varying z positions to optimise the focal plane.




Figure 3.25: Co-polar and cross-polar u - v plane grids of the 6m off-axis Gregorian optical configuration. Before (top) and after (bottom) tilting the feeds towards the secondary mirror. By tilting the feeds, the shape of the focal plane becomes more rectangular and some beam circularity is restored. The cross-polarisation of the central feeds is also reduced.

In contrast to the OAG, it was not possible to generate a full $1.5\text{m} \times 1.5\text{m}$ grid because the bottom row of feeds become obscured by the rim of the sub-reflector. This however does not mean the CXD has a poorer response. If we look at the Figures in 3.26 we see immediately that there is significantly more recoverable focal plane than

there is for the OAG. There is barely a visible warp in the beam positions and the focal plane is clearly flat. We can see that tilting the feeds does not improve the co-polar patterns much for the central pixels, but does flatten the beam positions slightly for the feeds located at the edges. We can also see that tilting causes some feeds in the grid to retain the distinctive cross-polar two-lobe symmetry. This is beneficial as it means there is a drop in cross-polar power at the centre directly beneath the main beam. These cross-polar patterns can be seen in more detail in Figure 3.27.

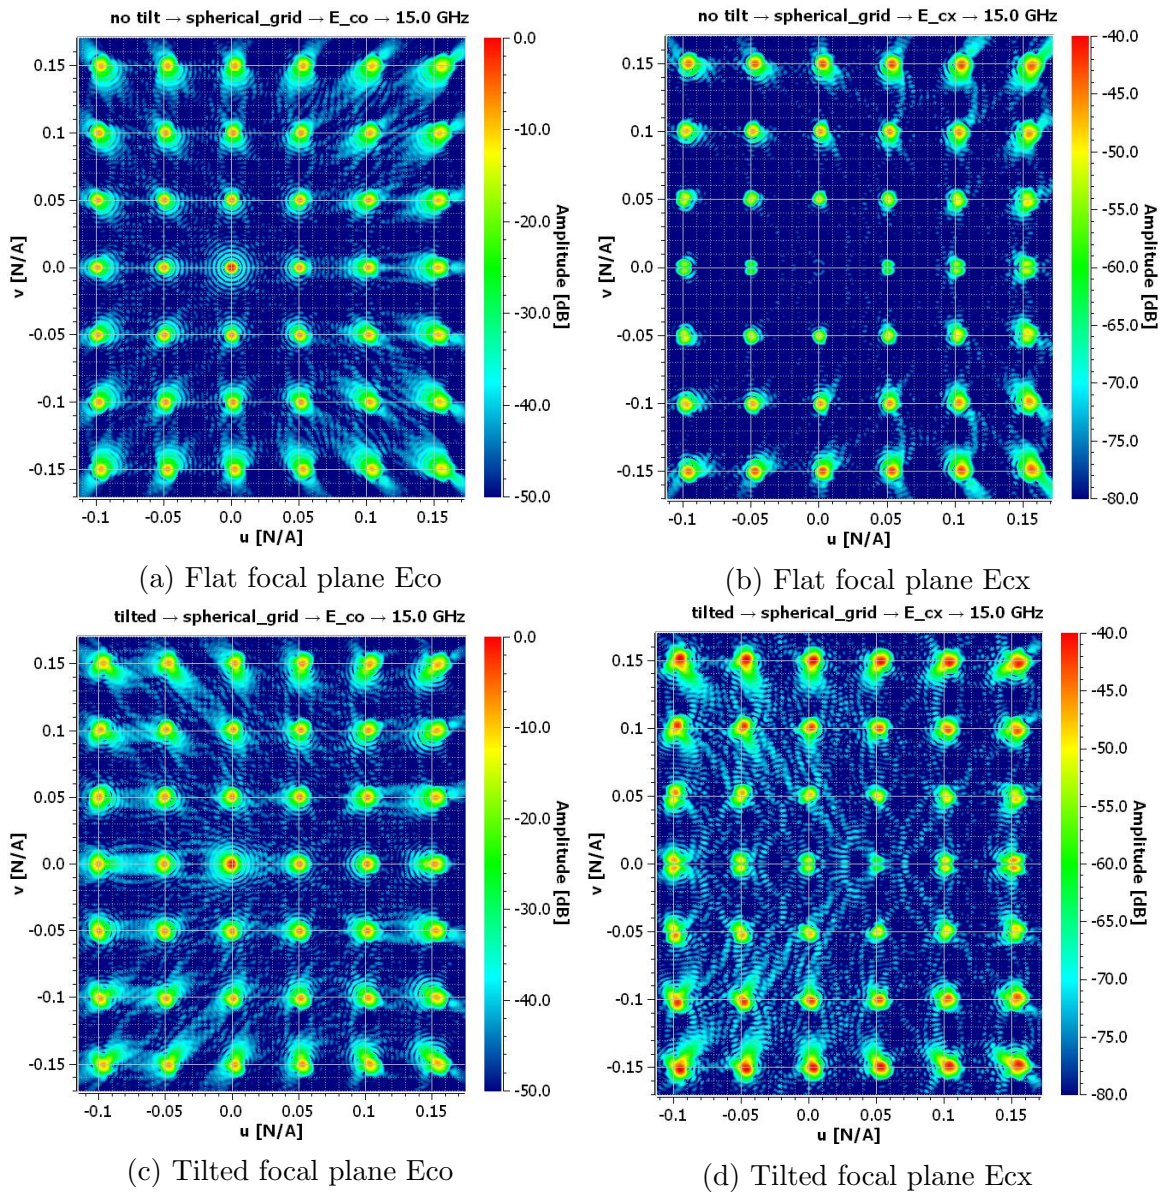


Figure 3.26: Focal plane of the 6m cross-Dragone configuration with and without tilting the feeds. The focal plane of the cross-Dragone is highly flat and the resulting focal points are evenly spaced. By tilting the feeds, there is a reduction in beam symmetry.

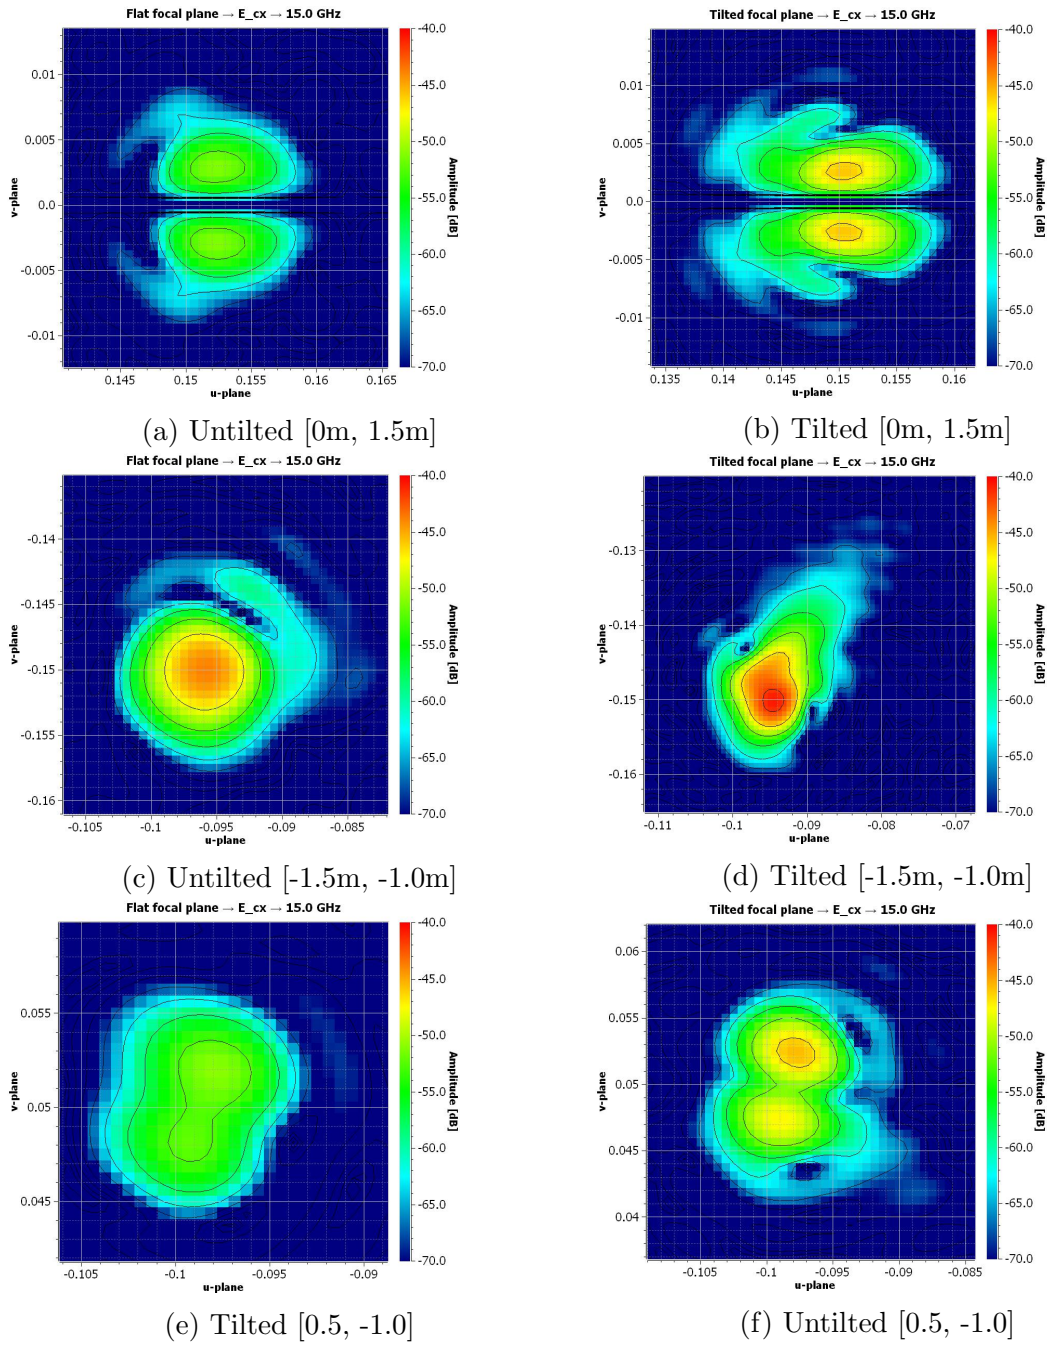


Figure 3.27: Cross-polarisation of a selection of feeds in the focal plane of the cross-Dragone telescope. The coordinates are indicated underneath each Figure. The feeds left flat in the focal-plane are given on the left-hand side, whilst the tilted feeds are on the right. The tilting resulting in the improvement of the separation of the two cross-polar lobes, reducing the cross-polar power at the centre of the beam.

To compare both the focal planes of the OAG and CXD designs, I have used the focal grids to take the maximum directivity and cross-polar peak of each individual beam to obtain the maximum peak-to-peak cross-polarisation. I have generated the

surface plots at 15 and 30 GHz on the same scale, shown in Figure 3.28. On the same scale we can see that the OAG has much poorer peak-to-peak cross-polarisation away from the central pixel. Even at the largest displacements, the CXD optics still maintains a cross-polarisation of below -30dB. There is also a less severe drop away from the central pixel along the y -displacement for the CXD design. This is due to the cross-Dragone optical arrangement having a larger sub-reflector and better symmetry than the off-axis Gregorian.




Figure 3.28: Surface plots of the peak-to-peak cross-polarisation of each beam in the focal grid outlined above. The colour scale is set to be the same for both the cross-Dragone (CXD) and off-axis Gregorian (OAG) axes, but note the 30dB offset in the vertical axis for the CXD. The plots show that there is a sharp peak of lowest cross-polarisation for the OAG, however the change in cross-polarisation more gradual towards the outer feeds for the CXD. The OAG has a much poorer cross-polarisation behaviour.

3.4.5 Optical Shielding

It is clear that the cross-Dragone-style instrument has the best potential to meet the requirements of NextBASS due to the excellent cross-polarisation performance across a wide focal-plane. The main systematic that arises in this instrument is the spillover lobes that arise from the secondary and primary reflectors. One way to control this systematic is to place the telescope in an absorbing housing. This housing is at the ambient temperature of the telescope. Any systematics will then be controlled as all

angles external to the main beam will be at a fixed temperature. This means ground calibration becomes much simpler resulting in more stable datasets. However, this comes at a risk of an increase in system temperature due to the presence of the absorbing material. In practice the material cannot be cooled but this is usually a fair trade-off for greater instrument stability.

To effectively capture the double-reflection sidelobe, a baffle may be used. This feature is located at an angle of 50° in the beam pattern, therefore the spillover angle, α can be defined as $\alpha = 90^\circ - \theta_{spill}$. If we construct a baffle to capture this radiation then the dimensions will be determined by $\tan \alpha = l/d$ where l is the length of the baffle and d is the size of the aperture, taken to be just larger than the primary mirror diameter, this is demonstrated in Figure 3.29. Using GRASP I have created a box around the optics that fits the dimensions of the mirror. Figure 3.30 shows the shape of the baffle and housing that would be necessary for the NextBASS CXD optics.




Figure 3.29: Diagram showing the calculation of the baffle length, given the spillover angle of the secondary lobe as $90^\circ - \alpha$.

To fully stabilise the outer most spillover lobes, it is necessary to house the optics in a box. I have designed a shield that will leave enough room to fit a $1\text{m} \times 1\text{m}$ feed array that is shaped to fit neatly to the mirrors (see Figure 3.30).

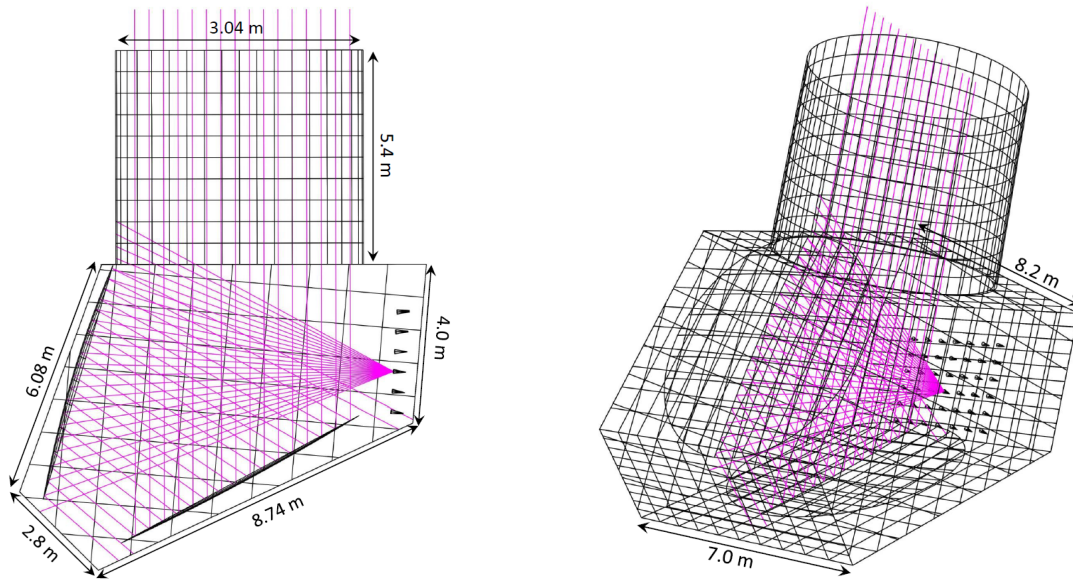


Figure 3.30: Dimensions of CXD housing and baffle for the 6 m CXD NextBASS optics. The largest panel beneath the primary mirror is 8.74m in length. To effectively remove the majority of the spillover from the double-reflection, the baffle must be 5.4m in length.

By performing the same set of GRASP commands as carried out for NextBASS previously in this section, it is possible to send the spillover lobes to the absorbing housing. Since this behaves as a perfect absorber, the difference between the beam pattern with and without the shielding will show exactly which sidelobes are being absorbed. Figure 3.31 plots the resulting beams at 7.0 GHz for the optics alone, the optics and the baffle and the optics with the full shielding. In this Figure, we can see that the addition of the baffle has effectively captured 10 dB of the secondary spillover lobe, as represented by the grey line. The black line shows the effect of the full shielding, which has captured a considerable amount of spillover from both the sub-reflector and the main reflector.




Figure 3.31: Comparison of three scenarios, the NextBASS optics alone, with the addition of a 5.4m length baffle and finally with the baffle and surrounded by a perfect absorbing box. These cuts were taken at 7.0 GHz, since the lowest frequencies are susceptible to poorer spillover. The baffle (grey) effectively captures the double-reflection at 50° (the red beam), the full shielding absorbs the majority of the secondary spillover (black line).

The beam pattern produced by GRASP for the shielding does not represent the true beam pattern one would get from the shielded optics. Instead, when compared to the beam pattern of the un-shielded optics, the difference in the power as a function of θ shows how much power is collected by the shielding. By calculating the beam efficiency of the shielded optics, we can estimate how much power is sent to the baffle or shielding, and hence find the effective efficiency of the shielding. Table 3.4 gives the beam efficiency (η_{beam}) of each stage of shielding calculated within a 1° beam. The difference between the beam efficiency of the shielded optics and the that of the un-shielded optics gives the shielding efficiency (η_{shield}). The remaining power left over gives how much residual power will pick up the ground signal.

Scenario	η_{beam}	η_{shield}	η_{ground}
Un-shielded optics	94.84%	0%	5.16%
Un-shielded optics + baffle	95.17%	0.33%	4.83%
Full optical shield	98.83%	3.99%	1.17%

Table 3.4: The effective beam efficiency within a 1° beam and the effective shielding and ground efficiencies given as η_{shield} and η_{ground} respectively. These efficiencies have been calculated such that $\eta_{\text{shield}} = (\eta_{\text{beam}}^{\text{shield}} - \eta_{\text{beam}}^{\text{optics}})$ and $\eta_{\text{ground}} = [1 - (\eta_{\text{beam}}^{\text{optics}} + \eta_{\text{shield}})]$. We can see that the shielded system has only 1.17% of power on the ground compared to the 5.16% of the un-shielded optics.

We can see that the shielded system captures $\approx 4\%$ of the power, with only $\approx 1.2\%$ of power remaining on the ground signal. This demonstrates that there is an increase in stability of the shielded system, as 4% of the power that was originally on the unstable ground, is now sent to the optical shielding. This offers an increase in stability in trade-off with a potential increase in antenna temperature. Since the shielding may be at an ambient temperature which is greater than the ground temperature, the overall antenna temperature will be higher. A higher antenna temperature means poorer sensitivity, I discuss this relationship further in Chapter 6.




Figure 3.32: Cumulative power of the CXD optics at increasing levels of optical shielding. The baffle captures from the main beam up to 70.0° , where as adding the full shielding captures all angles up to 130° where the majority of the power has now been sent to the cut.

Calculating these efficiencies does not show the angular response of the shielding, which can help isolate which sidelobes are captured by the shield and which will still retain some ground signal. The cumulative power curve for each level of shielding can demonstrate which angles are being captured by the optical shield. Figure 3.32 shows

the cumulative power at 7.0 GHz. We can see that the baffle has captured spillover between the main beam and 70° , but all angles after this match the standalone optics. The full shielding however, captures all angles between the main beam to 130° where there is very little power left in the beam.

3.5 Optical Design Review

In conclusion to this chapter, the X-BASS optics have shown to be suitable up to 15 GHz when including surface errors in the GRASP simulations. For NextBASS, the CXD design is better than the OAG due to its large, flat focal plane. The CXD optics can be placed in a ground-shield with an absorbing baffle, significantly reducing any systematics caused by ground temperature variations, at the trade-off of potentially increasing the antenna temperature. I discuss the antenna temperature further in Chapter 6. Both X-BASS and NextBASS optical arrangements have shown to perform up to a high standard meeting the requirements set out at the start of this chapter without further system optimisation. In Chapter 4, I will outline the optimised feed horn designs used in conjunction with the optical simulations in this chapter.

Chapter 4

The X-BASS and NextBASS Feed Horns

“ A clever mind is not a heart. Knowledge doesn't really care, wisdom does. ”

Benjamin Hoff, The Tao of Pooh

The optics for X-BASS and NextBASS achieve high beam efficiencies and low cross-polarisation across their respective frequency bands. The performance of the antennas must be maximised with the appropriate choice of feed horn. In this chapter, I will describe the process of designing the X-BASS and NextBASS feed horns. The choice of feed horn depends on a large range of criteria. The main criteria required for the two designs in question, is that the feed horns must have low cross-polarisation and low sidelobes consistent for the full bandwidth. The type of feed horn that best matches these criteria is the corrugated horn, the feed horns designed in this chapter will be based on corrugated horn design principles that I discuss at the start of this chapter.

The C-BASS optics adapted for X-BASS require an optimised horn to achieve low sidelobes and a low cross-polarisation. The shaped secondary is in near-field distance of the feed-horn, making horn matching more complex. Initially, I designed a preliminary horn of roughly the right length and aperture size to achieve a 40 dB edge illumination for X-BASS. Using the TICRA corrugated horn modelling software CHAMP ([TICRA \(2015a\)](#)), I ran an optimisation using a cubic spline interpolation to change the horn profile to match the beam-width of the original C-BASS feed horn beam pattern to ensure the correct edge illumination on the primary reflector. I also optimised the horn design to minimize the cross-polarisation. To check the final horn design, I used the ANSYS High Frequency Simulation Software (HFSS) to analyse the X-BASS horn with HFSS.

For the NextBASS band-1 and band-2 feeds, I outline a recipe for far-field horn design, using Gaussian optics to match the beam-waist of the feed horn to the beam-waist of an equivalent Cassegrain that describes the 6m cross-Dragone detailed in Chapter 3. I generated design curves to compare aperture-size and slant-lengths that yield a 25 dB edge taper on the primary reflector. I also compared different profiles to demonstrate that, without further optimisation, a hyperbolic profile achieves a suitably low cross-polarisation and consistent beam-widths across the band. I have also included a brief discussion about the band-2 feed array and ways of using circular packing to configure the feeds in the focal plane.

4.1 Feed Horn Principle

In this section I describe the basic principles behind using feed horns as receivers and compare the different types of feed design. A feed horn couples a propagating electromagnetic wave in free-space to a wave guide. It is placed at the focal point of a paraboloidal reflector and is the first part of the receiver. After the waveguide, a transducer is used to generate an electrical signal which is coupled to the electronics. The feed horn is designed to couple the correct modes to the optics. They must be tapered or have a surface that can smoothly transform the free-space wave into the desired modes, in particular the design must prevent generating higher order modes as these cause degradation of the beam.

4.1.1 Mode Generation

A feed horn transforms the propagating wave in a waveguide into the desired modes. A horn will support two orthogonal mode sets. For waveguides with perfectly conducting, smooth walls that are either hollow or filled with a homogeneous material, the mode sets consist of 'transverse magnetic' (TM) with $H_z = 0$ and 'transverse electric' (TE) and $E_z = 0$. There are infinite numbers of possible propagating modes, but the number of finite modes depends on the waveguide size relative to wavelengths as imposed by the boundary conditions of the waveguide walls. These finite modes are the designated TM_{nm} and TE_{nm} modes where n and m are integers and refer to the solutions of the wave equation in two directions.

If the cross-section of the waveguide is inhomogeneous, eg. corrugated, the TM and TE modes are coupled together and all field components can exist. The modes are now called 'hybrid modes'. The two hybrid modes that exist are orthogonal and are known as the HE_{nm} and EH_{nm} modes.

Pure mode horns radiate at a single mode; the radiation patterns are almost entirely determined by the transverse fields of the mode, not by the surrounding structure. The dominant mode in the waveguide determines the single mode of the horn. We can separate pure mode horns into their size, such that ‘small-size’ correspond to aperture diameters of $a < 1.5\lambda$, and ‘medium-size’ horns have $1.5\lambda < a < 6\lambda$. Small-size pure mode horns depend on the precise geometry of the horn as currents outside the horn can contribute to the radiation field. These are typically used in small diameter front-fed reflector antenna applications that require a reduced mass and volume, such as satellite antennas. Medium-size pure mode horns are used as separate antenna elements instead of a single receiver due to poor radiation patterns. Medium-size Pyramidal horns, for example, have poor radiation patterns but have reliable gain estimates due to the presence of only one dominant mode.

The boundary condition for TM modes in a cylindrical waveguide is such that $\phi_{TM} = 0$, as a result the dominant mode in a cylindrical waveguide is always the TE mode. The propagating transverse electric fields form the aperture electric fields E_a . When used in conjunction with the Fourier-transform integrals, these fields can be used to get the radiation pattern. These aperture fields in the x and y directions can be described by the following:

$$f_x = \int_{x'} \int_{y'} E_{ax}(x', y') \exp jk(x' \sin \theta \cos \phi + y' \sin \theta \sin \phi) dx' dy' \quad (4.1)$$

$$f_y = \int_{x'} \int_{y'} E_{ay}(x', y') \exp jk(x' \sin \theta \cos \phi + y' \sin \theta \sin \phi) dx' dy' \quad (4.2)$$

The radiation integral in x and y directions $f_x(\theta, \phi)$ and $f_y(\theta, \phi)$ can be used to determine the far-field E_θ and E_ϕ components, where θ and ϕ refer to the polar coordinates of the propagating EM wave.

$$E_\theta = f_x(\theta, \phi) \cos \phi + f_y(\theta, \phi) \sin \phi \quad (4.3)$$

$$E_\phi = \cos \phi f_y(\theta, \phi) \cos \phi - f_x(\theta, \phi) \sin \phi \quad (4.4)$$

The far-field co-polar and cross-polar fields (E_{co} and E_{cx} respectively) are related to the spherical field components using Ludwig’s third definition of polarisation (Olver et al. (1994)).

$$\begin{bmatrix} E_{co}(\theta, \phi) \\ E_{cx}(\theta, \phi) \end{bmatrix} = \begin{bmatrix} \sin \phi & \cos \phi \\ \cos \phi & \sin \phi \end{bmatrix} \begin{bmatrix} E_\theta \\ E_\phi \end{bmatrix} \quad (4.5)$$

E_{cx} will only vanish if the terms associated with f_x and f_y cancel over all values of θ and ϕ . Since both f_x and f_y must be non-zero for this to happen, the transverse electric fields in the aperture must have both x -directed and y -directed components. This is not the case in pure mode horns. For a hybrid mode horn, at the boundary imposed by the corrugations, we can write the impedance and admittance as X and Y respectively as (Olver et al. (1994)):

$$X = -j \frac{E_\phi}{H_z} y_0 \quad (4.6)$$

$$Y = +j \frac{H_\phi}{E_z} y_0 \quad (4.7)$$

where y_0 is the admittance of free-space. If $(X - Y) = 0$, the aperture field is independent from the angular variable ϕ , as a result, $E_y = 0$ meaning no cross-polarised field exists. This happens if X and Y are either finite or both zero. By corrugating the inside of a feed horn, the corrugations cause X and Y to become zero at the wall of the horn. This is described as a ‘balanced hybrid-mode’ because the electric and magnetic fields are exactly balanced to produce pattern symmetry and low cross-polarisation. The field lines must exhibit some degree of curvature as a purely linear aperture electric field will not radiate zero cross-polarisation (Clarricoats and Olver (1984), Olver et al. (1994)).

4.1.2 Types of Feed Horn

The four main types of feed horn are pyramidal, conical, corrugated and ridged. Pyramidal horns, as already mentioned, have a tapered pyramid shape with a rectangular cross-section, the antenna pattern is asymmetric. Conical horns generate curved E and H -plane wave-fronts. Corrugated horns contain evenly spaced grooves known as corrugations, which suppress any presence of higher order modes. The spacing of corrugations remains constant but the depth is gradually tapered from $d = \lambda/2$ at the throat of the horn, to $d = \lambda/4$ at the aperture. The tapering enables gradual generation of hybrid modes whilst the $\lambda/4$ slot depth cancels out any higher order modes. A flat and uniform distribution of the electric field across the aperture is produced, with curvature only appearing at the edges. They are sensitive to all polarisations due to their circular symmetry, making them an excellent choice for polarisation experiments. Finally, ridged horns use specially shaped, optimised ridges to control the generation of higher order modes. Commonly these come as either double or quad-ridged horns. They are often used in broad-band systems requiring a bandwidth of up to 3 octaves. In Figure 4.1 I show these feed horn designs diagrammatically.

The common Pyramidal and Conical horns are simple in concept but do not have good electrical characteristics as they are pure mode horns. Conical horns have a symmetric radiation pattern but poor cross-polarisation due to field curvature. Pyramidal horns have asymmetric beam patterns, they have poor cross-polarisation and high gain, making them unsuitable for a sensitive CMB foregrounds experiment.




Figure 4.1: Types of Feed Horn: The four different types of feed horn, each with the appropriate waveguide. The field lines of each horn are indicated on the aperture. The Pyramidal horn uses a rectangular waveguide and generates a linear field. The Conical horn has a cylindrical waveguide and generates a curved wavefront. Corrugated horns also use cylindrical waveguides, due to their $\lambda/4$ corrugations, the resulting wavefront is flat with curvature near the edges. Corrugated horns have low cross-polarisation and good beam patterns. The Ridged horn consists of profiled ridges that can be optimised to achieve a desirable beam pattern.

Ridged horns are often used in surveillance and radar where a single antenna is required to transmit or receive over a wide frequency. The specially shaped ridges of varying tapers achieve this. They are usually double ridged however the quad-ridge can be used to achieve dual or circular polarisation. The taper of the ridges determines the radiation pattern by varying the cut-off positions for higher order modes. A good polarisation sensitivity is possible when optimising a quad-ridge design. The main drawback of ridged horns is the asymmetric beam pattern; this makes them an unsuitable choice for a foregrounds mapping instrument as the angular power distribution will not be uniform across the beam.

Corrugated horns generate fairly constant beam widths for each frequency in the band. This is because the modal interactions are well-behaved due to the basic

symmetry of the radiation pattern. Since the field curvature is not precisely linear, this provides the desirable properties of axial beam symmetry, low sidelobes and low cross-polarisation, making corrugated feed horns the optimum choice for the X-BASS and NextBASS receivers.

4.2 Broad-band Corrugated Horn Design

As outlined above, the corrugated feed horn is capable of low cross-polarisation due to the cancellation of higher order modes and an excellent polarisation sensitivity due to symmetrical E and H wave generation. They are also capable of achieving a 2:1 bandwidth and can be profiled to optimise the beam. Below are some of the design principles required to design a broad-band corrugated feed ([Granet and Graeme \(2005\)](#)).

4.2.1 Waveguide and Mode-Converter

A broadband feed horn can be defined as having a frequency range $f_{max} \geq 2.4f_{min}$. To propagate the beam at the full range of frequencies, the input radius must satisfy the equality $a_i \frac{2\pi f_{min}}{c} \geq 1.841$ where c is the speed of light. More generally the input radius is given as:

$$a_i = \frac{3\lambda_c}{2\pi} \quad (4.8)$$

where λ_c is the central wavelength and is equal to $1.2 \frac{c}{f_{min}}$.

The first part of the corrugated horn after the waveguide is the mode-converter. The input waveguide is excited by the pure TE_{11} mode meaning a mode-converter is required to produce the HE_{11} mode. This is done over a specified number of slots. There are three types of mode converter. A variable-depth-slot mode converter is the most commonly used, this gradually tapers to the central wavelength, $\lambda_c/4$ slot depth. A variable-pitch-to-width-slot mode converter is less commonly used, it starts with a high pitch-to-width ratio (tall, narrow corrugations) and gradually decreases to shorter and broader corrugations. This type of mode-converter suits bandwidths of up to $f_{max} \leq 2.05f_{min}$. Finally, a ring-loaded slot mode-converter has the best response for broadband applications, performing well up to $f_{max} \leq 2.4f_{min}$. The ring-loaded slot mode converter has been the choice for all feed horns designed for NextBASS and X-BASS. The slot width and depth can be optimised to improve the suppression of higher order modes. The typical pitch-to-width ratio for broadband horns is $\lambda_c/10$. Figure 4.2 shows the structure of a ring-loaded-slot mode converter.




Figure 4.2: A ring-loaded slot mode converter: a_i indicates the waveguide radius, a_0 is the aperture radius, N_{MC} is the number of slots in the mode converter and L is the length of the horn. The slots gradually change from having narrow corrugations into a ring of corrugation width $\lambda/4$, to the standard corrugation width. This type of mode converter is suitable for applications such that $f_{max} \leq 2.4f_{min}$.

4.2.2 Profiling

Profiling is a method of feed horn design where the horn is shaped in such a way that achieves optimal beam properties. For a simple corrugated feed horn, the profile might be linear, but depending on the bandwidth and application, different profiles might be more suited. In Figure 4.3, I outline some of the profiles that can be used for horn design.




Figure 4.3: Examples of exaggerated horn profiles. Profiling can be used to improve the performance of a feed horn. The hyperbolic profile is very close to linear however the taper at the throat results in better cross-polarisation. The flare of a tangential horn profile can generate good broadband beam patterns, however worse cross-polarisation than the hyperbolic horn. The asymmetric sine profile has a poor broadband response but achieves high aperture efficiency due to the straightened taper at the aperture of the horn.

The hyperbolic profile gives a frequency consistent beam pattern across the band and has been the choice of horn profile for both the band-1 and band-2 NextBASS feed horns. The equation that describes the aperture diameter, $a(z)$, of a hyperbolic

horn at displacement z along the z -axis is given by:

$$a(z) = \sqrt{a_i + \frac{z^2(a_o - a_i)^2}{L^2}}, \quad (4.9)$$

where a_i is the input radius of the waveguide, a_o is the output radius, z is the displacement from the apex and L is the total length of the horn.

4.2.3 Phase Centre

The phase centre of the feed horn can vary significantly across the band. It is important to choose the correct phase centre because you want to achieve a wave-front that is in phase at the aperture for all frequencies in the band. This can cause beam de-focusing when used with the telescope optics.

Generally, a horn does not possess a unique phase-centre. This is particularly problematic for broadband applications. CHAMP employs a best-fit method to find a phase centre for each frequency specified in the analysis (TICRA (2015a)). Since the generated horn geometry is symmetric about an axis, the phase centre lies only on the z -axis. I have used CHAMP to identify the phase centre location by observing the maximum displacement across the band for each frequency and choosing the most central displacement. In the corrugated horn designers handbook Granet and Graeme (2005), they state that the phase centre of a horn can be estimated by:

$$\text{Phase Centre} = \alpha L = L \left(1 - \exp \left[-4.8 \left(\frac{k_c a_o^2}{4\pi L} \right)^2 \right] \right) \quad (4.10)$$

where α is a dimensionless quantity that describes the fractional distance from the throat to the aperture, such that $0 \leq \alpha \leq 1$, where $\alpha = 0$ is the throat of the horn and $\alpha = 1$ is the aperture. L is the length of the feed horn, k_c is the wavenumber at the design frequency and a_o is the aperture radius of the horn.

In cases where the feed horn may not be in the far-field of the sub-reflector, a more rigorous testing of the phase centre must be carried out. An investigation of this sort has been done to determine the best phase centre position for the X-BASS feed.

4.2.4 CHAMP Analysis

The CHAMP software (see TICRA (2015a)) allows for the design and analysis of waveguides and feed horns. It is particularly suited to designing corrugated feeds and has been used to aid design of the horns in this thesis. I will briefly outline the

methods used in its analysis algorithm and describe the optimisation process that was carried out for the X-BASS feed horn. The frequencies in the analysis are equally spaced between f_s and f_e .

$$f_i = f_s + (i - 1) \frac{f_e - f_s}{N_f - 1} \quad (4.11)$$

The excitation can be in either the fundamental TE_{11} mode or other tracking modes for example TM_{01} , TE_{21} or TE_{01} . The fundamental mode has been used for the analyses in this chapter. For linear polarisation in the x -axis, the far-field can be represented as:

$$\begin{aligned} E_\theta(\theta, \phi) &= f_\theta(\theta) \cos \phi \\ E_\phi(\theta, \phi) &= -f_\phi(\theta) \sin \phi \end{aligned} \quad (4.12)$$

The polarisation ratio can be defined along the axis as $p = p_x = E_y/E_x$ where $|p_x| \leq 1$.

$$\begin{aligned} E_\theta(\theta, \phi) &= f_\theta(\theta)(\cos \phi + p_x \sin \phi) \\ E_\phi(\theta, \phi) &= -f_\phi(\theta)(\sin \phi - p_x \cos \phi) \end{aligned} \quad (4.13)$$

The output parameters include return loss, on-axis directivity, aperture efficiency and best fit phase centres. The return loss for the fundamental mode can be given by (TICRA (2015a)):

$$RL = -20 \log_{10} \left(\frac{|b_1(TE_{11})|}{|a_1(TE_{11})|} \right), \quad (4.14)$$

where a_1 is and b_1 are the amplitudes of the incident and reflected modes at the throat of the horn respectively.

The CHAMP optimisation uses a ‘minmax’ algorithm to optimise features of the horn. The user specifies the horn design parameters, such as length, L , or aperture size, a_o , to optimise along with the type of optimisation to be run. The different types of optimisation possible are optimising the return loss, the horn aperture efficiency, the cross-polarisation levels relative to the co-polar peak, directionality of the beam, the power in the main beam for a given frequency, the phase centre location and the co-polar and cross-polar pattern templates. For the X-BASS feed, I have optimised the peak-to-peak cross-polarisation to below 40 dB, and the co-polar pattern template to match the C-BASS beam.

4.3 Far Field Horn Design

The electromagnetic field from a radiating element is given by its diffraction properties. When calculating the diffraction properties of an antenna or feed, we can define the Fraunhofer distance as $d_f = 2D^2/\lambda$, where D is the aperture size and λ is the radiating wavelength (Hecht (1997)). The near-field radiation pattern is calculated at distances closer than the Fraunhofer distance and the far-field radiation patterns are given at positions greater than this distance. For an feed horn aperture diameter of 2.5λ , which applies to the C-BASS and X-BASS feed horns, the Fraunhofer distance $d_f \approx 12.5\lambda$. At a wavelength of approximately 5cm, this distance is about $d_f \approx 625\text{mm}$. In the case of NextBASS cross-Dragone optics, the distance between the sub-reflector and the feed is much greater than this distance. However, in the C-BASS South optics, the sub-reflector is located 500mm from the focus, meaning the feed is functioning in the near-field regime. This means the performance of the lower end of the X-BASS channel will exhibit somewhat worse beam characteristics than the higher frequencies, which are closer to the far-field regime. This has been compensated for by optimising the profile of the X-BASS feed horn to match the beam pattern to that of the original C-BASS 5 GHz horn.

This is not a problem for the NextBASS feed horn meaning we can use a far-field design approach, however since the wavelengths are less than a tenth of the physical antenna reflector aperture sizes, a better description of these optical systems than using ray optics, is to apply quasi-optical, or Gaussian, optical design techniques. In this section, I describe the Gaussian optical principles used to model an equivalent system of the NextBASS antennas and use beam matching to establish a feed horn aperture size and length for a given edge taper.

4.3.1 Gaussian Optics

Radio, millimetre and sub-millimetre antenna systems produce electric field distributions over an aperture that is usually greater than or equal to a wavelength in size. To achieve a high antenna efficiency, the electric field distribution across the aperture must be fairly uniform with a significant edge illumination to maximise the amount of collection power. However, at long wavelengths, these fields are not usually uniform and therefore considerable diffraction effects are present and we find that classical ray-optics no longer offer the best description.




Figure 4.4: Diagram of a Gaussian beam where the focal point is the gaussian beam waist, the radius of curvature at some distance r and the confocal distance where the beam waist increases by a factor of $\sqrt{2}$, also indicated is the divergence angle.

Instead, Gaussian beams can be used to describe wave propagation in an optical system. A Gaussian beam is monochromatic. Its electric and magnetic field amplitudes are given by a Gaussian irradiance profile. The fundamental mode of the propagating electromagnetic wave, TEM_{00} , well describes the behaviour of a Gaussian beam. When the beam undergoes a transform by any optical element in its path, a new Gaussian beam with a different phase is produced. The field amplitudes of a Gaussian beam for a given wavelength and polarisation can be suitably described by the Gaussian beam waist, w_0 at a position, z relative to the focus.

Plane waves can describe electromagnetic wave propagation where the distribution of field amplitudes is independent of position along the propagating wave-front. When the plane wave undergoes diffraction, we can treat the plane-wave as a largely collimated Gaussian beam. This is described by the paraxial wave equation. The axially symmetric paraxial wave equation in cylindrical coordinates is given by:

$$\frac{\partial^2 u}{\partial r^2} + \frac{1}{r} \frac{\partial u}{\partial r} - 2jk \frac{\partial u}{\partial z} = 0. \quad (4.15)$$

This has the solution:

$$u(r, z) = A(z) \exp \left[\frac{-jkr^2}{2q(z)} \right] \quad (4.16)$$

where A and q are two complex functions of z . These can be determined by substituting the solution back into the paraxial wave equation (Goldsmith (1993)). Doing this, we find that q describes the complex beam parameter, the real part of which gives the radius of curvature, R , of the beam. These parameters are defined in Figure 4.4.

$$\left(\frac{1}{q} \right)_r = \frac{1}{R}. \quad (4.17)$$

The three equations that fully describe the nature of a Gaussian beam are given by equations 4.18, 4.19 and 4.20. The beam waist at distance, z , along the axis, $\omega(z)$ is located at the minimum of two divergent rays. The beam radius, $R(z)$ describes the curvature and hence the beam size as it propagates along the axis. (Goldsmith (1993))

$$\omega(z) = \omega_0 \left[1 + \left(\frac{\lambda z}{\pi \omega_0^2} \right)^2 \right]^{0.5} \quad (4.18)$$

$$R(z) = z + \frac{(\pi \omega_0^2)^2}{z \lambda^2} \quad (4.19)$$

$$z = \frac{R}{1 + \left[\frac{\lambda R}{\omega(z)^2} \right]} \quad (4.20)$$

An important definition of a Gaussian beam is the confocal distance, z_R . This describes separation between the beam waist and a position located where the beam cross-section has doubled (see Figure 4.4). This is also known as the Rayleigh length and is give by equation 4.21.

$$z_R = \frac{\pi \omega_0^2}{\lambda} \quad (4.21)$$

4.3.2 Edge Taper

In Radio astronomy, when observing a point source, the resulting illumination on the primary reflector is much less than the total area available. This is such that the total illumination area is given by $A_{geom} = \pi r^2$, where r is the radius of the primary mirror. If observing an extended source, then the full telescope aperture is completely filled. This means that any power illuminating the outer dish-edge could cause additional un-wanted noise. To reduce this noise, one can under-illuminate the dish, this means an overall loss in power to compensate for an increase in sensitivity. This can be described in Gaussian optics as the edge taper of a system (Goldsmith (1993)).

The edge taper, T_e , is the relative power at a radial distance, r_e . The edge taper in decibels is given by $T_e(dB) = -10 \log_{10}(T_e)$ where the edge taper T_e is defined as the ratio of the power at the edge of the aperture to the power at the centre of the aperture, $T_e = \frac{P(r_e)}{P(0)}$. In terms of the radial distance, r_e and the beam waist ω_a this becomes:

$$T_e(r_e) = \exp\left(\frac{-2r_e^2}{\omega_a^2}\right). \quad (4.22)$$

The edge radius of the beam is obtained from the edge taper using the ratio of radial distance r_e to the aperture beam waist ω_a . If we consider this in terms of some truncation at the aperture edge, it is common to write the edge taper of the optics as:

$$T_e(dB) = 8.686 \left(\frac{r_e}{\omega_a}\right)^2. \quad (4.23)$$

A dual-reflector system can be simplified and described as an equivalent paraboloid. The ideal condition that must be satisfied is that the output beam of this paraboloid has a planar wave-front. We can treat the antenna system as a single focusing element of equivalent focal length f_e , this has been described for a dual-reflector antenna in Equation 3.5. The equivalent Gaussian beam waist of such an optical system is given by:

$$\omega_{0illum} = \frac{\lambda f_e}{\pi \omega_a} \quad (4.24)$$

where ω_a is the beam waist that meets the planar wave-front condition. This is defined such that $z = f_e$ (Goldsmith (1993)).

The waist position shifts between the optimum waist location and the focal point of the dual-reflector system. This shift, Δz , is negligible for large antennas that satisfy $\pi \omega_a^2 / \lambda \gg f_e$. We can say that a Gaussian beam located at the telescope focus with beam waist ω_{0illum} will be well focused. The illumination waist of the telescope can then be written in terms of the edge taper in dB as:

$$\omega_{0illum} = 0.216 (T_e(dB))^{\frac{1}{2}} \frac{f_e}{D_m} \lambda_c. \quad (4.25)$$

where D_m is the diameter of the main reflector and λ_c is the central wavelength of the band. We can then define the divergence angle of this ideal Gaussian optical configuration using the equation:

$$\theta_0 = \frac{\lambda}{\pi \omega_{0illum}}. \quad (4.26)$$

The full-width-half-maximum (FWHM) of such a beam is hence given by $\theta_{FWHM} = 1.18\theta_0$.

4.3.3 Beam Matching

From the definition of a Gaussian beam, a corrugated feed horn's beam can be estimated as a Gaussian by the following two equations. A good feed design must be well matched to the optical configuration it will be used with. I have matched the illumination beam waist of the horn aperture to the beam waist I achieve at the focal point of telescope with a given edge taper and λ_c , the central wavelength of the feed horn. The illumination beam waist of the optical configuration has been described in Equation 4.25. When the Gaussian beam waist of the feed horn matches the beam waist of optics, then the feed and optics are well-matched and the feed horn should provide the desired edge taper. We can write the beam waist of the horn in terms of the horn aperture radius a_o as:

$$\omega_0^{horn} = \frac{a_o \rho}{\sqrt{1 + \left(\frac{\pi \rho^2 a_o^2}{\lambda R_h}\right)}}, \quad (4.27)$$

$$z = \frac{R_h}{\left(1 + \left[\frac{\lambda R_h}{\rho^2 a_o^2}\right]\right)^2} \quad (4.28)$$

where ω_o^{horn} is the horn beam waist at the phase centre, R_h is the horn slant length and ρ is some design constant, which for corrugations $\rho = 0.644$ (Goldsmith (1993)).

4.3.4 Gaussian beam approximation

The first stage in horn matching is to match the feed beam waist to the beam waist of the optical configuration for a given edge taper. This can be done by treating the system as an equivalent paraboloid, as described in the previous section. From equation 4.25, I have found the beam waists of X-BASS and the Cross-Dragone NextBASS telescopes. This beam waist will change over the band, so I have used the design frequency of the corrugated horn as $1.2f_{min}$ where f_{min} is the lowest frequency in the band (Granet and Graeme (2005)). The factor of 1.2 comes from the limitation that a corrugated broadband horn can perform well up to a maximum bandwidth of $2.4f_{min}$ and therefore, for a broadband application, the central frequency can be taken to be half this value. I have also included the angle subtended by the feed to the sub-reflector as this is important for the X-BASS system as the secondary mirror is in the near-field of the horn.

In Table 4.1 I list the edge tapers chosen for the X-BASS and NextBASS optics and the feed half angle that fully illuminates the secondary mirror. I have also calculated

the Gaussian beam waist of each optical system which is given by ω_{0illum} and hence the divergence angles that must match that of the feed horn if the horn and optical systems are to be well-matched. The edge taper of X-BASS is up to 40dB because the system does not have a method of reducing ground spillover so under-illuminating the primary reflector prevents larger spillover lobes. The NextBASS optics will be shielded and hence can benefit from larger edge tapers. I found an edge taper of 25 dB to perform best when designing the feed horns.

Telescope	f_c	Edge Taper	Feed Half angle	ω_{0illum}	θ_{0illum}
X-BASS	8.4 GHz	40dB	46.3°	28.51mm	26.96°
NextBASS ¹	8.4 GHz	25dB	25°	63.45mm	10.27°
NextBASS ²	18.0 GHz	25dB	25°	29.61mm	12.11°

Table 4.1: This Table shows the design edge taper, feed half angle, the Gaussian beam waist and the resulting feed horn FWHM for X-BASS and both bands of the NextBASS experiment.

4.4 An Optimised Horn for X-BASS

In the previous section I mentioned two factors that are limitations of the optical system of C-BASS that results in difficulties in the design of a feed horn. These are the fact that the sub-reflector is in the near-field of the horn and that a large edge taper has to be chosen in order to reduce the effects from ground-spillover. The designs in this section reflect an optimised feed that matches the C-BASS optics with an edge taper of 40 dB.

4.4.1 Recreating C-BASS horn design

I started with the 2.5λ aperture diameter and 9.7° semi-flare which was used in the design of the C-BASS horn. I created a hyperbolic profile horn to match these specifications at a design frequency of 8.4 GHz. This was chosen to match the $f_c = 1.2f_{min}$ required for a broadband corrugated horn, where $f_{min} = 7.0$ GHz for X-BASS. Using CHAMP I generated the beam pattern of the 2.5λ horn design. Using the resulting beam pattern, I then tested this horn with the full optical model of X-BASS in GRASP to check both the beam pattern and the edge illumination on the primary mirror. I found that the edge illumination on the primary mirror was close to -60 dB, which is far too high of an edge illumination for X-BASS.

I then redesigned the horn in CHAMP with the same semi-flare angle but now with an aperture radius of 2.0λ . In the GRASP simulations of the 2.0λ horn, I found that

this aperture size introduced higher spillover lobes from the secondary, resulting in an edge illumination of -30 dB, which is under the target edge illumination of -40 dB. Repeating this method of trial and error, I found that the best choice of aperture size to use for the X-BASS horn was 2.2λ , that without profile optimisation, gave close to the required primary edge taper of -40 dB, which is necessary to minimise ground spillover for the X-BASS antenna.

4.4.2 Optimisation

I then began a routine of optimisations. The X-BASS horn must ideally have a consistent edge illumination across the band, have low sidelobes, a cross-polarisation target of 40 dB and preferably a design that means the FWHM is not frequency dependent. This is difficult to achieve and requires careful consideration into the optimisation method. The optimisation routines must also be ordered in a way that will not cause drastic changes to the aperture size, length or semi-flare angle of the horn as we know that these parameters have already been well matched with the optics. I began with a hyperbolic horn profile that matched the aperture sizes mentioned in the previous section, I then added 10 points on the horn profile such that the optimisation routine could perform a cubic spline interpolation between each point on the surface of the horn to change the profile to match the optimisation aims.

By looking at the edge taper plots across the band, I was able to create a feed horn radiation pattern template that would ensure a -40 dB edge illumination. Using this pattern as the optimisation goal, it was possible to run the optimiser to change the surface profile to achieve this goal across the band. I then optimised the cross-polarisation to -40 dB across the band. I began the optimisation routines on the hyperbolic profile, investigating profiling with different combinations of optimisation goals.

I performed an optimisation of the feed horn using the radiation pattern template and the cross-polarisation goal of -40 dB, both separately then consecutively. I found that when both optimisation goals were run consecutively, the resulting feed horn gave a better performance in first sidelobe level and cross-polarisation, than when the optimiser only considered one of these goals.

I also investigated improving the aperture efficiency of the horn. If the aperture efficiency can be optimised across the band, then a similar illumination is achieved at each frequency and therefore a more frequency independent edge illumination would be possible. I optimised the aperture radius and horn length to match an aperture efficiency of 70%. This initial optimisation also meant that the following pattern

optimisation ran faster since the beam sizes at each 1 GHz frequency interval were similar in size.


The different optimisation routines are shown in Table 4.2, the surface points refers to an optimisation changing the position of the cubic spline points. The column that specifies the previous stage indicates the previous optimisation routine that was carried out and acts as the starting profile for the next optimisation. The ordering of these optimisations can change the profile in different ways.

Label	Parameter	Optimisation Aim	Previous stage
Opt _{CX}	surface points	Cross-polarisation, minimisation, 40dB	Hyperbolic horn
Opt _{PT}	surface points	Pattern template, target ideal beam	Hyperbolic horn
Opt _{PTCX}	surface points	Cross-polarisation, minimisation, 40dB	Opt _{PT}
Opt _{Aeff}	aperture size and length	Aperture efficiency, maximise, 70%	Hyperbolic horn
Opt _{AeffPT}	surface points	Pattern template, target ideal beam	Opt _{Aeff}
Opt _{AeffPTCX}	surface points	Cross-polarisation, minimisation, 40dB	Opt _{AeffPT}


Table 4.2: This Table gives the nomenclature and types of optimisations trialled for the X-BASS feed horn. The parameters that have been optimised has either been the cubic spline surface points placed on the profile of the horn, or the aperture size and horn length. The Optimisation Aim lists the target goals set in the CHAMP optimiser. Finally, the previous stage indicates the starting horn profile used for that optimisation.

From the results in Figure 4.5 it is clear that the best performing X-BASS horn is the Opt_{PTCX} optimised horn (the optimisation process is outlined by Opt_{PTCX}). This gives consistently low sidelobes and cross-polarisation across the band. It is not frequency independent but when considering the weighting of constraints, the low sidelobes and low cross-polarisation are more necessary so the selection of this horn will mean slightly under illuminated beams at the upper end of the band. ¹


¹A common processing technique performed in CMB map making is smoothing the beams to a fixed beam size or resolution, which would imply that the extra resolution (or smaller beam sizes) at high frequencies is not of concern. However, maintaining a consistent beam-size across the band can ensure a similar edge illumination across the band. This constraint on the design of the feed is less important than achieving low first sidelobes or a low cross-polarisation.




(a) Edge Illumination of Primary



(b) Maximum Peak-to-peak Cross-polarisation



(c) First Sidelobe Peak



(d) 12dB beam width

Figure 4.5: Feed horn optimisation routines and their performance with the realistic C-BASS optical model. The results in these Figures are the Primary Edge Illumination, the Maximum Peak-to-Peak Cross-polarisation, the first sidelobe peak and the 12dB beam width. A good choice of horn for X-BASS needs an edge illumination of approximately 40 dB, a cross-polarisation below at least 35 dB across the band, low sidelobe levels and a fairly uniform beam width across the band. Achieving all of these targets is difficult, but it is evident that the PTCX optimised horn gives the best cross-polarisation and first-sidelobe level across the band, at a trade-off to a frequency dependent beam width.

The horn profiles of the resulting optimisations are given in Figure 4.6. The design that performed best is given by horn c. in this Figure. The rest of this section will focus on the analysis of this feed horn. Remarking on the resulting profiles, it is interesting that in order to achieve an improved cross-polarisation, it is the flare angle at the throat of the horn that changes drastically, whereas when optimising the horn to match the pattern of C-BASS, the profile curvature changes more across the length of the horn. The final design resembles a straight waveguide with only some

departure of linearity after the mode converter and before the aperture. The aperture efficiency optimised horn does not give good beam properties for a broadband design.

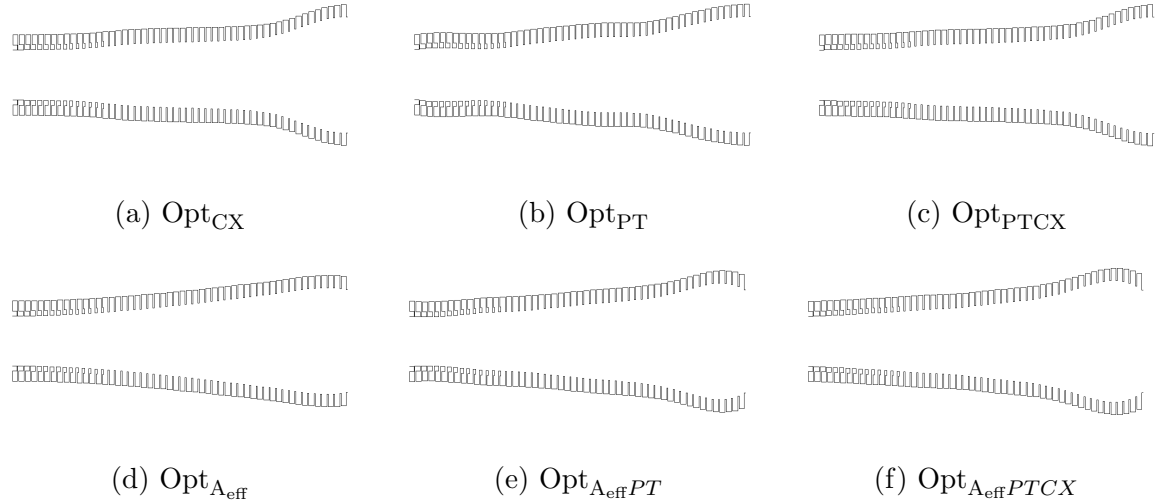


Figure 4.6: The resulting horn profiles of different optimisation routines of the X-BASS feed horn design. The Figures are: a. optimised hyperbolic horn to obtain a cross-polarisation of 40 dB; b. optimised hyperbolic horn to match C-BASS pattern template; c. optimised cross-polar horn to match C-BASS pattern template; d. optimised hyperbolic horn to achieve high aperture efficiency; e. aperture efficiency and pattern template optimised horn and f. aperture efficiency, pattern template and cross-polarisation optimized horn. The final X-BASS horn is horn c.

4.4.3 CHAMP results

Figure 4.7 shows the outcome of the CHAMP parameters. This includes the return loss, which is sufficiently low across the band, the on-axis directivity, the peak-to-peak cross-polarisation, the phase centre location and the beam width at half power. The phase centre location is important as it indicates the best position of the phase centre, this is required in positioning the feed at the best focal point to achieve the best performance with the optics. This will vary considerably over a broad bandwidth and I investigate this later in Section 4.4.5. The phase centre varies by 30 mm between 8.0 GHz and 15.0 GHz, considering the length of the horn is $\approx 228\text{mm}$, the 30mm phase variation will not cause considerable defocusing of the optics.

In Figure 4.7, (c), there is a spike in the peak-to-peak cross-polarisation. This may be due to a resonance in the horn that causes it to break down between 7.0 and 8.0 GHz. Cross-referencing this beam feature with part (d) of Figure 4.7, we can see that the spike is present for the $\phi = 0.0^\circ$ and $\phi = 45.0^\circ$ cuts however absent in the $\phi = 90.0^\circ$ cut. The ϕ cuts represent the direction of the radiation pattern coming

out of the horn, since the resonance is at a maximum at the 0.0° cut, the orthogonal 90.0° undergoes a minimum of this feature.

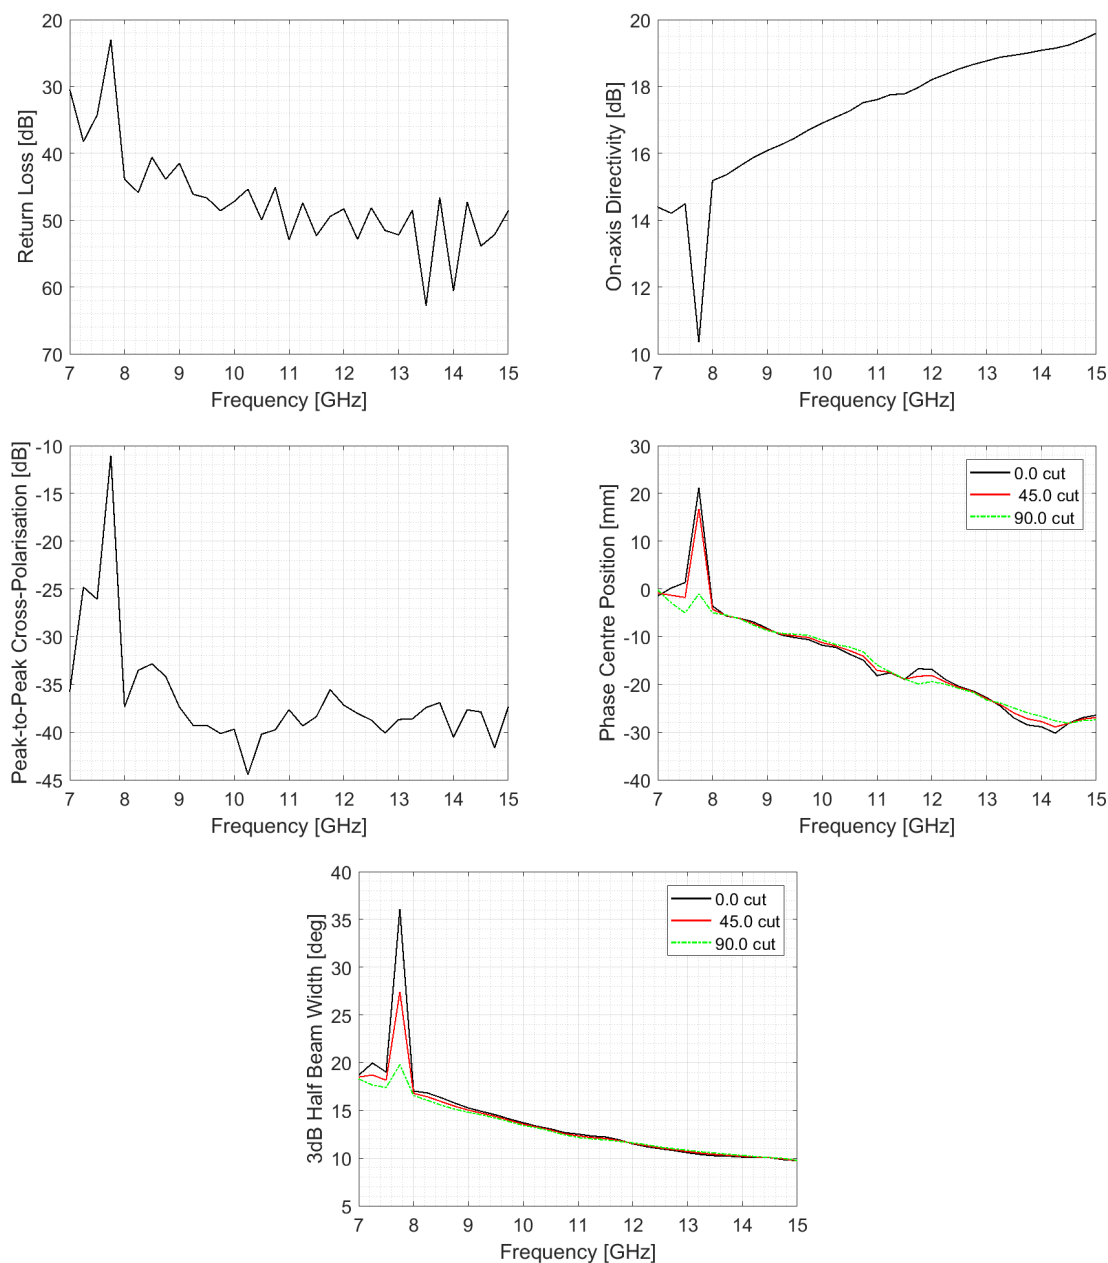


Figure 4.7: The design output of the CHAMP analysis of the X-BASS Feed horn. This shows an excellent return loss below 20 dB across the band. These results were taken at steps of 0.25 GHz. The majority of the band has an excellent cross-polarisation, however there is a spike at around 7.75 GHz where the performance drops. This may be due to some resonance within the feed. CHAMP does not have sophisticated enough modelling to identify where this problem comes from. None the less, the majority of the band is high performing and meets the requirements needed for X-BASS.

In Figure 4.8, I have plotted the co-polar and cross-polar beam patterns of the X-BASS horn. The cross-polarisation is at least 40 dB across the band. The black line indicates the template I have used to optimise the beam pattern. This shows that the 7.0 and 8.0 GHz frequencies are well matched to the pattern. The presence of the pattern has also suppressed the location of the first sidelobes for the higher frequencies.






Figure 4.8: X-BASS Beam Pattern - generated in CHAMP with the pattern template used to optimise it superimposed in black. For all intervals of 1 GHz between 7.0 and 15.0 GHz, the cross-polarisation is at 40 dB across the band.

4.4.4 Analysis in HFSS


To investigate the feed performance further, I have modelled the feed horn in HFSS. This will achieve a more accurate co-polar beam pattern than CHAMP, however due to the method that HFSS meshes surfaces, there will be some asymmetries in the cross-polar pattern. In Figure 4.9, I have plotted the HFSS beams at 1GHz intervals across the band. The cross-polarisation performance is good and the 0.0° , 45.0° and 90.0° cuts are well aligned. The CAD horn profile is given in Figure 4.10.




(a) 7.0 GHz




(b) 8.0 GHz



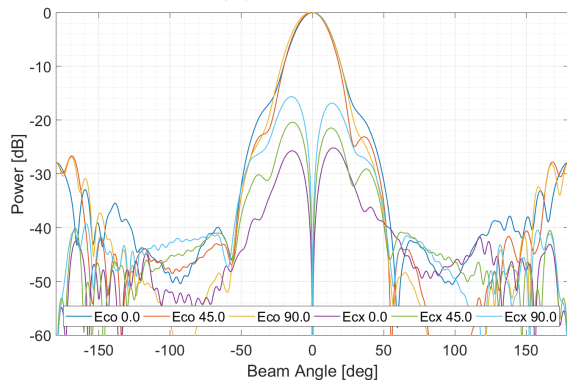
(c) 9.0 GHz



(d) 10.0 GHz



(e) 11.0 GHz



(f) 12.0 GHz




Figure 4.9: X-BASS horn HFSS Analysis of 0.0, 45.0 and 90.0 cuts for both co and cross-polarisation. The asymmetries in the cross-polarisation pattern are a result of the HFSS method of calculating the propagating modes, however the resulting amplitudes of the cross-polar peaks are well representative of what the true horn behaviour should be like.




Figure 4.10: Model of horn Opt_{PTCX} horn HFSS. An HFSS analysis is carried out by creating a circular waveguide attached to the horn throat and creating a horn surface as a perfect conductor. HFSS then propagates the electromagnetic field through the horn and generates a radiation pattern from the propagating modes.

4.4.5 Phase Centre

The best fit phase centre for a broadband horn is difficult to find. Each frequency will have a phase centre location slightly different within the horn. The phase centre location begins at the aperture and works its way back towards the apex of the horn, this is the same as moving the horn closer to the optics. I have found that the higher frequencies favour phase centre locations further towards the apex of the horn. In Figure 4.11, I have changed the position of the feed horn from 8 mm to 20 mm in intervals of 4 mm. I have demonstrated that the low frequency channels favour phase centres closer to the aperture, however it is the higher frequencies that suffer more from a poor choice of phase centre. Towards the upper end of the band, the sidelobes become larger if the feed is not in focus for that frequency. As a result, I recommend a phase centre between 14 mm – 16 mm, that will not cause shallow first nulls in the low frequency beams and will not cause an increase in first sidelobe gain for the higher frequencies.

4.4.6 Performance with Optics

I have tested the horns performance with the optics and find that the main beam achieves a cross-polarisation of below -35 dB across the band. Even when modelled with the real optics, the asymmetric nature has not hindered the performance of the main beam and the sidelobe levels are still low. In Figure 4.12 I have plotted the main beam of the X-BASS horn modelled with the real dish surface.

4.4.7 Edge Tapers

To ensure that the feed design will be suitable for the C-BASS optics, I have taken surface cuts of the primary and secondary mirrors. These have been taken with the ideal optics modelling to ensure a symmetric beam. The primary edge taper is between -40 dB to -50 dB for most of the band, indicating that the feed will successfully under illuminate the dish, achieving low spillover.

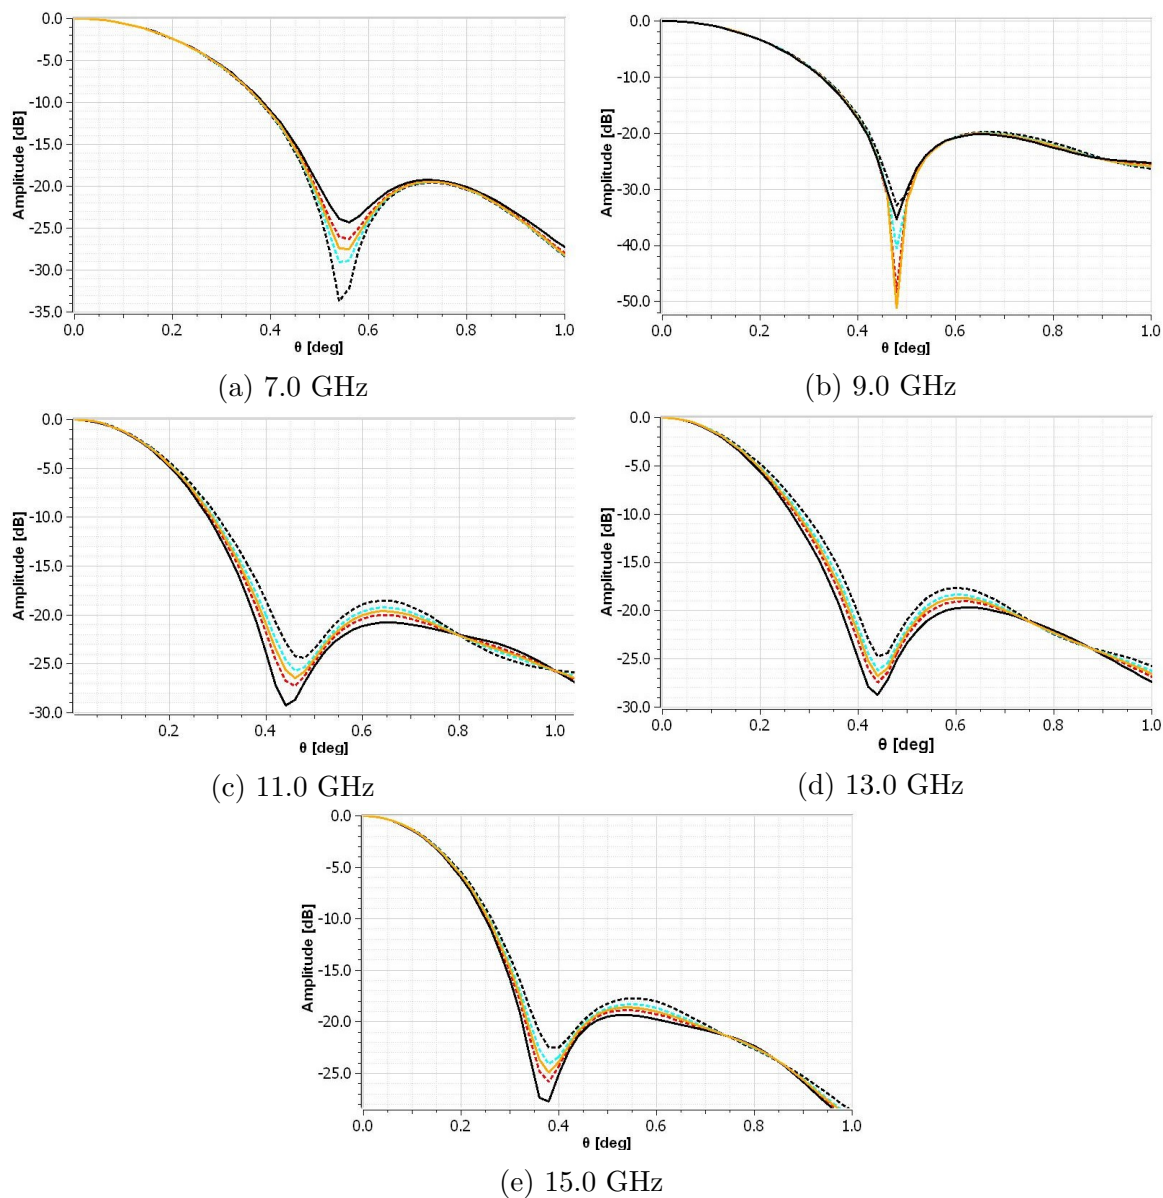


Figure 4.11: Phase centres of feed: black dashed = 8.0 mm, cyan dashed = 12.0 mm, orange solid = 14 mm, red dashed = 16 mm, black solid = 20 mm. The lower frequencies have poorer nulls at larger phase centres, however the sidelobes increase in the middle and upper end of the band. The best location is towards middle of the band (≈ 14 mm). Note the different y -axes to show how the beam nulls and first sidelobes change with frequency.

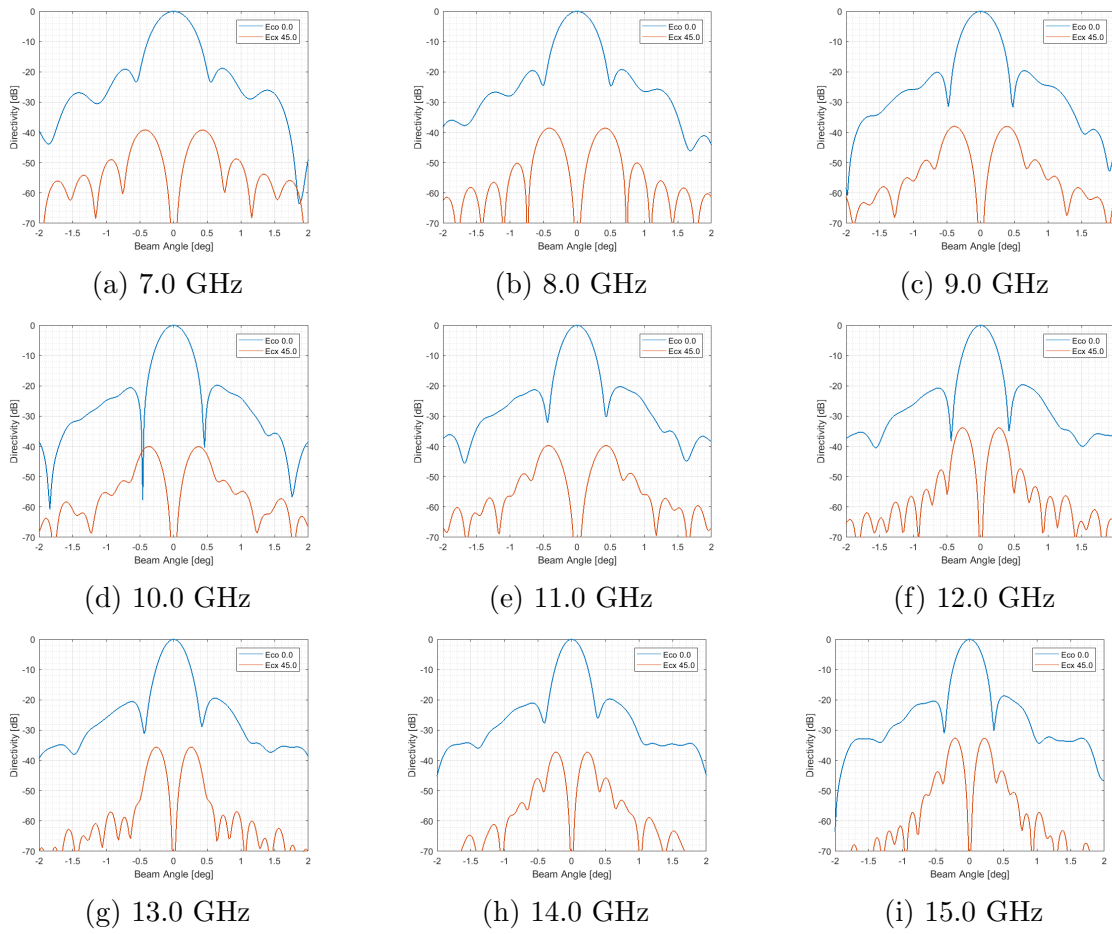


Figure 4.12: X-BASS main beam performance with the C-BASS antenna using the photogrammetry data to model the surface distortions of the primary reflector. The maximum peak-to-peak cross-polarisation is below 35 dB across the band. Using a phase centre of 14 mm, each beam exhibits sharp first nulls meaning the horn position is well focused.





Figure 4.13: Edge Taper on the primary and secondary reflectors using the X-BASS Opt_{PTCX} feed horn with the C-BASS GRASP model containing surface distortions. The edge tapers vary between -40 dB to -60 dB, and -20 dB to -40 dB for the primary and secondary reflectors respectively. The original C-BASS feed horn was designed for the 6.1 m C-BASS North primary mirror and X-BASS is being used with the larger, C-BASS South 7.6 m primary, meaning the edge tapers are well matched to the C-BASS feed, and should achieve similar levels of edge illumination to C-BASS South.


4.5 The NextBASS Feed Horns

4.5.1 Design Curves

Using the Gaussian beam matching method and Equation 4.25, I have developed design curves for different slant lengths and aperture sizes. This let me determine the best semi-flare angle to use. In Figure, 4.14, we can see that for a 25 dB taper, we require a semi-flare angle of between $14^\circ - 16^\circ$, which is on the boundary between a narrow and wide-angle feed horn. I then choose a range of length and aperture sizes matching that semi-flare angle and edge taper. Figure 4.15 shows similar design curves for the 15 – 30 GHz frequency range.




(a) Aperture radius vs. edge taper




(b) Semi-flare angle vs. edge taper

Figure 4.14: Design curves for 7 – 15 GHz at a design frequency of 8.4 GHz for a 2:1 bandwidth. The plots show how semi-flare angles and aperture size vary with the design edge taper.



(a) Aperture radius vs. edge taper




(b) Semi-flare angle vs. edge taper

Figure 4.15: Design curves for 15 – 30 GHz at a design frequency of 18 GHz for a 2:1 bandwidth. The plots show how Semi-Flare angles and aperture size vary with the design edge taper.


4.5.2 Profiles

I carried out a brief investigation on profiling using the three types of profile outlined in Figure 4.1. I chose to test these profiles as they have been referenced in papers as a good choice for broad-band feed horns. Figures 4.3b, 4.3c and 4.17 give the beam patterns from 7.0 GHz to 15.0 GHz of the tangential, asymmetric sine and hyperbolic horn profiles respectively. I found that the asymmetric sine profile causes defocusing at frequencies higher than the design frequency, making it unsuitable for a 2:1 bandwidth. The tangential profile performed well across the band with low

sidelobes. However the broad beamwidth at 7.0 GHz could become problematic when used with the CXD optics. We know from the previous Chapter that the lower frequencies are more susceptible to spillover, and that the CXD optics have a sidelobe caused by the secondary at $\approx 40.0^\circ$. Ideally, we want the beam to stay narrow with minimal power in the sidelobes. The tangential horn has 10 dB more power at 40.0° than the hyperbolic profile, which suggests it would cause larger problems with spillover when used with the full NextBASS optics.



(a) Asymmetric Horn Profile



(b) Tangential Beam Profile

Figure 4.16: The Asymmetric and Tangential horn profile beam patterns for the NextBASS 7 – 15 GHz feed horn designed for a horn of aperture size matching a 25 dB taper on the primary.

4.5.3 Band-1 corrugated horn design

To find a suitable feed horn design, I began by using the horn matching method outlined above. Using a 25 dB edge taper I established a selection of horn slant lengths and aperture radii that matched the required full-width half maximum. Using CHAMP, I generated these horns, allowing the CHAMP wizard to calculate the appropriate number of slots and the necessary slot-depth. As demonstrated broadband feeds perform best with a ring-loaded slot and a hyperbolic profile.

I then performed an analysis of 1 GHz interval steps between 7.0 – 15.0 GHz to investigate the performance of each feed. Once the best design had been chosen, I used an optimisation process to refine the length and aperture size of the horn to achieve a cross-polarisation of -40 dB. I then tested this horn with the optical arrangement using the CXD telescope outlined in the previous chapter.

Using the phase centre location generated by the CHAMP analysis, I specified a range of positions to try with the optics. I found that it is best to use a phase centre that favours the lower frequencies of the band, as they tend to deteriorate more than the higher frequencies in terms of beam circularity and cross-polarisation. The final dimensions of the hyperbolic feed horn for band-1 and band-2 is given in Table 4.3.

Horn Parameter	NextBASS ¹	NextBASS ²
Length	698.506 mm	349.253 mm
Aperture size	149.310 mm	74.655 mm
Input aperture	17.041 mm	8.520 mm
Mode converter	ring-loaded slots	ring-loaded slots
No. Slots	157	157
No. slots in MC	16	16
Profile	hyperbolic	hyperbolic
width-to-pitch ratio	0.8	0.8
Design Frequency	8.4 GHz	18.0 GHz
Output Frequency	9.24 GHz	19.8 GHz
Phase Centre	250 mm	150 mm

Table 4.3: Specifications of the hyperbolic corrugated feed horn design for both band-1 (NextBASS¹) and band-2 (NextBASS²). The horns are almost scaled versions of one another as they are designed for the same optical configuration and only shifted in frequency. The output frequency represents the frequency that defines the slot depth of the last corrugation such that slot depth = $\lambda_o/4$.

In Figures 4.17 and 4.18, I present here the horn beam pattern and analysis of

the final feed horn design for a horn that would be suitable for band 1 of NextBASS. The hyperbolic horn has a maximum cross-polarisation of ≈ -35 dB at a frequency of 9.0 GHz and the return loss is below -40 dB for the majority of the band. Since it is not priority to have a return loss this low, future optimisation routines may be able to compromise return loss for a better cross-polarisation. There is only a 1.5° variation in beam width at the 3dB point, guaranteeing a frequency dominated resolution across the band.

The performance of this horn with the symmetrical and low cross-polarisation CXD optical configuration demonstrates its suitability for a future foregrounds experiment. The beam patterns have been taken at a phase centre location of 250 mm, as this gives the sharpest nulls after the first sidelobe. In Figure 4.19, we see that the beam width does not change significantly across the band, there is only some small variation in beam shape.

I list the beam properties and performance with the optics in Table 4.4, including an estimate of the beam power in a 1° beam and a calculation of the aperture efficiency. I have also included the beam-ellipticity parameter to demonstrate the excellent circular and symmetric response of the horn and CXD optical configuration.

ν [GHz]	Gain [dB]	E_{CX} [dB]	1^{st} SL [dB]	Θ_{3dB} [deg]	A_{eff} [%]	sky_{eff} (1°) [%]	e^{3dB}
7.0	50.90	-39.50	-29.45	0.53	39.6	94.82	0.0004
8.0	51.64	-45.14	-30.31	0.48	36.0	97.04	0.0190
9.0	52.13	-49.84	-30.79	0.45	31.8	97.76	0.0001
10.0	52.43	-45.27	-34.41	0.43	27.6	98.15	0.0001
11.0	52.59	-36.68	-31.92	0.42	23.7	98.56	0.0002
12.0	52.89	-31.61	-35.39	0.40	21.3	99.14	0.0004
13.0	52.46	-32.91	-34.95	0.44	16.4	99.13	0.0002
14.0	52.33	-37.84	-34.17	0.46	13.8	99.25	0.0003
15.0	52.21	-40.67	-33.07	0.48	11.7	99.43	0.0005

Table 4.4: The NextBASS¹ feed horn results at the focus of the cross-Dragone optics. The hyperbolic feed horn achieves an exceptionally high performance in terms of its low first sidelobes and cross-polarisation. The beam also mains approximately one-half degree resolution across the band. The analytical methods used to derive results have been explained in Chapter 3, Section 3.2.2.




Figure 4.17: Beam pattern of the Band-1 horn design, generated in CHAMP. The beam pattern is similar in beam width for all frequencies in the band and the cross-polarisation is ≈ 40 dB.

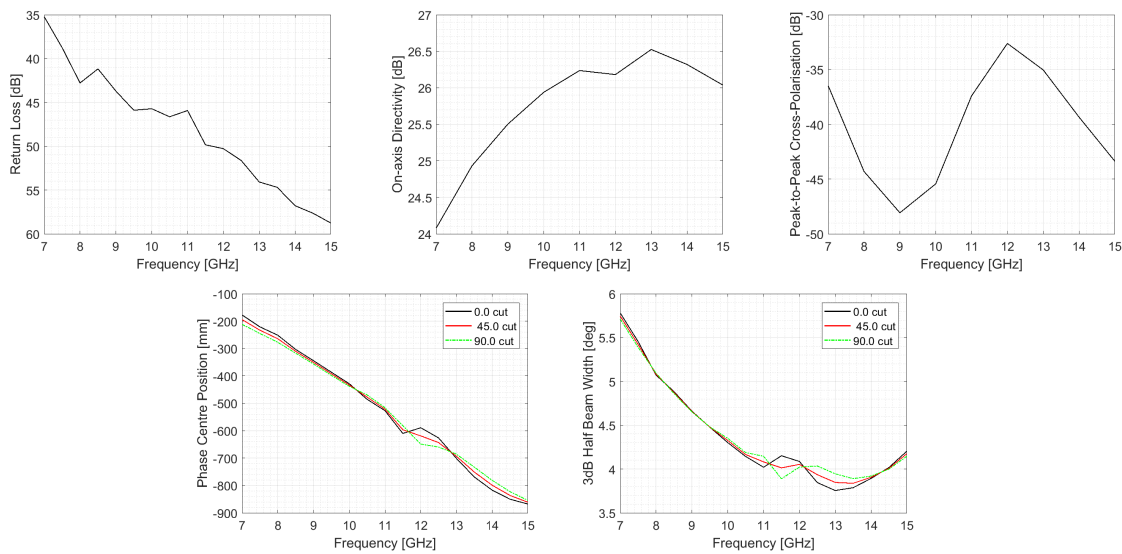


Figure 4.18: NextBASS Band 1 CHAMP Parameters: These show the return loss is below 35dB across the band, the on-axis directivity does not vary more than 1.5 dB and the cross-polarisation is < 35 dB for the majority of the band. This means the hyperbolic horn profile will meet the requirements for the NextBASS instrument. The plot also shows that there is a variation of 600 mm in phase centre location across the band. This is a large variation, covering nearly the entire length of the horn. A good choice of phase centre location is the phase centre of the design frequency. The 0.0° , 45.0° , and 90.0° cuts are very similar meaning the beam is nearly purely circular in the ϕ direction.

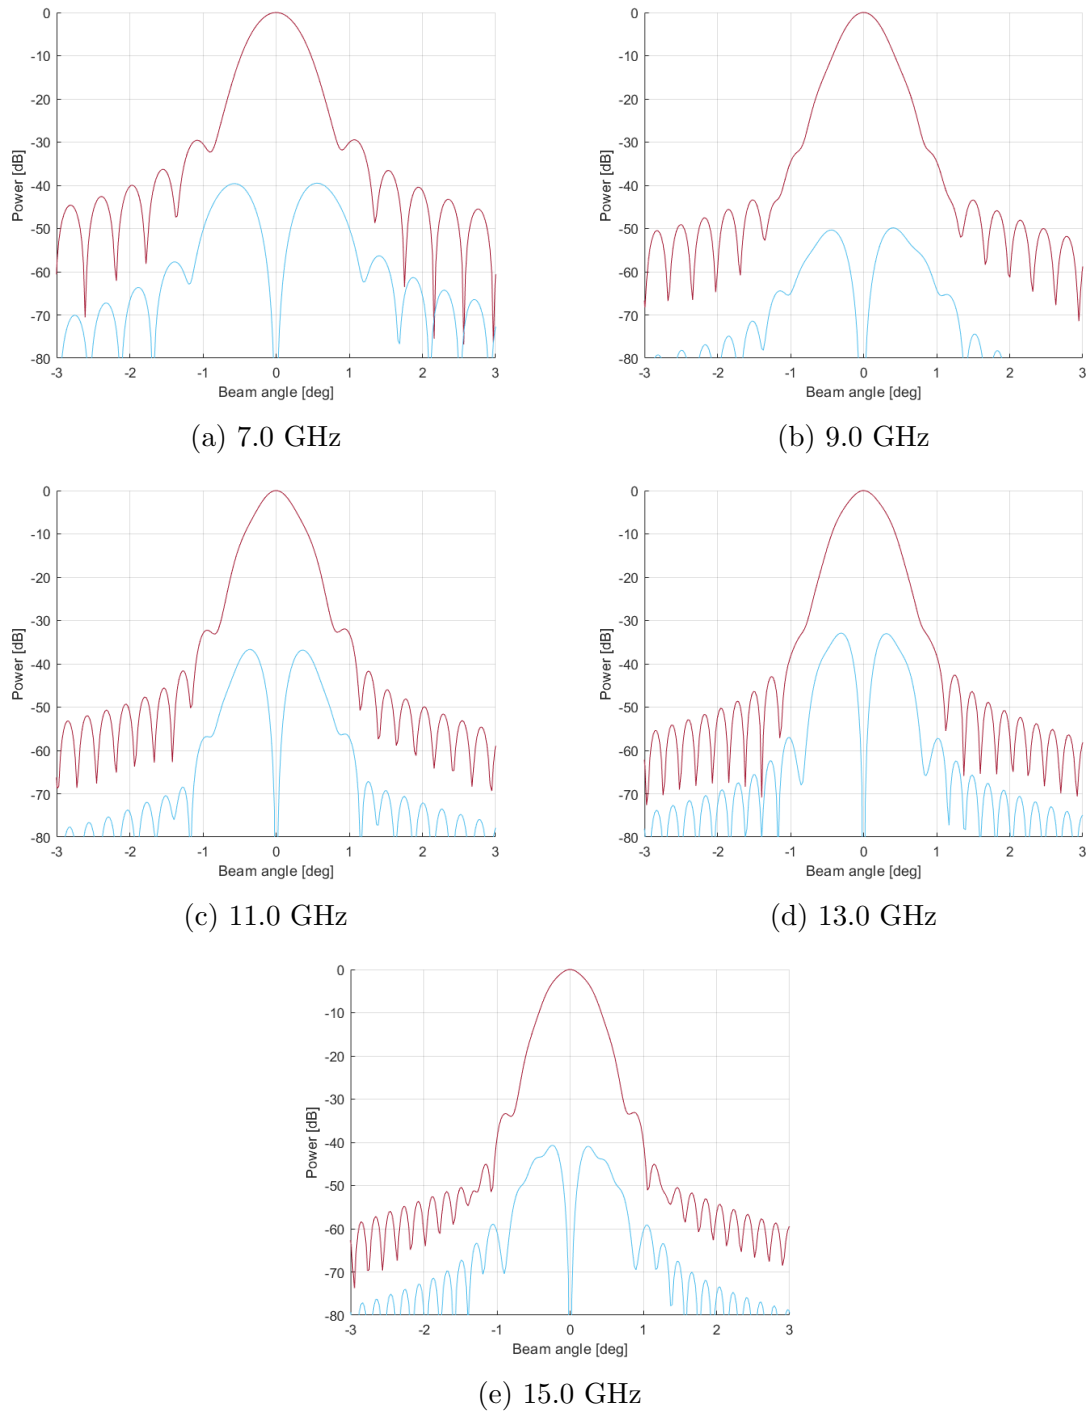


Figure 4.19: The main beam pattern of the 7 – 15 GHz hyperbolic NextBASS¹ feed horn the with the 6m cross-Dragone Optics. Each frequency has low first sidelobes, excellent beam symmetry and low cross-polarisation. The 7.0 GHz beam pattern has less well-defined first-nulls due to the phase centre location favouring the higher frequencies. The phase centre location has been chosen to match the design frequency. The design of the optics means that the sub-reflector is in the far-field of the beam, which means the varying phase centre of the horn causes little de-focusing, as the overall optical system behaves similar to ray optics.

4.5.4 Band-2 corrugated horn design

I generated the high frequency feed with the process as outlined above, only scaling the centre frequency. I used the 25 dB edge taper and found the best horn design to almost match a scaling factor of 0.5, as one would expect when we double the frequency and band-width. The resulting horn dimensions can be found in Table 4.3.

The beam pattern of this horn is similar in shape and cross-polarisation levels to the band-1 horn. The beam pattern at 15.0 GHz shows a higher gain at larger angles, in practice this will result in a larger ground-spillover at this frequency. The lower frequencies of the horn are always susceptible to higher spillover, provided that the performance of the horn and the optics still achieves a suitable level of power in the main beam, this increase in spillover is not problematic. It should also be noted that there is an overlap at 15.0 GHz from band-1, which achieves much lower sidelobes at this frequency.

In Figure 4.20, I have plotted the beam pattern for the NextBASS² hyperbolic feed horn for 3GHz intervals across the band. Furthermore, Figure 4.21, shows the return loss, the directivity, peak-to-peak cross-polarisation, 3dB beam widths and the 3dB phase centres for this horn. The performance is of the same standard as the band-1 horn.




Figure 4.20: the NextBASS² 15 – 30 GHz hyperbolic feed horn beam pattern generated in CHAMP. In this plot we can see that the beam widths are similar across the band and that the feed has low cross-polarisation. At 16.0 GHz, the beam has slightly larger sidelobes before becoming more stable at the higher frequencies.

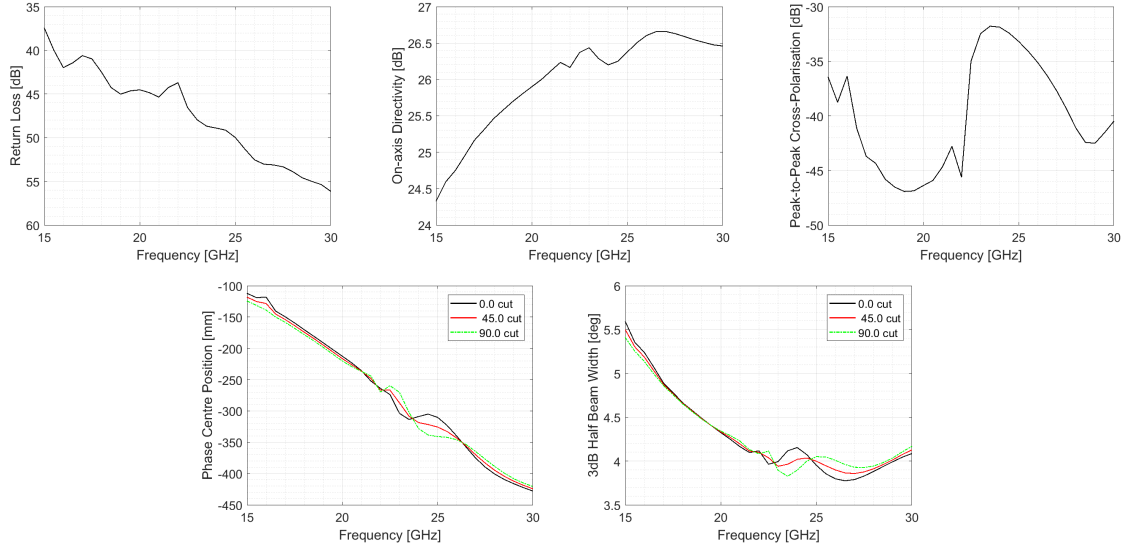


Figure 4.21: The CHAMP analysis of the 15 – 30 GHz, NextBASS² hyperbolic feed horn. There is sufficiently low cross-polarisation across the band along with good beam symmetry.

Table 4.5 gives the properties of the beam with the NextBASS cross-Dragone optics. The feed maintains high beam circularity, low cross-polarisation and low sidelobes across the band. In Figure 4.22, the main beam for each 3 GHz interval is shown with the co-polar 0.0° cut and the cross-polar 45.0° cut.

ν [GHz]	Gain [dB]	E_{CX} [dB]	1 st SL [dB]	Θ_{3dB} [deg]	A_{eff} [%]	sky _{eff} (1°) [%]	e^{3dB}
15.0	57.49	-36.83	-27.48	0.26	39.38	95.33	0.04
18.0	59.06	-38.04	-26.34	0.22	39.28	95.78	0.05
21.0	60.50	-40.26	-29.13	0.18	40.18	96.73	0.05
24.0	61.58	-41.26	-28.42	0.18	39.50	97.32	0.06
27.0	62.51	39.77	-27.77	0.14	38.68	97.26	0.07
30.0	63.32	-43.73	-28.59	0.14	37.72	97.92	0.08

Table 4.5: NextBASS band-2 feed horn results in conjunction with the cross-Dragone optics. This Table demonstrates that the corrugated horn with a hyperbolic profile is capable of achieving cross-polarisation of less than -36 dB with very low sidelobes. The beam also produces an efficiency of > 95% across the band. The results in this Table correspond to the beam patterns in Figure 4.22.

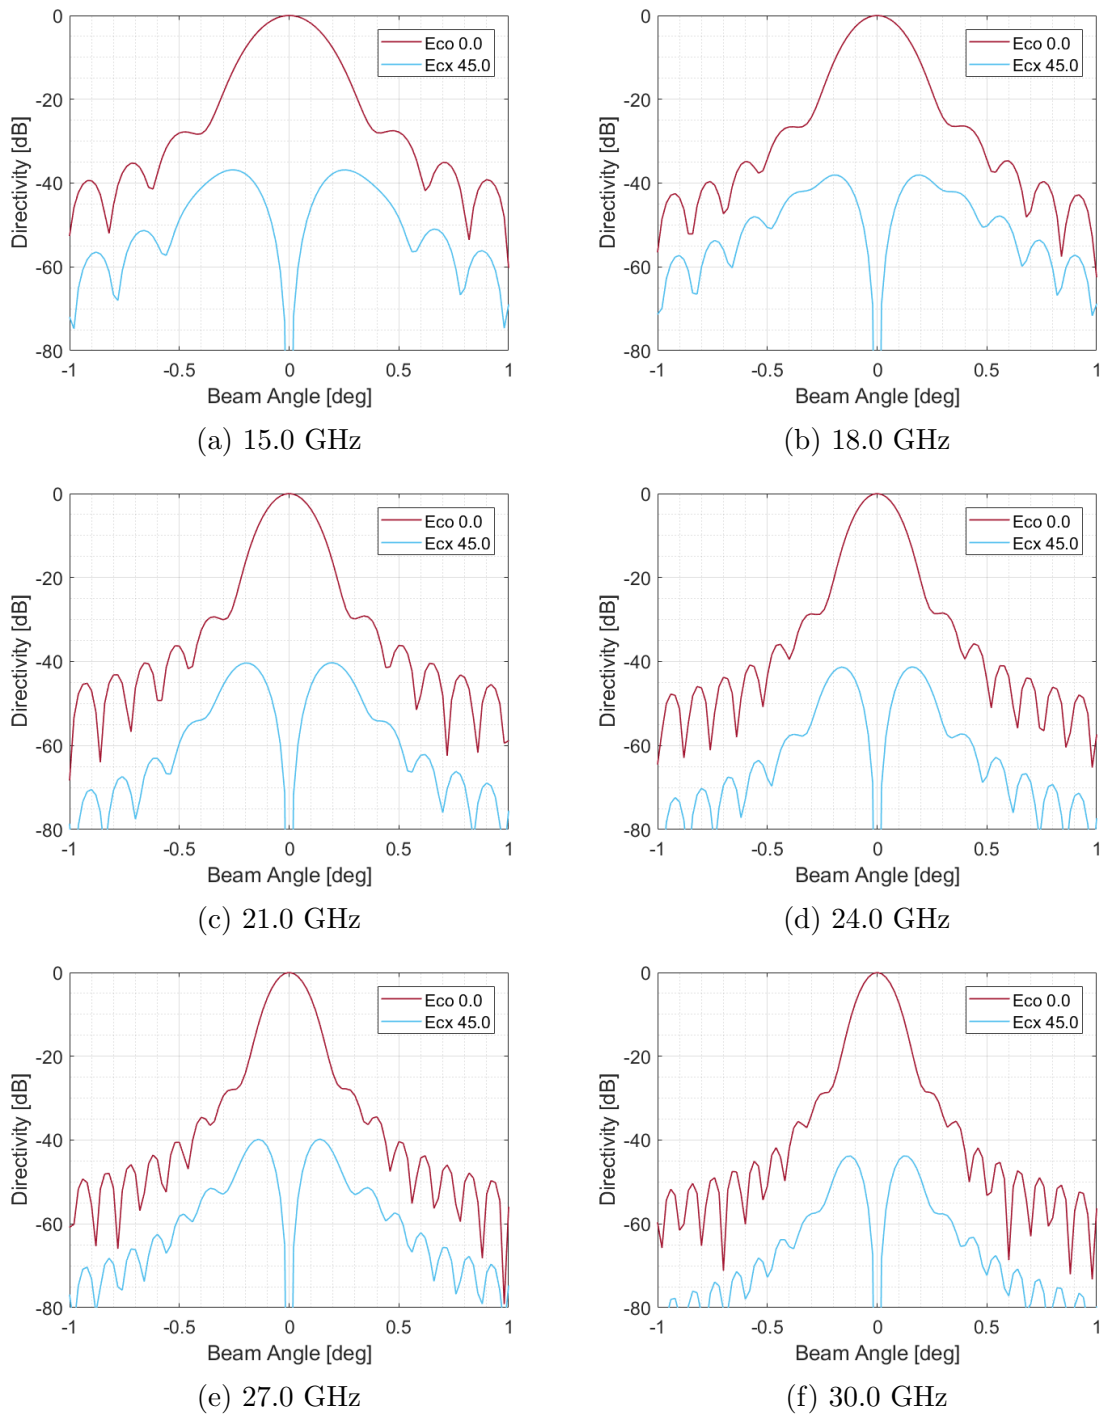


Figure 4.22: Performance of the band-2 hyperbolic feed horn with the NextBASS cross-Dragnone optics. There is < -40 dB cross-polarisation across the band.

4.6 NextBASS Edge Taper

When checking that the band-1 and band-2 feeds create the desired edge taper on the primary, I found that the edge taper varies from between -20 to -30 dB for band-1, and -25 to -35 for band-2. Due to the asymmetry of the off-axis optics, the left- and right-hand side illuminations differ. In Figure 4.23, the primary and secondary reflectors illuminations are shown for both band-1 and band-2.

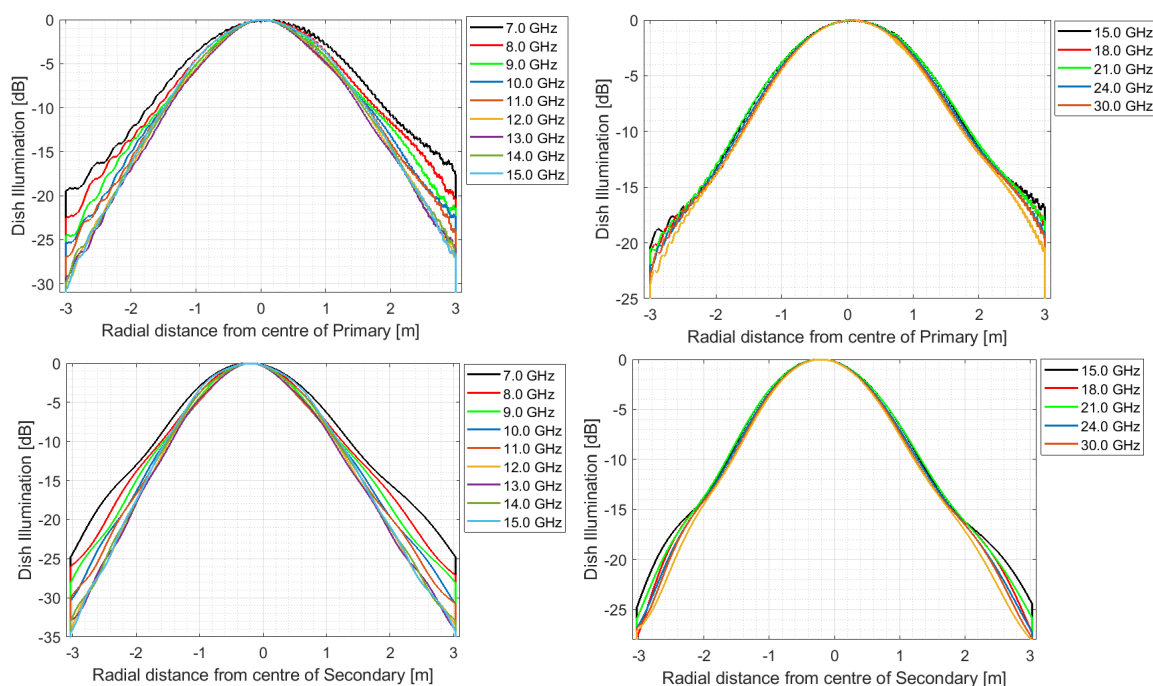


Figure 4.23: The band-1 (left) and band-2 (right), primary (top) and secondary (bottom) edge tapers. The designs match the -25dB edge taper, with some differences due to changes in beam width as a function of frequency. The ripples in the band show the diffraction edge effects, which only affect the lower frequencies. Both the primary and secondary mirror have similar edge tapers.

4.7 Feed Array

In this section, I discuss the options for circular packing the band-2 feed arrays. The feed-array requires at least 33 feed horns to meet the target sensitivity of NextBASS at the higher frequencies. As a result, the array must be designed so that these feeds are packed together as closely as possible to avoid moving into regions of the focal plane that are more distorted. I then discuss the possible manufacturing methods to reduce the costs and time-scales in designing such an array.

4.7.1 Circular Packing

In Chapter 6 I demonstrate that at least 33 feeds are required to reach an equivalent sensitivity to synchrotron of $\approx 1\mu\text{Karcmin}$ for band-2 of NextBASS. I will outline two possible array configurations that use 30+ feed horns. The first of which is based around using hexagonal modules each consisting of seven feeds, this design allows an easy removal of feed modules if it is required that the focal plane is to be smaller, or replaced with different feed horns. The disadvantages of this design is that it does not utilise the best possible packing arrangement. The second design concept uses a tight circular packing of 33 feeds with two removed in order to fit the larger, band-1 feed. The band-1 feed would have to be placed at the edge of the focal plane. Figure 4.24 shows the configuration which utilises circular packing.




Figure 4.24: Possible array configuration using the measurements of the feed horns designed in this chapter.

4.7.2 Feed Positioning

Figure 4.24 shows the placing of the higher frequency feeds in the centre of the focal plane, and the one 7-15 GHz feed on the edge of the focal plane. This is commonly done in CMB radiometer experiments. The focal plane on *Planck* placed the High Frequency Instrument (HFI) in the centre of the focal plane, and the Low Frequency Instrument (LFI) outside of this (?). They did this as the HFI instrument required a

much lower system temperature. Similarly for NextBASS, the higher frequencies can suffer from greater losses in sensitivity arising from an increased system temperature, both from the sky and the analogue electronics of the receiver. As a result, it becomes necessary to minimise this at earlier stages. By having the higher frequencies at the centre of the focal plane, the overall increase in spillover from the angular positions of the feed will be minimised. This is a logical decision as the introduction of the feed array is to increase sensitivity, so the positioning of the feeds on the focal plane must compliment this.

4.7.3 Manufacturing Methods

I will now comment briefly on methods of manufacturing for the NextBASS feed horns. Traditional manufacturing methods include moulding or etching a large piece of metal. However, making a horn out of one piece of metal is costly and not necessary at long wavelengths, as the joins are much less than a wavelength in thickness. The C-BASS feed horn designed by the group was produced by the etching of aluminium. The horns are split into sections, each machined and then bolted together. This method is costly and time consuming and for a system that requires many corrugated feed horns, new methods may be required.

In the past decade, 3D printing has gone from being state of the art, to becoming a cheap way to manufacture parts in high precision out of a range of materials. There are two options for the 3D printing of feed horns, these are 3D-printing a dielectric and providing a metal coating to cover it, or additive manufacturing, which 3D-prints the metal immediately.

[Samuel Chieh et al. \(2014\)](#) developed a *Ku*-Band corrugated conical horn using stereolithography. The antenna was printed using a thermoplastic which was then coated with conductive aerosol paint. They demonstrate that the feed operates as expected however with more losses than the simulated results. They suggest this may be due to the dielectric surface beneath. [Zhang et al. \(2018\)](#) state that 3D printing of feed horns with dielectric may be susceptible to electrostatic discharge causing losses. However, feed horns produced using this technique are extremely cheap and lightweight compared to their metal counter-parts and would be an excellent option for large focal plane arrays or satellite focal planes.

The advantages of additive manufacturing include faster prototyping and production of metal parts, a reduction in cost and weight reduction. It is likely that the X-BASS horn will be produced in a similar method to the C-BASS horn, but additive manufacturing remains a possibility for NextBASS and the Band-2 feed array.

4.8 Future Feed Development

In conclusion to this chapter, I have shown that high performing, low cross-polarisation feed horns with a 2:1 bandwidth can be designed for the two optical systems of X-BASS and NextBASS. By using optimisation routines I have developed a feed that is well matched to the C-BASS optics. Similarly, I have demonstrated that it is possible to design far-field horns for the cross-Dragone optics using a Gaussian optics, far-field ‘recipe’, and that the beam patterns of the resulting horns perform well with the NextBASS, cross-Dragone optics. The future work that would improve upon the feed designs in this thesis includes getting the horns machined and the real radiation patterns measured in an anechoic chamber. The feeds must then be tested with a suitable Orthomode-transducer (OMT) that matches the waveguide radius and bandwidth of each horn. This would make up the first part of the cryogenic stage of the receiver, this is discussed in detail in the following chapter. Further more, the NextBASS horns will require further optimisation to minimise their size and improve cross-polarisation. The reduction in cross-polarisation is particularly important at wide angles so that when using the feed horns off-axis in the large focal plane of the cross-Dragone optics, the performance of the beam is still up to a high standard.

Chapter 5

A Conceptual 2:1 Bandwidth Polarisation Receiver

“ The wise know their limitations, the foolish do not. ”

Benjamin Hoff, The Tao of Pooh

In this chapter, I describe the X-BASS and NextBASS receiver architectures. As NextBASS is still in its conceptual design phase, the Digital Back-end (DBE) technology processing speeds available today, are likely to be outdated by the time of NextBASS' implementation. None the less, it is important to make decisions at this stage that use parts with a high Technology Readiness Level (TRL). I base the receiver architecture of NextBASS on the specifications of a 3.5 bit flash converter Analogue-to-Digital Converter (ADC) with 20 GHz of bandwidth.

I will compare and contrast different receiver architectures and methods of gain stabilisation to minimise pink noise, also known as $1/f$ noise. Both X-BASS and NextBASS receivers will use a Continuous Wave (CW) stabilised system. This is based on the gain stabilisation testing done by Pollak in [Pollak \(2018\)](#), which is being further developed by [Pollak and Holler \(in prep.\)](#) in prep. The CW architecture has fewer RF and IF channels than a Continuous Comparison receiver, such as C-BASS, making it less susceptible to sensitivity degradation. The CW-stabilised system can also exploit the faster processing speeds available with modern FPGA's and perform gain stabilisation in real-time.

The X-BASS receiver is to be implemented on the C-BASS optics, on a time-scale shorter than that of NextBASS, therefore it is more convenient to use a DBE that has been implemented previously on experiments with similar frequencies and bandwidths. Therefore, the choice of FPGA will be the RFSoc, which has two inbuilt

ADCs with 2x2 GHz bandwidth. To use these ADCs with the X-BASS frequencies requires two separate LO's, four IQ mixers and a 90 degree hybrid to generate a double sideband IF channel. This technology has a high TRL since the similar style receivers have been implemented on GHY-3, Goonhilly and the SKA Band 5 (Pollak (2018), XILINX (2018)).

The large bandwidth of the NextBASS ADC means the IF channels can be designed so as to not reduce the bandwidth. By choosing the local oscillator (LO) frequencies of 16.5 GHz and 31.5 GHz for band 1 (7 – 15 GHz) and band 2 (15 – 30 GHz) respectively, the same mixer and anti-aliasing filter can be used for both bands.

In this Chapter, I will also discuss the process of designing the anti-aliasing and bandpass filters required for NextBASS. The proto-type anti-aliasing filter (AAF) is a stepped-impedance low-pass filter (LPF) with a sharp roll-off at 9.5 GHz. The proto-type filter has been realised on RT/duroid R04530, (ROGERS corporation (2018)), and measured to give a 20 dB return loss across the band, and only a small shift in cut-off frequency due to Ohmic losses. Using Chebyshev filters, I have also designed 7 – 15 GHz and 15 – 30 GHz bandpass filters. The prototypes of these still require further development as the track widths are very narrow and difficult to manufacture.

5.1 Radio Receiver

A radio receiver takes an incoming radio frequency (RF) signal and converts it into bands that can be processed by the digital back-end (DBE). To do this, a receiver must perform some basic tasks. These include filtering out all unwanted frequencies, amplifying the signal with a low-noise amplifier (LNA) and sometimes frequency shifting with a mixer and an local-oscillator (LO).

When the frequency has to be shifted, the receiver is super-heterodyne. The RF signal is combined with the LO at a specific frequency using a mixer, as a result an intermediate frequency (IF) band is produced. The IF band limits are the sum and differences of the LO and RF frequencies. The main advantage of this receiver is that performing filtering and amplification over an IF band is less power intensive. Another advantage is that the IF filter can be fixed and the LO can be tuned to achieve different RF channels (Burke and Graham-Smith (2010)).

5.1.1 Receiver Front End

The first component of the receiver is the feed horn. This converts the propagating EM wave into the desired modes which can be transformed to an electrical signal by a

transducer. The feed horn is attached to a waveguide, which is cooled in a cryostat. As a result the feed, waveguide and transducer will be cooled to different stages, starting at ambient temperature ($\approx 300\text{K}$), to an intermediate stage ($\approx 70\text{K}$) and a cool stage of usually 10 – 15 K. A radiometer-style instrument does not require stages cooler than this, because they measure the incoming radiative flux via analogue electronics. The cryostat is required to keep the first components of the receiver at a low noise level. As a result, radiometers have the advantage over a bolometer style instrument as they require considerable cooling since they measure power by the incident EM waves heating a material with a temperature-dependent electrical resistance. The low noise amplifiers (LNA) and connectors are placed in the coolest cryogenic stage. The intermediate stage also houses the first filters and then the following RF chain is at ambient temperatures.

An orthomode transducer (OMT) is used to extract two orthogonally polarised modes from a waveguide, and transfer them into separate coaxial cables. The OMT is positioned at the feed horn throat, which has already transformed the received EM signal to the desired propagating modes. For a circular waveguide, the orthogonal polarisations are produced from the TE_{11} mode. As a result, the system generates linear polarisation in the E_x and E_y directions. The OMT must perform this function well across the full bandwidth. Typical OMT designs use either two (double-ridge) or four (quad-ridged) ridges that are tapered to achieve impedance matching between unloaded waveguide and the 50Ω coaxial impedance. The ridges can be profiled to optimise performance over a large bandwidth. [Pollak \(2018\)](#) designed a quad-ridged OMT that is suitable for use over a 2:1 bandwidth.

5.1.2 Analogue Components

The analogue components of the receiver begin at the coaxial cables at the OMT and include all analogue electronics until the ADC, where the signal is digitized.

The LNA is a high gain amplifier that boosts the signal gain immediately after the OMT. The first component in the receiver is the largest contributor to system noise. Since the LNA is high gain, any lossy components coming after it will be negligible. To achieve a good signal-to-noise-ratio (SNR), thermal noise from the LNA must be minimised by cooling the LNA in a cryostat. Often, additional amplification is needed in the warm, RF and IF stages. These amplifiers can be stacked with bandpass filters to improve out-of-band rejection.

Filtering can be done before and after the LNA. Having a filter before the LNA is known as band pre-selection. This blocks large out-of-band signals, reducing power

and helps prevent noise arising from the image signal. This is not ideal for astronomical applications because it limits the power of incoming signals and any loss due to the first bandpass filter will result in a lower SNR. It is best to amplify all incoming signals and filter after the LNA (Wilson (2012)). One of the major requirements of RF filtering is to attenuate any unwanted Radio Frequency Interference (RFI). Image frequencies from the LO must also be filtered out. A sharper cut-off can be achieved by stacking amplifier and filters in series. I describe the processes involved in microwave filter design in the Section 5.6.

The LO can be fixed or tuned to specify the band limits of the IF band. When choosing the LO frequency, care must be taken to avoid contributions from image frequencies as these can interfere with the original signal. The image or Nyquist frequency arises from the Nyquist-Shannon sampling theorem such that a perfect reconstruction of a sampled signal gives a bandwidth, $B < f_s/2$ (Burke and Graham-Smith (2010)).

Once the LO frequency has been chosen, a mixer is used to combine the LO and RF frequencies. The resulting IF band is then amplified and filtered using an anti-aliasing filter. For broadband applications, the mixer must have a good response in both the RF and IF bands. The limits of an IF band are given by $f_{IF} = |f_{LO} \pm f_{RF}|$.

The generation of IF frequencies creates two distinct IF bands due to the sum and difference of RF and LO frequencies. These are known as ‘sidebands’. When the LO frequency is within the RF band, the system uses a double sideband (DSB). Both sidebands are then required in the reconstruction of the original RF signal by using a complex Fourier transform (FT). If the LO is outside the RF band, two sidebands are still produced, but only one is required to reconstruct the original signal. This is done using the real part of the FT in the DBE. Figure 5.1 shows the single and double sideband configurations in frequency space (Burke and Graham-Smith (2010)).

To obtain the complex FT, the signals must be split and mixed using an IQ mixer. This means splitting the IF channels, resulting in lower power levels at the ADC input. Fewer channels means the SSB can achieve a higher SNR. The DSB also results in a reduction of bandwidth where as the SSB is good for recovering large bandwidths. As a result, it is more computationally expensive to process the SSB.




Figure 5.1: Double and single sideband generation. The double sideband requires the complex part of the Fourier transform to extract the RF band.

5.1.3 Digital Backend

The last component in the analogue receiver chain, and the first part of the DBE is the analogue-to-digital converter (ADC). This takes the IF band and digitises the signal. This process is known as quantisation, where the continuous analogue signal is approximated a number of discrete levels. The number of digital levels or counts, N , is determined by the number of bits, m available, $N = \pm 2^m$ (Kester (2008b)). Figure 5.2 demonstrates how a 3-bit ADC can quantise a sinusoidal signal into $2^3 = 8$ discrete levels.




Figure 5.2: Example of quantized sine wave for a 3-bit, 1Mbps ADC with a maximum ADC peak-to-peak voltage of $V_{peak} = 1.0V$. The value of the least significant bit (LSB) is indicated.

The ADC speed determines how many samples it can process per second. Modern ADCs can sample in tens of giga-samples per second (GSPS). The bias voltage, V_{peak} , is the maximum electrical bias generation possible for the ADC. An important parameter to define is the least-significant bit (LSB) of an ADC. This is the smallest level of voltage that an ADC can convert. (Kester (2008b))

$$\text{LSB} = \frac{V_{peak}}{2^m} \quad (5.1)$$

The quantisation noise power tells us the amount of power in any level (Kester (2008a)). This depends on the terminating impedance, Z_0 , of the system. This is typically $Z_0 = 50\Omega$. This is usually given in units of dBm and can be converted using $Q_P(\text{dBm}) = 10 \log_{10}(1000 \times Q_P(W))$.

$$Q_P = \frac{\text{LSB}^2}{12Z_0} \quad (5.2)$$

If we consider a signal that fills half the dynamic-range of the ADC, the root-mean square voltage, (V_{rms}) is given by equation 5.3. This can be used to obtain the *rms* power of the ADC. By knowing P_{rms} , it is possible to design the receiver chain to meet the input power required by the ADC. The dynamic range is the full power an ADC can process before saturating, it is given by $DR(\text{dBm}) = Q_P - P_{rms}$.

$$V_{rms} = \frac{V_{peak}}{2\sqrt{2}} \quad (5.3)$$

The P_{rms} of the ADC can be calculated by:

$$P_{rms} = \frac{V_{rms}^2}{Z_0}. \quad (5.4)$$

A Field-Programmable Gate Array (FPGA) is an integrated circuit which can be customised after manufacture. It consists of programmable logic blocks and reconfigurable interconnects. These can be configured to perform complicated combinational logical functions or simple logic gates such as *AND*, *OR*, *XOR*, *NOT*, *NAND*, *NOR* and *XNOR*. They also contain memory blocks as either flip-flops or more complicated memory (XILINX (2019)).

In radio astronomy, once the EM signal of the sky has propagated through the receiver and converted to a digital signal, it must then be transformed from a time-dependent signal into a frequency dependent signal. The receiver has been designed to have a finite bandwidth and has down converted the incoming frequency. A Fourier transform must be performed to obtain the frequency information of the incoming EM wave.

FPGAs are extremely useful for this application as they can perform Fourier transforms on the incoming signal in real-time, along with breaking the signal into smaller frequency channels to create spectral data. Their fast processing speeds make them ideal for processing large datasets.

Since multi-GSPS FPGAs are becoming more widely accessible, processing large datasets is a lot easier. As a result, spectral back-ends can be implemented. In order to create a power spectrum from the digitised signal, the incoming signal must be filtered into narrowband channels, channelised and then Fourier transformed (Price (2016)). The channelisation creates smaller, narrower bandwidths, which when converted into frequency space result in a higher resolution spectral image. This is a practical benefit for radio astronomy as interesting spectral features can be reconstructed from the data. The combination of this channelisation followed by the transformation into the frequency domain is an essential part of radio astronomy digital back-end design, and a number of pre-written algorithms exist to do this.

A commonly used algorithm is the Poly-phase Filter Bank (PFB). A PFB is a computationally efficient implementation of a filter-bank with very low spectral leakage. A filter-bank is an array of bandpass filters that split the input signal into narrowband components. The input sequence is then decomposed into a set of sub-sequences, each of which is known as a poly-phase decomposition. This separates the channels further into multiple narrow sub-bands using a fast Fourier transform (FFT) based filter-bank. A windowing function is applied to the channelised data to reduce the side-lobes and broaden the channel response. Having a spectral back-end means that narrow band phenomena can be observed even with a broadband receiver architecture, and calibration techniques using a continuous wave signal can be implemented. Another advantage is that any narrow band signal interference, for example satellite RFI, can be isolated in the post-processing stage (Price (2016), Pollak (2018)).

Polyphase filter banks have been extensively developed and used in radio astronomy, particularly by the CASPER open source technology who develop ROACH boards and the algorithms required for analysis (Collaboration for Astronomy Signal Processing and Electronics Research (2009)). The CASPER style technology has been used on the Greenbank telescope, the Very Large Array (VLA) and the Atacama Large Millimeter/submillimeter Array (ALMA)(Morris and Luce (2010)).

5.2 Sources of Noise

Astronomical signals have random noise characteristics. The time-domain amplitude is usually described by a Gaussian random noise distribution. Due to their random nature, the astronomical signals are difficult to isolate and hence the accuracy of a

receiver depends on the statistical nature of the voltage amplitude generated in the receiver (Burke and Graham-Smith (2010)).

There are two types of noise in the analogue chain of a receiver. Each can be characterised by their frequency spectrum. The motion of thermal electrons in the conducting elements produces uncorrelated noise with a flat frequency response. Correlated noise sources show a $1/f$ relationship with a greater response at lower frequencies.

5.2.1 Thermal Noise

Thermal noise, also known as Johnson-Nyquist noise, is generated by the resistance of any components in the receiver chain. In the time domain, it can be thought of as the random motion of electrons in the circuit. The noise introduced by a component is described by (Burke and Graham-Smith (2010)):

$$P = k_B T B, \quad (5.5)$$

where P is the maximum noise power, k_B is the Boltzmann constant, B is the bandwidth and T is the thermal temperature of the component.

The receiver's thermal noise will be well characterised by the random noise of the first amplifier. Since this gain variation happens on timescales less than the coherence time of radiation, the receiver gain can be described as a fixed random variable. The probability of the voltage amplitude, $v(t)$, is given by the probability density function:

$$P(v(t)|\mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-(v(t) - \mu)^2}{2\sigma^2}\right), \quad (5.6)$$

where σ is the standard deviation of voltage amplitude and μ is the mean signal.

Since the statistical properties of white noise can be described by a Gaussian distribution around zero, in equation 5.6, the standard deviation describes the *rms* voltage amplitude, $v_{rms}(t)$ and the mean can be considered as $\mu = 0V$ for a noise voltage (Wilson et al. (2008)).

5.2.2 Pink Noise

Pink noise, commonly referred to as $1/f$ noise has a $1/f$ distribution in the noise power spectrum. The point at which the $1/f$ signal falls below the thermal noise of the receiver is known as the knee-frequency, f_k . The $1/f$ noise represents the correlated gain fluctuations of a system. As a signal propagates through the receiver chain, each

temperature fluctuation becomes correlated and hence contributes to the $1/f$ noise. By tracking gain fluctuations on time-scales faster than the correlations, $1/f$ noise can be reduced. Correlated noise is less of a concern for polarisation measurements as the generation of Q and U Stokes parameters removes most correlated noise.




Figure 5.3: The power spectra of data from the C-BASS North receiver showing the $1/f$ noise. The black trace shows the sky data channel with a knee frequency of $f_{knee} \approx 1\text{Hz}$. The red traces shows the gain calibrated channel after the load signal is subtracted from the sky signal. The Q and U channels (green and blue respectively), show a much lower knee frequency below 10 mHz. This image has been reproduced from Figure 14 of [King et al. \(2014\)](#).

5.2.3 Quantisation Noise

The accuracy of digitisation depends on the number of available bits. The error that is introduced during analogue-to-digital conversion can be described as a rounding error between the analogue signal and the digital waveform. In an ideal ADC, the quantisation error is distributed between $-1/2$ LSB and $+1/2$ LSB and the signal has a uniform distribution covering all quantisation levels. Quantisation noise is additive to the white noise of the receiver and can be treated as a random noise signal with an *rms* power. The signal-to-noise ratio of an ADC is given by ([Kester \(2008b\)](#)):

$$SNR \approx 1.761 + 6.02Q_P(\text{dB}), \quad (5.7)$$

where the value of 1.761 accounts for a full-amplitude sine-wave, whose distribution is non-uniform, and the quantisation noise SNR_Q is given by $SNR_Q = 20 \log_{10}(2Q_P) \approx 6.02Q_P$.

5.3 Methods of Gain Stabilisation

In a simple radiometer, there is no way of distinguishing receiver noise from the astronomical signal. This would rely on the source being dominant, however we know this is not the case. Receiver calibration performs two tasks, to characterise the white noise level of the LNA and to reduce the correlated receiver noise.

5.3.1 Total Power Radiometer

In astronomy, we need the radio receiver to measure the total power of the incoming signal. We refer to power measuring receivers as radiometers. The most basic design can be seen in Figure 5.4. This is known as a total power radiometer. This system would be able to collect, amplify and shift an incoming signal, however it has no means to characterise noise (Burke and Graham-Smith (2010)).




Figure 5.4: Total Power Radiometer

The receiver gain and noise fluctuations are independent random processes and their sum yields the total fluctuations of a total-power radiometer. This described the total-power radiometer equation, given by equation 5.8 (National Radio Astronomy Observatory (NRAO) (2017)). The noise term is given by the noise variance which depends on the RF bandwidth, $\Delta\nu$, and the integration time, τ . The gain fluctuation, ΔG , produces a signal that cannot be distinguished from the noise temperature of an astronomical source. As a result, the variance is given by the ratio of the gain fluctuation, ΔG , to the overall gain of the receiver, G , such that (Wilson et al. (2008)):

$$\sigma_T = T_{sys} \left[\frac{1}{\Delta\nu\tau} + \left(\frac{\Delta G}{G} \right)^2 \right]^{1/2} \quad (5.8)$$

5.3.2 Dicke Switch Radiometer

The earliest form of gain monitoring uses a Dicke Switch. This rapidly switches between a calibration load and the sky signal to characterise the gain as a function of time. By using a lock-in-amplifier, the two phase switch positions return either a positive or a negative sign (Wilson (2012)). When both the load and signal are equal, the lock-in-amplifier returns zero and hence a measurement of the gain can be made. The resulting temperature change from the gain fluctuations is given by:

$$\Delta T = \frac{\Delta G}{G}(T_S - T_{ref}), \quad (5.9)$$

where G is the total gain, ΔG is the gain fluctuation, T_S is the signal noise and T_{ref} is the blackbody temperature of the reference load.




Figure 5.5: Dicke Switch Radiometer

Figure 5.5 shows a Dicke Switch Radiometer that uses a lock-in amplifier and power detector to control the switch. This is controlled by a square wave with a duty cycle of 50%. Compared to a total-power radiometer, there is a degradation in performance by a factor of 2. This comes from a reduction in sensitivity of $\sqrt{2}$ since the antenna is only being measured for half the time, and a further $\sqrt{2}$ from the difference between two the uncorrelated random noise signals. The limit of obtainable sensitivity for a Dicke receiver is given by Equation 5.10. Usually this factor of two degradation in performance is a fair trade off for greater gain stability.

$$\sigma_T = 2 \times T_{sys} \left[\frac{1}{\Delta \nu \tau} + \left(\frac{\Delta G}{G} \right)^2 \right]^{1/2} \quad (5.10)$$

5.3.3 Continuous Comparison Radiometer

A continuous comparison, or differential radiometer uses a hybrid circuit to measure the difference between antenna noise power and a reference load. The C-BASS receiver is of this design. It uses the coherent nature of radiation and phase differences to achieve gain stabilisation. The input source and reference signals are connected to

two amplifiers through an 180° hybrid, which introduces a phase shift between them. The amplified signals are then multiplied together (Burke and Graham-Smith (2010)).

When looking at more extended sources such as the galactic background, an absolute measurement of antenna temperature must be carried out. This is done by replacing the antenna by a dummy load at a known temperature. The load and gain signals are added in quadrature and the resulting sums and differences can be used to subtract the gain and measure the astronomical signal. This measurement is difficult to perform as the hybrids may be asymmetric and there may be losses in the receiver due to impedance mismatching.




Figure 5.6: Continuous Comparison

The continuous comparison architecture has an advantage over the Dicke method as the receiver continuously measures the antenna signal and the stabilisation does not depend on the levelling of reference and sky signals. However, this stabilisation architecture requires twice the number of components as it uses two parallel signal chains to process both output signals of the hybrid. Its sensitivity hence decreases due to the additional noise of the reference signal by a factor of $\sqrt{2}$ compared to the total power receiver. There is also an additional increase in system temperature, T_{sys} , due to positioning the hybrid before the first amplifier (King et al. (2014)).

5.3.4 Noise Injection Radiometer

Noise injection radiometers are simpler and do not require the use of 180° hybrid circuits. By using a series of single pulses at fast sampling rates, they perform continuous gain monitoring of the receiver. They can also be used with slower pulses to perform calibration measurements of the LNA. They operate by introducing a reference signal before the first amplifier, this is also known as a white-noise stabilised receiver. This uses a noise diode of reference temperature T_{ref} which is computer controlled to generate the pulses.




Figure 5.7: White noise Stabilised

The fast firing white-noise stabilisation method allows the user to determine the current gain of the receiver. This is done by comparing the output power, with and without the noise diode. Since this method of stabilisation requires the reference signal to be turned on and off, there is still a decrease in overall power. If the diode has to be on 50% of the time, then the overall decrease in sensitivity would be the same as a Dicke Switched system. However, often the dynamic range of the receiver is usually considerably less than the noise diode signal, meaning much shorter duty signals can be used. The decrease in sensitivity will now lie between 1 and $\sqrt{2}$. The Goonhilly GHY-3 receiver utilises a white-noise stabilised wave architecture due to the multiple IF channels (Pollak (2018)).

5.3.5 Continuous Wave

A CW radiometer uses the same architecture as the white-noise stabilised system, however instead of a noise diode, a narrow-band signal is used. If we make the assumption that gain variations are not frequency dependent, measuring the gain at a specific frequency chosen randomly, allows for the correction of fluctuations over the full observation bandwidth. This is advantageous as continuous gain monitoring can take place without losing sensitivity.




Figure 5.8: Continuous Wave Stabilised

Another advantage of CW stabilisation is that the reference signal has a narrow bandwidth so the band is not contaminated by a broadband signal, like in the

white-stabilised noise architecture. As a result, CW stabilisation is a good choice for broadband receivers. In order for this method to work, a spectral back-end will have to be implemented in the FPGA using a polyphase filter bank (PFB), for example. The recovered gain of the system for a continuous wave signal can be described by :

$$G_{system} = \sqrt{\frac{P_{cw}}{P_{cw,R}}}. \quad (5.11)$$

where, P_{cw} is the detected power of the reference signal and $P_{cw,R}$ is the power of the reference signal at the input. The signal-to-noise ratio, of a CW system is the ratio of the detected power, P_{cw} to the power produced by the receiver system temperature, P_{sys} plus the power of the received antenna, P_A such that:

$$SNR = \frac{P_{cw}}{P_{sys} + P_A}. \quad (5.12)$$

This determines the level of CW signal required to calibrate the receiver. The strength of the CW signal at the output must be limited to less than the dynamic range of the ADC, otherwise the signal would be saturated and gain tracking would not be possible. Ideally, the CW signal would fill at least half of the dynamic range (see equation 5.3). The gain stabilisation techniques for the GHY-3 broadband receiver as developed by Pollak (2018), are applicable to the design for NextBASS. Further investigation into the application of CW stabilisation at higher frequencies along with checking gain variations as a function of frequency is part of the continued work required for NextBASS.

5.4 Polarisation

A polarimeter requires an expansion on the radiometer architectures shown previously. It perform an additional task to just measuring power, it must introduce a phase difference so that the I , Q and U Stokes parameters can be extracted. The easiest way to do this is by generating circular polarisation and then adding the circularly polarised channels in quadrature.

5.4.1 Stokes Parameters

In order to measure the polarisation, the receiver has to extract the Stokes parameters. The Stokes parameters fully describe the polarisation state of a propagating electromagnetic wave. They are given by the complex amplitudes of the propagating

electric field is given in either a Cartesian or circular basis (Burke and Graham-Smith (2010), et al (2014)).

$$I = |E_x|^2 + |E_y|^2 = |E_L|^2 + |E_R|^2 \quad (5.13)$$

$$Q = |E_x|^2 - |E_y|^2 = 2\Re(E_L E_R^*) \quad (5.14)$$

$$U = 2\Re(E_x E_y^*) = -\Im(E_L E_R^*) \quad (5.15)$$

$$V = -2\Im(E_x E_y^*) = |E_L|^2 - |E_R|^2 \quad (5.16)$$

The I Stokes parameter represents the total power of the combined polarisation states. The Q parameter, depending on the directionality represents the linear polarisation states. The U parameter is the linear polarisation, 45° to the Q channel. Finally the V parameter represents either purely left-hand or right-hand circular polarisation, depending on the sign. Most astronomical sources do not generate the Stokes parameter V , as they do not emit circularly polarised light.

Obtaining a value of Q from the purely measured E_x and E_y can result in large errors due to random noise signals. It is better to generate left-hand, E_L , and right-hand, E_R , circular polarisation, which then removes any uncorrelated noise signals. The uncorrelated noise is removed because when generating the circular polarisation, both the E_x and E_y signals are added in quadrature, so any correlations between the two are also removed. Generating circular polarisation in this way also makes the generation of Stokes Parameters simpler in terms of computation.

The circular polarisation signals are mathematically described by:

$$\begin{aligned} E_L &= \frac{E_x + iE_y}{\sqrt{2}} \\ E_R &= \frac{E_x - iE_y}{\sqrt{2}} \end{aligned} \quad (5.17)$$

In both the C-BASS South and North receivers, the analogue receiver architectures use 90° hybrids to introduce phase shifts between E_x and E_y , generating left-handed and right-handed circular polarisations, E_L and E_R . The 90° hybrids typically are designed to operate over a certain bandwidth and should perform the task with minimal polarisation leakage over their operational range. The gain tracking in C-BASS is carried out by a reference which is coupled using an 180° hybrid before the sky signal and after the linear-to-circular conversion. This allows the diode to follow the

same path through the receiver chain. Two load signals are added in quadrature to the circular polarisations, which are then split into equal parts into a radiometer and a polarimeter. The radiometer extracts I by generating $|E_R|^2$ and $|E_L|^2$ and adding them (et al (2014)). The matrix multiplication of these channels can be given by:

$$\begin{bmatrix} \langle E_R E_R^* \rangle & \langle E_R E_L^* \rangle \\ \langle E_L E_R^* \rangle & \langle E_L E_L^* \rangle \end{bmatrix} = \begin{bmatrix} I + V & Q + iU \\ Q - iU & I - V \end{bmatrix} \quad (5.18)$$

This means that at the output of the four channels, it is just a case of finding the sum and difference to obtain the Stokes parameters. Since it is unlikely to receive signals in Stokes parameter V , the $I + V$ and $I - V$ channels are added together to produce a $2I$ output. The $Q + iU$ and $Q - iU$ channels are added and subtracted to obtain $2Q$ and $2iU$ respectively.

5.5 Receiver Architectures

Both the X-BASS and NextBASS receivers will be super-heterodyne, and measure linear polarisation, implementing a continuous wave (CW) signal to monitor the gain drift of the receiver. As a result, the two polarisation signals follow the same signal paths making gain tracking more reliable. This design is different to the continuous-comparison radiometer implemented in C-BASS North and South, which requires the use of hybrid couplers to perform gain tracking. A continuous comparison radiometer is not ideal for X-BASS or NextBASS because of the degradation in sensitivity due to the multiple RF channels, and the difficulty in implementing analogue hybrids over a broad bandwidth. The best option is to use the CW calibrated system and implement a spectral back-end in the FPGA, in order to identify the narrowband signal. This means X-BASS and NextBASS will not be susceptible to broad-band noise, and the receivers can be kept as simple as possible in order to maintain a low system temperature, hence an improved sensitivity. Furthermore, fewer analogue components means less polarisation leakage and therefore, more reliable polarisation measurements.

5.5.1 Friis Equation

The process of deciding on the number of amplifiers and filters required along with the ordering of components in the chain requires knowledge of the noise for each component in order to estimate how the gain will propagate through the receiver chain. This process is described by the Friis formulas of a receiver, which can be used

to find the resulting power output and system temperature at each component in the chain (Balanis (2005)).

The Friis formulas for noise provide a relationship between the temperature and gain of each successive component. The Friis formula states that the overall noise figure of a radio receiver is determined by the noise figure of the LNA, and that the following stages only lessen the SNR by a small amount. With more modern systems, the LNA noise is better controlled. The Friis formula states that the overall receiver noise figure can be expressed in terms of the LNA noise figure, plus the noise figure for all subsequent components in the receiver. This is described by (Burke and Graham-Smith (2010)):

$$F_r = F_{\text{LNA}} + \frac{F_{\text{chain}} - 1}{G_{\text{LNA}}} \quad (5.19)$$

where F_r is the receiver noise figure, F_{LNA} is the noise figure of the LNA, F_{chain} is the noise figure of all other components in the chain and G_{LNA} is the gain of the LNA. Expanding on this definition further, if we assume that the impedances are matched at each stage, the equivalent system temperature can be expressed as:

$$T_{\text{sys}} = \frac{T_1}{G_1} + \frac{T_2}{G_1 \times G_2} + \frac{T_3}{G_1 \times G_2 \times G_3} + \dots \quad (5.20)$$

The subsequent running gain, R_G is calculated at each n^{th} stage by:

$$R_G(n+1) = G(n+1) * R_G(n), \quad (5.21)$$

where $G(n+1)$ is the gain of the $(n+1)^{\text{th}}$ component in the receiver chain, and $R_G(n)$ is the running gain at the previous stage in the component chain.

Each component contributes an added noise power, independent of its point in the chain. This is given by N and is dependent on the physical temperature of the component (T_{physical}), its gain G measured in Watts and its noise temperature T_{noise} .

$$N = T_{\text{physical}} * (1 - G) + T_{\text{noise}} * G \quad (5.22)$$

Note, that the physical temperature of any high gain component such as the LNA, is set to zero as it is the noise temperature that contributes to the overall noise power. As a result, when the gain, $G > 1$ then the physical temperature will be zero, meaning only the noise temperature term contributes. The value we are interested in is the output power at the n^{th} stage. The resulting power at the end of chain gives the receiver output power. This is the value that must match the *rms* power attainable by the ADC.

$$P_n = 10 \log_{10}(N_n k \Delta f) \quad (5.23)$$

The Y -factor method combines the gain and noise temperature of each component in the chain to find an equivalent noise temperature. The Y factor after the final component in the chain is analogous to the final system temperature, T_{sys} .

$$Y_{n+1} = \sum_1^{n+1} N_{n+1}/G_{n+1} + T_n \quad (5.24)$$

The value of T_{sys} can hence be used to obtain a sensitivity level estimate of the receiver. A detailed receiver model is described in the following Chapter on system temperature generation.

5.5.2 Components

To investigate the proof of concept, I have identified a number of available components that can be used for two potential receiver designs. These are based on pre-existing systems such as C-BASS, GHY-3 and the SKA. In Table 5.1, I list the components that can be used for the X-BASS and NextBASS receivers.

Using the data sheets for the Low Noise Factory 6 – 20 GHz and 15 – 29 GHz amplifiers I have the predicted noise temperature and gain of the amplifiers. These values make up the frequency dependent noise figures for the LNA used in the receiver model detailed in Chapter 6, Section 6.2.1. Figure 5.9 shows the LNA noise temperature taken from the Low Noise Factory Data sheets for each LNA in the NextBASS receiver (Low Noise and Factory (2016), Low Noise Factory (2017)).

The X-BASS FPGA will be the RFSoc, which has been implemented on GHY-3 (XILINX (2018), Pollak (2018)). The ADC that accompanies this is a 8×4 GSPS 12-bit ADC, providing 16 GHz of bandwidth, which is twice that required for X-BASS. One option for the NextBASS FPGA is the Intel Stratix GX (see INTEL (2018)), an FPGA being developed for use with the SKA by Del Rizzo and Garstin (2018). An ADC that could accompany this FPGA is the HMCAD5831 ADC (Hittite Microwave Corporation (2016)), operating with 3.5 bit at 26 Gbps ADC. This is advantageous for fast real-time processing. The 3.5 bit sampling is sufficient for the low SNR made by a CMB radiometer experiment. It has approximately 8 levels, when half of those are filled by the continuous wave calibration signal, then there are enough levels to make a detection. However, due to the small number of bits, there may be additional quantisation noise losses. By ensuring the receiver generates the correct

Part	X-BASS	NextBASS ¹	NextBASS ²
LNA	LNF-LNC6-20C	LNF-LNC6-20C	LNF-LNC15-29B
f_{RF}	7-15 GHz	7-15 GHz	15 - 30 GHz
BPF	(7-15) GHz	(7-15) GHz	(15-30) GHz
RF AMP	AOX 20	AOX 20	HMC994A
MIXER	HMC1048A	HMC773A	HMC292
LO	(0.05-20) GHz ASY-0050-2000-10	(0.05-20)GHz ASY-0050-2000-10	ASY (up to 32 GHz)
90° Hybrid Coupler	(4-18) GHz AS7984	N/A	N/A
f_{LO}	(9 & 13) GHz	16.5 GHz	(16.5 & 31.5) GHz
f_{IF}	$\pm 2 \times 2$ (8 GHz IF)	0-10 GHz SSB	$2 \times (1.5-9.5)$ GHz
AAF	(0-2)GHz	(0-9.5)GHz	(0-9.5)GHz
IF AMP	C-BASS IF	AOX 20	AOX 20
DBE	RF Soc	Intel 10 Stratix FPGA	Intel 10 Stratix FPGA
ADC	(RF SoC) 12 bit, 4 Gbps	(HMCAD5831) 3.5 bit, 26 Gbps	(HMCAD5831) 3.5 bit, 26 Gbps

Table 5.1: X-BASS and NextBASS receiver components based on the C-BASS and Goonhilly GHY-3 architectures (See [King et al. \(2014\)](#) and [Pollak \(2018\)](#)). The other components have been chosen based on data sheets of available parts (see [Low Noise and Factory \(2016\)](#), [Low Noise Factory \(2017\)](#)), [AtlanTecRF \(2019b\)](#), [Analogue Devies \(2015\)](#), [Analogue Devies \(2018\)](#), [AtlanTecRF \(2019a\)](#), [Hittite Microwave Corporation \(2016\)](#), [INTEL \(2018\)](#), [XILINX \(2018\)](#), [Analog Devices \(2015\)](#)). The FPGA choice for NextBASS comes from work currently being carried out for the SKA (see [Del Rizzo and Garstin \(2018\)](#)).




Figure 5.9: The Noise temperature and gain measurements taken from the Low Noise Factory Data Sheets for the two LNA's required for the two NextBASS Channels. These are the LNF-LNC6-20C, 6 to 20 GHz LNA and the LNF-LNC15-29B, 15 to 29 GHz LNA. (See [Low Noise and Factory \(2016\)](#) and [Low Noise Factory \(2017\)](#)).

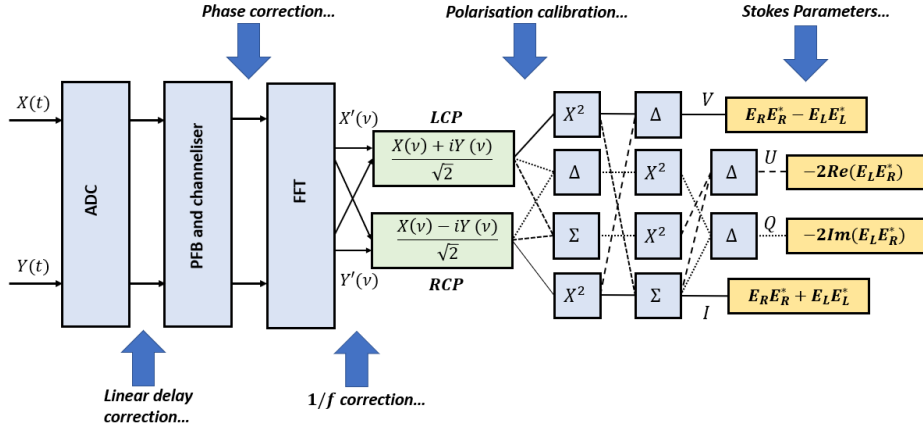


Figure 5.10: A conceptual FPGA architecture that could be used for the X-BASS and NextBASS DBE. The FPGA must channelise, filter and calibrate the receiver. In the X-BASS and NextBASS architectures, the FPGA must also generate the Stokes parameters. The generation of Stokes parameters can be done after the gain calibration. The calibration should be done in real-time, however the generation of the Stokes parameters can be done in post-processing.

output power, then saturation can be avoided, hence reducing quantisation noise. The additional quantisation noise behaves as an additional increase in system temperature and should not result in data degradation. This is a trade-off between faster sampling rates to process broad bandwidth data in real-time.

I have calculated the estimated output power, P_{rms} required to effectively sample the chain. In Table 5.2, I list the specifications of both ADCs for X-BASS and NextBASS. The difference in the dynamic range of the ADC will result in differences in quantisation noise between X-BASS and NextBASS. However, NextBASS will be able to achieve real-time data calibration where as X-BASS will not be able to. These designs outline two different approaches that can be made for the NextBASS experiments. The X-BASS approach which makes use of a large dynamic range to make more sensitive measurements. The NextBASS approach which utilises smaller bits, meaning more data can be taken and as a result, faster map making along with more accurate data processing and calibration. NextBASS has been chosen to use the faster FPGA to demonstrate the need for better calibration methods during the map making process.

	X-BASS	NextBASS
DBE	Xilinx RFSoc	HMCAD5831
Bits	12 bit	3.5 bit
Sample rate	4 Gbps	26 Gbps
Bias Voltage	925mV	256 mV
Total dynamic range	74 dBm	20 dBm
P_{rms}	0.3 dBm	-8 dBm

Table 5.2: ADC specifications for the Xilinx RFSoc, 8×4 Gbps, 12 bit ADC (XILINX (2018)) and the HMCAD5831, 3.5 bit, 26 Gbps ADC (Hittite Microwave Corporation (2016)). The Xilinx RFSoc is the choice of ADC for the X-BASS receiver, whereas the HMCAD5831 offers 10 GHz of bandwidth and is the prime candidate for the NextBASS receiver.

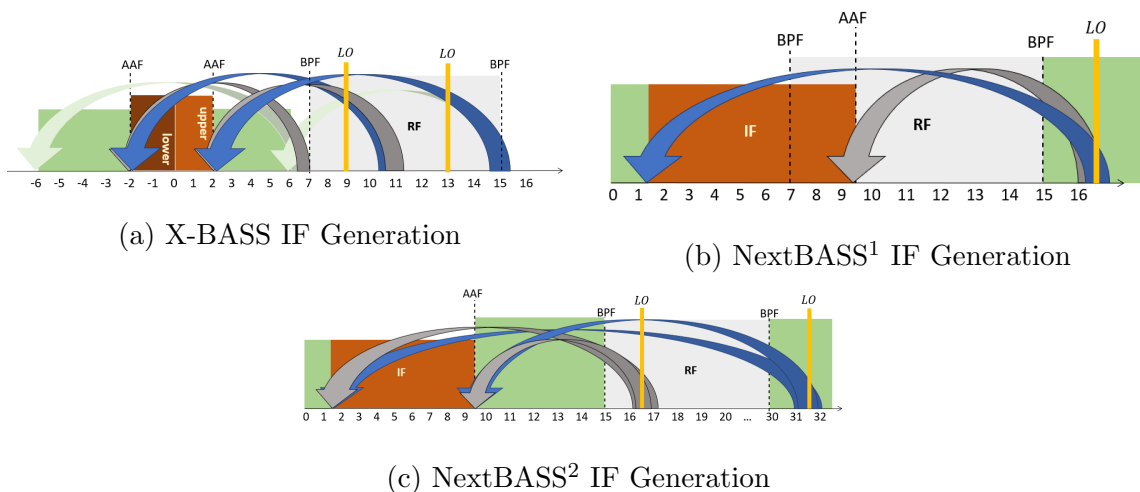


Figure 5.11: The generation of the IF channels for X-BASS, NextBASS¹ and NextBASS². The X-BASS channel generates DSB, whereas both NextBASS channels use a SSB. The LO frequencies for X-BASS are at 9 and 13 GHz. For NextBASS the LO is 16.5 GHz for band-1, and for band-2, the band is split and 2 LO's are used, at frequencies 16.5 and 31.5 GHz. The orange squares indicate the IF bandwidth, the grey squares indicate the RF bandwidth and the green bands indicate the out-of-band frequency regions.

The IF bands for both X-BASS and NextBASS are given in Table 5.1. X-BASS utilises the $2 \times \pm 2$ GHz of bandwidth available by the RFSoc ADCs by employing a DSB architecture. There will be two mixing frequencies, one at 9 GHz and another at 11 GHz. NextBASS will generate a SSB, generating an IF band of 1.5 - 9.5 GHz. The anti-aliasing filter will therefore be a low-pass filter with a cut-off frequency of 9.5 GHz. NextBASS band-2 will have to use two LO frequencies to sample the 15 GHz of RF bandwidth. Figure 5.11 shows the IF generation for X-BASS and NextBASS,

including the IF band and AAF cut-off frequencies, along with the RF bands, LO frequencies and the BPF cut-off.

5.5.3 X-BASS Architecture

The DSB IF generation is only possible if the power output of the RF channel is split in order to be combined with the two LOs. The splitting of the RF channels results in an overall $\frac{1}{2}$ amplitude loss, which corresponds to an $\frac{1}{4}$ of full power over each side-band. However, since the ADC has a two input channel, both side-bands are then combined and the amplitude reduction is limited to $\frac{1}{2}$. I give the full schematic of this receiver architecture in Figure 5.12.

5.5.4 NextBASS Architecture

The NextBASS architecture differs from X-BASS as it will use an SSB architecture. The NextBASS ADC ([Hittite Microwave Corporation \(2016\)](#)) has up to 10 GHz of bandwidth which matches the DC-10 GHz IF band of NextBASS. The choice of IF band for NextBASS was kept fixed for both bands so that the same anti-aliasing filter can be used. For NextBASS¹, a mixing frequency of 16.5 GHz results in an IF band of 1.5 – 9.5 GHz. This is ideal as the LO frequency is both out of the RF and IF band.

The difficulty arises with NextBASS² as it becomes problematic to decrease the bandwidth and find a mixing frequency that is out of the RF and IF bands. To solve this, I suggest splitting the channel into two and use different local-oscillator (LO) frequencies to produce two side-bands with the same IF response. An $f_{LO} = 13.5$ GHz and $f_{LO} = 31.5$ GHz produces the desired IF band for 15-23 GHz and 22-30 GHz respectively. These choices of LO do not experience overlap from lower or higher frequencies, as they will be filtered out by both the 15-30 GHz BPFs and AAFs. The splitting can either take place as two separate observations or through an additional power split with two different LO's. The result will be a factor of $\sqrt{2}$ loss in sensitivity due to splitting the observations.

For simplicity, I have assumed that there will be no power splitting and hence the receiver architecture can remain the same for band 1 and band 2. In Figure 5.13 I have drafted a schematic diagram of the NextBASS¹ and NextBASS² architectures. This includes the use of attenuators to match the P_{rms} of -8.0 dBm necessary for the HMCAD5831 ADC.

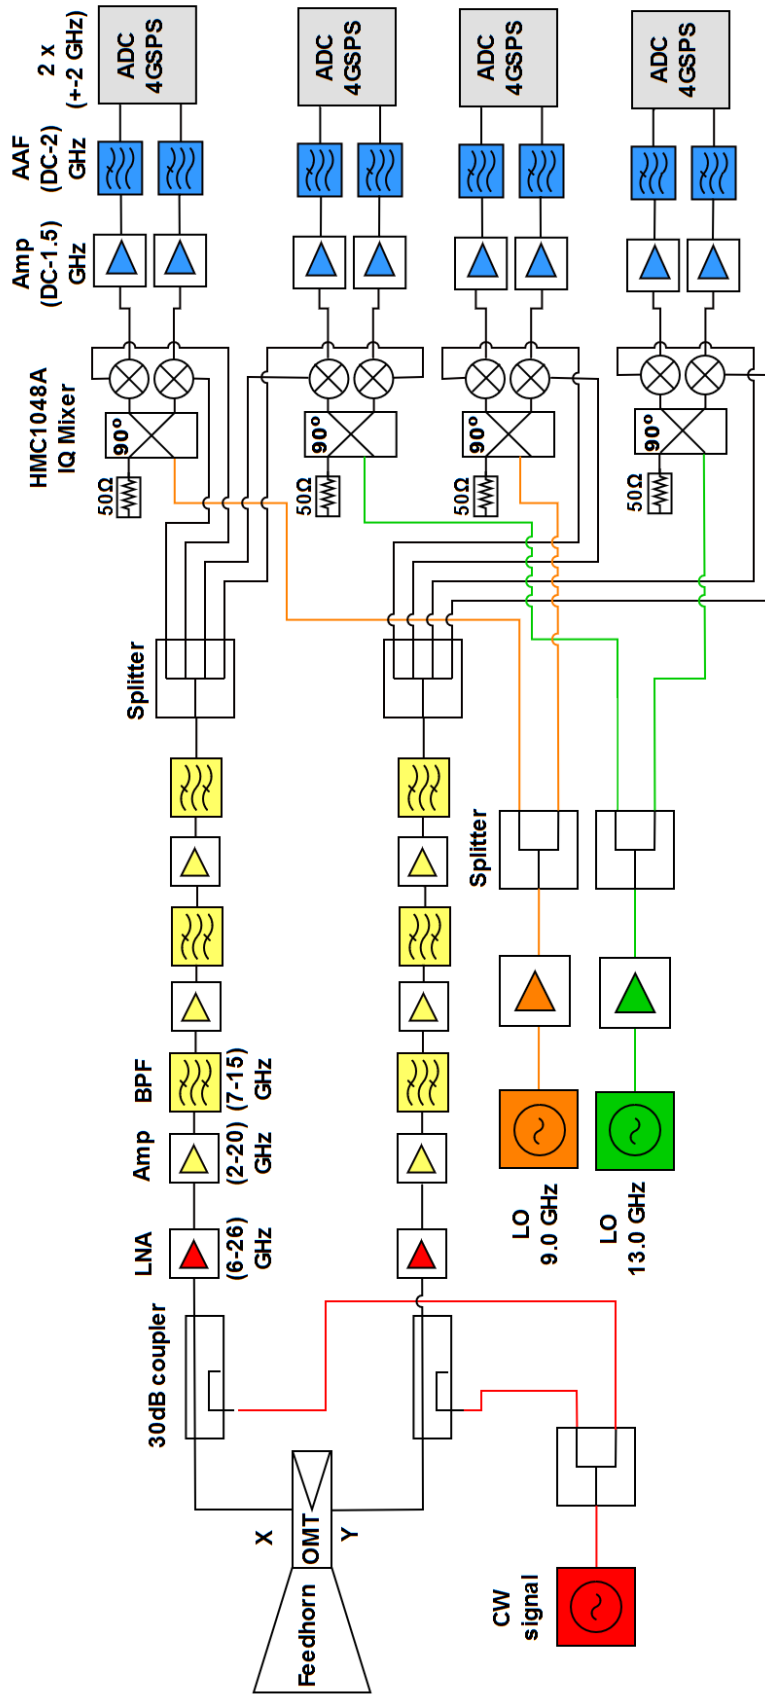
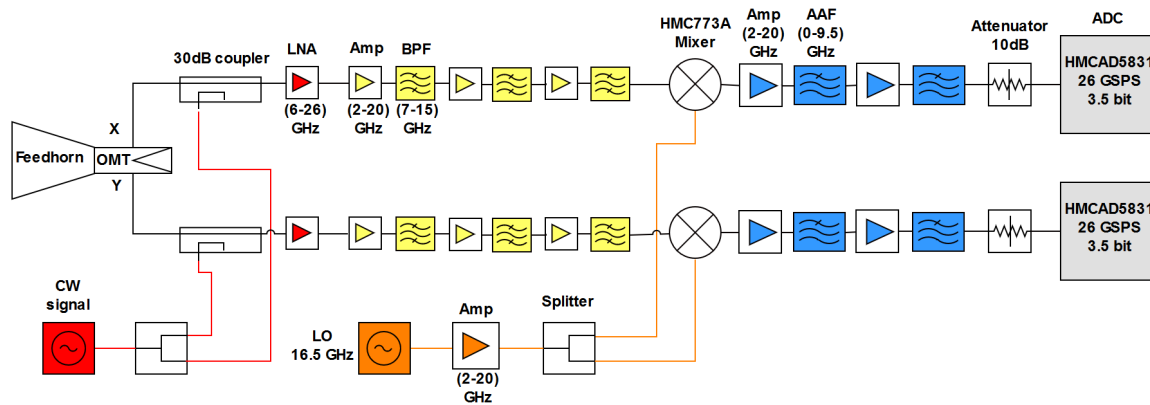
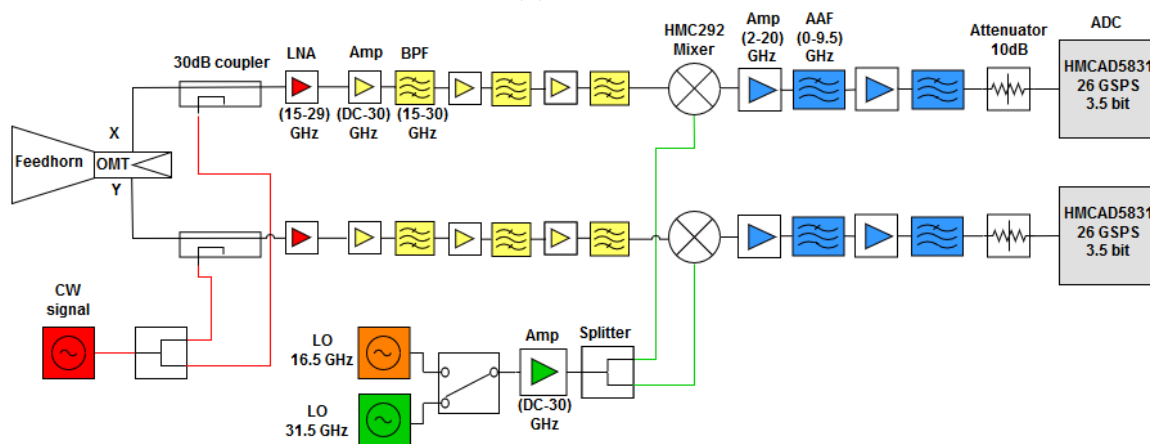


Figure 5.12: X-BASS Receiver Architecture. This employs a continuous wave method of gain stabilisation indicated by the CW source coupled in before the LNA. Each RF polarisation chain is split into four and mixed with the two LO signals using two IQ mixers and a 90° hybrid. The two IF channels from each set of mixers are then amplified and anti-aliased before entering the 2 channel ADCs.



(a) Band 1



(b) Band 2

Figure 5.13: NextBASS Receiver Architectures: Both channels use the continuous wave method of gain stabilisation coupled into the RF chain before the LNA. There is no power splitting for the RF or IF chains. The NextBASS² receiver requires a switch to control the LO in order to sample the full band to match the IF bandwidth of 10 GHz.

5.5.5 Cryostat

The cryogenic cooling stage must begin at the throat of the feed horn and the OMT will be within the cryostat. The feed array of NextBASS creates a problem that each individual feed cannot have its own cryostat. To solve this problem it is possible to create a cryostat that holds the full feed array. I suggest placing the feed array into a vacuum sealed cryostat with the OMT connected to the coolest cryogenic stage. This would require the use of a vacuum window that will not bow under atmospheric pressure and will also remain transparent to radio signals. A commonly used material for this purpose is high density polyethylene (HDPE). A similar design has been carried out in the QUIET receiver, which uses an Ultra-high molecular

weight polyethylene (UHMWPE) which was chosen for its high stress-tolerance. The cryostat consists of the window with an anti-reflection coating of expanded teflon and an array of 91 corrugated W-band feeds and a further 19 horns in the Q-band (Bischoff (2013)). The QUIET cryostats are ≈ 56 cm in diameter, for a focal-plane of ≈ 1.2 m as required for NextBASS, this will have to be twice the size. The greater the volume of the cryostat, the longer the pump-down time from atmospheric pressure. The surface area of a cryostat must be minimised in order to reduce pressure drops due to water molecules that have been bound to the inside surfaces (Ekin (2007)). In Figure 5.14 I have presented a conceptual design of the NextBASS receiver with a side profile view of the band-2 feed array and the single, band-1 feed.




Figure 5.14: Conceptual design for the NextBASS cryostat. The feed array is housed in a box cryostat with the end of each horn entering the cryogenic stage. The horns are located within a vacuum using a radio-transparent window of HDPE.

5.6 Filter Design

The broadband receiver architectures demonstrated above require good filtering to achieve a sharp cut-off and prevent spurious signals from out-of-band frequencies. Since the bands are specialised, it is not trivial to produce large bandwidth filters. In this section, I describe the process of designing the filters required for the NextBASS experiment.

5.6.1 Microwave Filters

A microwave filter is a way of controlling the frequency response in a circuit providing transmission at passband frequencies and attenuation elsewhere. By considering the insertion loss, it is possible to design a filter with an ideal response. This method generates a periodic structure of lumped LC elements.

A good receiver should be linear, meaning that it amplifies signals within the bandpass with no distortion. For a simple linear receiver, the input signal is amplified and becomes the input voltage, $v_i(t)$, this is then filtered by a passive filter, creating a new signal attenuated at frequencies outside of the band. Finally, the signal is amplified once more to give an output voltage, $v_o(t)$. The loss of a system is minimised when the impedance of the input load is well-matched to the output load. A matched impedance occurs when the impedances are the complex conjugates of one another, $Z_i = Z_o^*$ (Balanis (2005)).

If the impedances are not well matched, then there will be a reduction in power across the functional bandwidth of the component. This loss is characterised by the insertion loss of a system. The insertion loss of a component defines the ratio of input power to output power. Losses may be due to charge carriers reflected at the input of a component, dissipated around the conducting surfaces of a connector or within the dielectric materials (Pozar (2005a)).

$$\text{Insertion Loss (dB)} = 10 * \log_{10}\left(\frac{P_o}{P_i}\right) \quad (5.25)$$

A lumped element model is a simplified description of an electrical system. This makes assumptions about the component and creates an idealised equivalent resistor, capacitor and inductor connected by perfectly conducting wires. At high frequencies, standard circuit theory is just an approximation, therefore lumped element circuits are not applicable for microwave frequencies. Instead, microwave components are often distributed elements, described by transmission line theory.

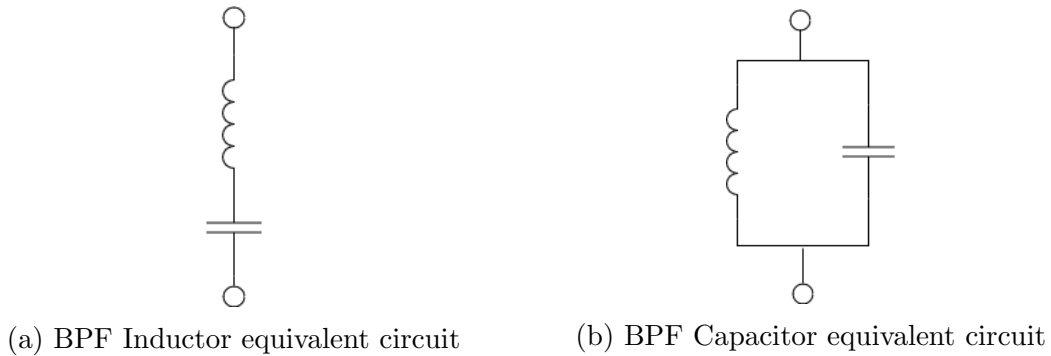


Figure 5.15: Bandpass filter lumped element models.

The difference between circuit theory and transmission line theory is the electrical size. The former assumes that the physical dimensions of a network are much smaller than the electrical wavelength, while the latter uses transmission lines that are close to a wavelength or many wavelengths in size. Distributed networks have voltages and currents that can vary in magnitude and phase.

An infinitely long transmission line loaded with loading elements is a periodic structure. The loading elements for a repeating transmission line are in actuality discontinuities in the line, however they can be modelled as lumped elements. The periodic structures support slow-wave propagation giving them passband and stop-band characteristics. This makes them a good choice for microwave bandpass filters (Pozar (2005a)).

Microstrip is one of the most popular types of planar transmission lines as it can be manufactured easily and can be integrated into both passive and active circuits. The structure of a section of microstrip is a conductor of width W printed on a thin, grounded dielectric substrate of thickness d and relative permittivity, ϵ_r . The effective dielectric constant of a microstrip line is given by:

$$\epsilon_e = \frac{\epsilon_r + 1}{2} \frac{\epsilon_r - 1}{2} \frac{1}{\sqrt{1 + 12d/W}}. \quad (5.26)$$

By applying a transform (such as the Richard's transform and Kuroda identities, see (Pozar (2005b))), it is possible to convert the lumped element circuit into transmission lines and shorted stubs, which can be realised on microstrip. As a result, for a lowpass filter, an inductor can be the equivalent of an short-circuited stub and a capacitor, an open-circuited stub. For a bandpass filter, the short and open stubs are defined as their equivalent inductor/capacitor circuits.

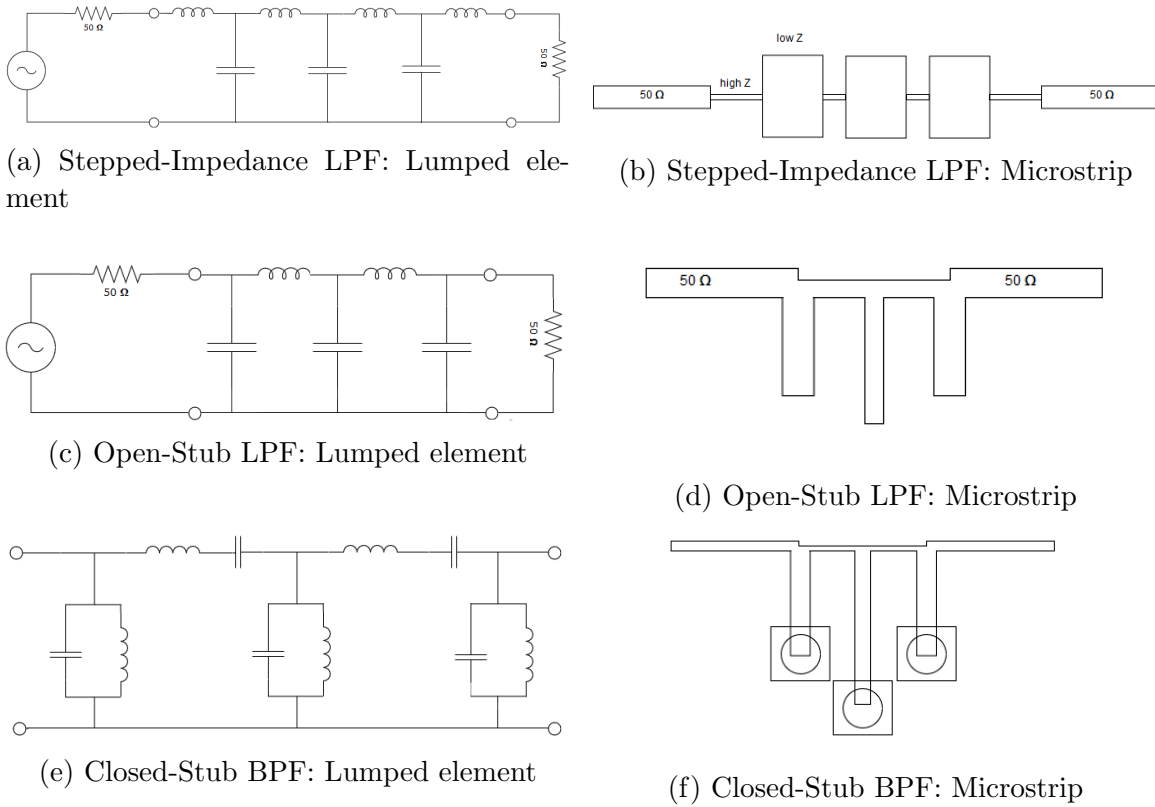


Figure 5.16: Equivalent lumped element model and the corresponding microstrip filter design (Pozar (2005b)).

5.6.2 Design Process

The ideal filter choice for broadband filters are Chebyshev filters. They optimise the bandwidth by introducing a passband ripple. The effects of ripples can be overcome by stacking multiple poles to improve the insertion loss across the band. This is a common method used in radio receivers to improve the band (Pozar (2005b)). The Chebyshev polynomial of the first kind defines the structure of the filter, and is a set of orthogonal polynomials defined as the solutions to Chebyshev differential equation, T_n . It is given by the recurrence relation such that:

$$T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x) \tag{5.27}$$

where T_{n+1} is the value of the polynomial at the $(n + 1)^{\text{th}}$ iteration, n is a positive integer and x is any real number. The insertion loss of this filter is then given by:

$$P_{LR} = 1 + k^2 T_n^2 \left(\frac{\omega}{\omega_c} \right), \tag{5.28}$$

where P_{LR} is the insertion loss, k is the characteristic impedance, T_n is the Chebyshev polynomial of order n , and ω_c is the cut-off frequency.

The filter design on microstrip requires shorted stubs separated by transmission line. The natural frequency of the filter design is given by $\omega = \sqrt{(\omega_1\omega_2)}$. The filter is designed for a bandwidth of $\omega_2 - \omega_1$. Using open stubs gives a wider bandwidth along with a better return loss however, resonances are introduced at lower frequencies with open stops. To overcome this, a closed stub design has been implemented, this is done by halving the length of the open stub, and using a via to close the stop.

The standard substrate used in the manufacturing of microstrip filters is a dielectric material coated on a thin layer of copper. The process of etching a filter design onto this substrate requires that the tracks are no thinner than the thickness of the copper layer. For NextBASS, I have designed a 7-15 GHz bandpass filter, a 15-30 GHz bandpass filter and a single anti-aliasing filters that, through careful choice of the mixing frequency, matches both lower and higher channels. Due to the broadband of the bandpass filters, I found that the impedance of the transmission line sections was either much too low or high to be manufactured. To compensate for this I chose a substrate with a higher dielectric constant. In Table 5.3 I list the substrate used for each filter design.

Filter	Dielectric constant	Material	Thickness	Thickness (Cu)
AAF	3.66	Rogers R04350	0.254 mm	0.017 5mm
Band 1	10.2	Rogers 6010-2LM	0.254 mm	0.0178 mm
Band 2	10.2	Rogers 6010-2LM	0.254 mm	0.0178 mm

Table 5.3: Materials and thicknesses used in the manufacturing of microwave filters.

5.6.3 Anti-Aliasing Filter

The anti-aliasing filter makes use of the fact that the cut-off for the 2-20 GHz LNA is about 1.5 GHz meaning that the AAF can be designed to be a sharp-roll-off low pass filter. For this filter I chose a stepped impedance design which alternates between high, Z_H , and low, Z_L , impedances. By setting $Z_H = 100\Omega$ and $Z_L = 20\Omega$, I was able to obtain the maximum difference in impedance possible for manufacturing on the Rogers R04350 substrate. A high order, n , filter means a sharper roll-off. The anti-aliasing filter has been designed as a low pass filter with a cut-off frequency of 9.5 GHz.




Figure 5.17: The anti-aliasing filter realised in microstrip placed on the testing jig.

The final design uses a Chebyshev order of 21 and a ripple of 0.01 dB, achieving -2.6 dB at 9.5 GHz and -21 dB at 10 GHz with a return loss of less than -20 dB across the whole band. A comparison between the simulated and measured performances are given in Figure 5.18. The return loss is good across the band however the prototype filter has a less steep cut off at a shifted frequency of 10.0 GHz. This shift is due to Ohmic losses from the copper material. This shift should not cause problems within the IF band and can be reducing by stacking the filters.




Figure 5.18: Simulated and measured data for the NextBASS prototype anti-aliasing filter. This measurement was carried out using a test jig and the Anritsu 37369C VNA calibrated to its coaxial ports. The resonance feature at 12 GHz could be minimised by placing the filter in a box containing absorbing material. The S -parameters S_{11} and S_{12} represent the reflected component, or return loss and the insertion loss respectively. Ideally a target return loss is below -15 dB and the insertion loss no less than -3 dB. The proto-type AAF achieves these levels.

5.6.4 7 – 15 GHz Bandpass Filter

The NextBASS bandpass filters consist of a Chebyshev BPF filter and a LPF to cut-off any Nyquist frequencies. I found that at higher frequencies, it was difficult to create a stepped-impedance style LPF since the value of Z_H became too high to manufacture. Instead, I used stubs separated by transmission lines, terminated at 50Ω . I was able to generate an LPF design using the filter wizard in ANSYS Designer 8.0 (ANSYS Inc. (2012)). The impedance of the last two transmission lines was above the manufacturing limit, so I removed these, leaving a order, $n = 4$, filter with a cut-off frequency of 28 GHz.



Figure 5.19: The 7-15 GHz bandpass filter realised in micro-strip.

The prototype 7 – 15 GHz bandpass filter has been made and measured with the Anritsu 37369C VNA up to 20 GHz. This measurement showed a good alignment of the passband frequencies however the results showed that there were ripples present in the passband. Since the measurement was carried out in the jig, some resonances would be expected to be present but these ripples are likely to be a result of manufacturing fault. The thin transmission lines mean there is the possibility of faults during manufacturing. In Figure 5.21 I present a collection of faulty tracks taken using a KEYENCE VHX-500F Digital microscope. The majority of these are breaks in the tracks which meant that particular filter was unusable. In image 5.21a, the track is uneven which could be the cause of the absorption features present in the measurement. The presence of these faulty tracks means that more work needs to be done on the filter design to obtain better prototypes. This may involve either changing the type of filter or using a different manufacturing method.




Figure 5.20: 7 – 15 GHz simulated and VNA data. The large spikes may be due to the lack of absorbing material places around the filter, however it is likely that the poor performance arises from the thin track-widths, this could be overcome by using a substrate with a thinner layer of Cu during the etching process. Otherwise the band cut-off is well matched to the design. The S -parameters S_{11} and S_{12} represent the return loss and the insertion loss respectively. Ideally a target return loss is below -15 dB and the insertion loss no less than -3 dB. This proto-type 7 – 15 GHz bandpass filter does not achieve this and hence should be re-designed to use a better substrate to minimise poor performance due to manufacturing limitations.




Figure 5.21: Faulty tracks on 7 – 15 GHz filter prototypes. This is due to limitations in the etching process. An improvement that could be made is to use a material with a thinner copper layer to improve etching precision.

5.6.5 15 – 30 GHz Bandpass Filter

The 15-30 GHz bandpass filter has been designed with the same method as the 7-15 GHz filter. It is a tenth order Chebyshev bandpass filter designed between 15.5 to 31 GHz. The low-pass filter that goes with it is a 6 order Chebyshev with a 0.0001 ripple and a 28 GHz cut-off. The ripple defines the level of periodic variation. A higher ripple can allow for broader bandwidths at a risk of introducing a ripple in the passband, reducing the insertion loss periodically across the band. This filter has not been realised on microstrip after the results from the 7 – 15 GHz bandpass filter. The design of the filter can remain unchanged, however it is now evident that a substrate with a thinner layer of Cu is necessary and hence a different material to the Rogers 6010-2LM must be used as this does not come available with a Cu cladding less than $18\mu\text{m}$ thick (ROGERS corporation (2018)). the simulated results are given in Figure 5.22.




Figure 5.22: The simulated performance of the 15 – 30 GHz bandpass filter. The S -parameters S_{11} and S_{12} represent the return loss and the insertion loss respectively. Designed using ANSYS Inc. (2012). This filter has not yet been realised onto microstrip as the material available is not suitable for the narrow track widths in this filter.

5.7 Further Receiver Development and testing

This chapter has contributed to the initial research required in developing a receiver architecture. This has been based on current methods and technology that is available today, and has explored new technologies that need further development, such as using

a continuous wave gain stabilisation in conjunction with a fast processing FPGA. The development of this method of gain stabilisation for a broad bandwidth can lead to improved FPGA and poly-phase filter algorithms that can optimally calibrate the data in real-time. The algorithms and Digital-Signal-Processing (DSP) techniques required for these receivers is at the forefront of receiver design, and is applicable to all broad-band width applications, and is therefore, a very important piece of technology development.

Unfortunately, due to time constraints in this project, the testing and development of the analogue receiver electronics has been limited. The microstrip filters may need different manufacturing or designs in order to develop functioning proto-types. Furthermore, the other parts that have been mentioned are off-the-shelf components and should be ordered, calibrated and then tested in the field with use of an appropriate feed horn and OMT that match the frequency of NextBASS. The cryostat for NextBASS that has been introduced in this chapter requires much more research regarding the dimensions of these types of cryostat, the type of cooler required, the vacuum and temperature properties of the cryostat and a comparison of the different appropriate window materials.

In the following chapter, I perform an investigation into how the physical properties of the analogue electronics can be used to generate an accurate model of the receiver temperature. The work in Chapter 6 follows on from the Friis equation and receiver design theories discussed in this Chapter. The receiver model in Chapter 6 confirms that the conceptual designs of the receivers for X-BASS and NextBASS are high-performing with low noise figures.

Chapter 6

An Accurate Estimate of Instrument Sensitivity

“ You can’t save time. You can only spend it, but you can spend it wisely or foolishly.”

Benjamin Hoff, *The Tao of Pooh*

It is extremely important that the performance of an instrument is well modelled before carrying out forecasting and simulations to investigate the instruments impact on science. If an instrument is not accurately modelled then there may be unseen problems that undermine the science capabilities. The instrument may be modelled as either under-performing or modelled with sensitivities that are unattainable in practice. In the case of ground-based survey radiometer experiments, there are two main causes of sensitivity degradation, the ambient noise from the atmospheric and ground contributions, and the receiver noise, which is dominated by the LNA.

If too simple a model of either the sky or receiver is assumed, then the instrument may be modelled to seem ideal. In reality, it will not be capable of achieving such a level of performance, especially across a large bandwidth. It is therefore important to accurately model frequency-dependent effects and understand the limitations to forecasting.

In this chapter, I obtain an accurate estimate of the sensitivity by addressing the most important parameters that will considerably affect performance. I begin with atmospheric modelling, where I discuss the idea that even at frequencies with good atmospheric transmission, a simple model should not be assumed as the varying levels of precipitable water vapour in the atmosphere can have a large impact on the overall sensitivity.

I then go on to model the receiver chain, using frequency data for the LNA as covered in the previous chapter. I also identify components that can affect the output system temperature. Using the ADC *rms* power level, derived in Chapter 5, Section 5.1.3, as the input power level of the receiver, I model the receiver chain of NextBASS and investigate ways of maintaining of achieving consistent power levels across the band and investigate ways reducing the overall system temperature.

The choice of scanning strategy, total observing time and other survey parameters such as the number of feeds and the visible sky fraction at each location, all impact the sensitivity of the experiment. The relationship between each of these factors has to be understood in order to correctly model a realistic instrument and find ways of improving the performance. There are also other factors that will make sensitivity prediction even more realistic. One that is often overlooked until the calibration stage of an observation, is the estimation of spillover temperature taken from the actual beam of the telescope optics. To model this, we require a function that accurately describes the discontinuity between sky and ground temperatures. I try to create this model and investigate whether a model this complex will have a considerable change on the overall sensitivity of the X-BASS and NextBASS instruments.

6.1 Atmospheric Modelling

In this section, I describe the radio frequency window and why it is possible to observe between 7 – 30 GHz successfully from the ground. Furthermore, I also address the presence of the 22 GHz water vapour line, which, depending on the amount of precipitable water vapour in the atmosphere for a given location, can cause a rise in the overall system temperature of the receiver, and hence a considerable loss in sensitivity. Therefore, factoring in the water column density becomes crucial to finding a location where the 22 GHz water line is weak, meaning observations at this frequency unaffected. I then describe the mathematics behind estimating the shape and intensity of this water line.

The atmospheric models for the X-BASS and NextBASS site have been obtained using the software AATM, an adaptation of the atmospheric models developed by ALMA for their beam calibration (Nikolic (2005), Pardo et al. (2001) and Curtis et al. (2009)). The resulting sky temperatures have been cross-checked with a model I have developed, based on the paper ITU (2012). I explore the differences in these models for a selection of potential locations and discuss why Klerefontein is a poor location for NextBASS, but a good choice for X-BASS.

6.1.1 Atmospheric Radio Spectrum

The atmospheric radio window allows for ground observations up to 60 GHz, at which begins the Oxygen resonance lines. There is another window between 80 and 118 GHz, where the next O₂ line is located. After the 183 GHz H₂O line, water vapour then dominates making ground observations at sub mm frequencies extremely difficult. The radio window below 60 GHz is still susceptible to the effects of water vapour present in the atmosphere. Although atmospheric conditions are more favourable at lower GHz frequencies, there is still a weak, H₂O absorption line at 23 GHz. Depending on the altitude, local temperature and pressure and the precipitable water vapour in the atmosphere, this line can change intensity and breadth.




Figure 6.1: Atmospheric transmission curves for varying values of water density. Calculated with an atmospheric pressure of $1.013 \times 10^5 \text{ Pa}$ and a temperature of 15°C . The path distance has been set to 1000 km as the majority of water vapour lies beneath this level. This shows that the atmosphere becomes opaque as frequencies approach sub mm.

The specific attenuation describes the attenuation of the atmosphere for the contribution of both dry, and wet air. Between frequencies of 1 and 1000 GHz, it is easy to obtain a fairly accurate estimate of the specific attenuation at a given value of pressure, temperature and humidity from contributions of dry air and water vapour. This can be done by summing the individual resonance lines from Oxygen and water vapour. Near 60 GHz, many Oxygen absorption lines merge together at sea-level pressures to form a single broad absorption band. When going up to higher altitudes, the individual lines become resolvable as the pressure decreases (ITU (2016)). Figure 6.2 shows the specific attenuation of the Earth's atmosphere along with the

attenuation of dry air. This demonstrates that the water vapour content causes the atmosphere to become opaque to microwave radiation after 60 GHz.




Figure 6.2: The specific attenuation of the atmosphere up to 150 GHz, the dry air attenuation is shown in red. This was plotted using the **MATLAB** function **gaspl**. The O_2 lines make it difficult to observe above 60.0 GHz, which is why higher frequency CMB observations must be done in space.

Since the resonance lines and water absorption features are well modelled between 1-1000 GHz, a frequency dependent spectrum can be estimated. In Figure 6.1, I plot the transmission window for different values of water vapour density, ρ_v g/m³. The typical value of ρ_v at sea level is 7.5 g/m³. As expected, transmission is improved for lower water column densities. The atmospheric attenuation, A , is hence related to transmission, t , by $A = -\log t$. The spectra plotted in Figures 6.1 and 6.2 were obtained using the **gaspl** function in MATLAB, a function that is based on the International Telecommunication Union Attenuation by Atmospheric Gases briefing (ITU (2016)).

At these frequencies, the atmospheric absorption features can be modelled by contributions from H_2O and O_2 only. The specific attenuation of the atmosphere, γ , can be expressed as the sum of the dry and wet attenuations as a function of frequency, $\gamma_0(f)$ and $\gamma_w(f)$ respectively. This is related to the complex atmospheric refractive index, $N''(f)$, by (ITU (2016)):

$$\gamma = \gamma_0(f) + \gamma_w(f) = 0.1820fN''(f) \quad (6.1)$$

The complex refractive index is the multiple of the sum of discrete spectral terms within a frequency bandwidth function, and the spectral line strength, S_i . For atmospheric Oxygen, each spectral line strength is given by the relation:

$$S_i(\text{O}_2) = a_1 \times 10^{-7} \left(\frac{300}{T} \right)^3 \exp \left[a_2 \left(1 - \left(\frac{300}{T} \right) \right) \right] P. \quad (6.2)$$

Similarly, for atmospheric water vapour, each spectral line strength is defined by:

$$S_i(\text{H}_2\text{O}) = b_1 \times 10^{-1} \left(\frac{300}{T} \right)^{3.5} \exp \left[b_2 \left(1 - \left(\frac{300}{T} \right) \right) \right] W, \quad (6.3)$$

where P is the atmospheric pressure, W is the water vapour density and T is the ambient temperature.

For each oxygen line, $S_i(\text{O}_2)$ depends on constant a_1 and a_2 , likewise, the water vapour line and the constant b_1 and b_2 . These are frequency dependent constants, the specific values of each are given in the paper ([ITU \(2016\)](#)). To compute the total attenuation for narrowband signals along a path, the specific attenuation is multiplied by the path length R . The total attenuation is then given by $L_g = R(\gamma_0 + \gamma_w)$, where γ_0 and γ_w represent the attenuation of dry and wet air respectively.

6.1.2 Water Vapour

The properties of water vapour in the atmosphere vary from location to location, with temperature, altitude and pressure. Therefore, it is non-trivial to develop a model that correctly estimates the water vapour content. Location information on the water vapour content is usually taken by measurements throughout the year, and at different times of the day. I will now discuss how water vapour contributes to the spectral properties of the atmosphere and how it is modelled and how we can define the amount of water in a column of air to help us improve these models.

The density of dry-air, ρ_0 (g/m^3) can be calculated using the ideal gas law such that $\rho_0 = \frac{p_d}{R_{\text{specific}}T}$. This is in terms of the partial pressure of dry air, p_d (Pa), absolute temperature, T (K) and the specific gas constant for dry air, R_{specific} ($\text{J}/(\text{kg}\cdot\text{K})$). The typical value for the specific gas constant is $287.058\text{J}/(\text{kg}\cdot\text{K})$ ([COESA \(1976\)](#)). The dry air component of the atmosphere can be well described at frequencies below 1000 GHz by the resonant Oxygen lines, pressure-induced Nitrogen and non-resonant Debye spectrum of Oxygen.

The addition of water vapour to air reduces the density of air since the molar mass of water is less than the molar mass of dry air. An ideal gas at a given temperature

and pressure has a constant number of molecules in a particular volume, hence when water molecules are added, the dry air molecules must decrease by the same amount to maintain a constant pressure and temperature. This means the mass per unit volume of the gas decreases. The water vapour density describes the mass of an air column above a certain point. It is measured in g/m^3 and has a value of approximately 7.5g/m^3 for a dry atmosphere at Standard Temperature and Pressure (STP: 1013 Pa, 15°C).

According to the ideal gas law, the density of water vapour, $\rho_v = \frac{e}{R_v T}$, where e , is the partial pressure, T is the temperature at a given altitude and R_v is the specific gas constant for water vapour, $461.495 \text{ J}/(\text{kg}\cdot\text{K})$ (ITU (2016)). The density of humid air can then be described as the mixing of ideal gases by summing the density of dry and wet air components $\rho = \rho_0 + \rho_v$. The partial vapour pressure of water, e , may be calculated from the saturation vapour pressure and relative humidity, W_a at a given site by $e = W_a e^*$.

The relative humidity is not trivial to calculate and often requires making a measurement of conditions at the site. For example, having the same amount of water vapour in cool air results in a relative humidity higher than that in warm air. It is often given as a percentage and can be defined by the ratio of dry, r to wet, r^* mixing ratios. These are the mass of water vapour per unit mass of dry or saturated air respectively, such that $W_a = 100r/r^*$ (Dominguez (2011)).

There are also many published approximations of saturated water vapour pressure of water. One value is the Tetens equation as described in Monteith and Unsworth (2007), which gives an accurate description of the behaviour of saturated air as a function of temperature, T , as:

$$e^* = 0.61078 \exp\left(\frac{17.27T}{T + 237.3}\right). \quad (6.4)$$

The effect of water vapour in the Earth's atmosphere is of particular concern for instruments that require a low noise. On days with high water vapour content, there is a thicker, denser layer of water vapour above the instrument which can result in much higher antenna temperatures and hence a lower sensitivity during that particular observation. When water vapour is present, more radiation is absorbed, which causes the sky temperature to increase. The antenna is then looking at a higher temperature at a given frequency and therefore there is an overall increase in T_{sys} . The precipitable water vapour (PWV) content, measured in mm, can give

an indication into the amount of water vapour in a column of air above a certain location. It is defined as:

$$PWV(z) = \int_z^{z_{max}} \rho_v(z) dz. \quad (6.5)$$

where $\rho_v(z)$ is the density of water vapour at a given height, z , and is given by:

$$\rho_v(z) = \rho_v \exp\left(\frac{-(z - z_0)}{wvl}\right), \quad (6.6)$$

where wvl is the water vapour scale height, the distance at which the water vapour pressure decreases by a factor of e . The water vapour scale height, wvl differs from the atmospheric scale height, which is usually taken as $H = 8.5$ km for a standard atmosphere, wvl is much less than this. As water condenses in the lower layers of the atmosphere, other gases such as Carbon dioxide and methane rise above the water vapour, as a result, the typical value for $wvl = 2$ km.

The total precipitable water vapour can also be expressed in commonly expressed in units of kg/m^2 , however can also be expressed in mm. Assuming a vertical column of air from the ground to the top of the atmosphere with a base of 1 m^2 , the PWV content of this column equal the height of water if all water vapour was condensed such that $1 \text{ kg}/\text{m}^2 = 1 \text{ mm}$.

Since the level of water vapour can change day to day, season to season, using PWV measured at a location is a much more accurate way measuring the atmospheric absorption and hence the sky temperature. As it is so dependent on location and local conditions, it is best measured at a given site over years and then average values taken. It is arguably best to select a site that does not suffer such an extreme variation in water vapour column density.

Measurements of PWV value can be done with Radiosonde and GPS methods ([Angel et al. \(2011\)](#), [Choy et al. \(2015\)](#)). The first uses humidity sensors that are deployed from a weather balloon which can track the relative humidity as a function of altitude. A more accurate method is to use GPS methods which relies on the refractive index of the troposphere, which depends on the air pressure, temperature and water vapour pressure. Variations of this can affect the propagation of electromagnetic waves, inducing a delay in any signal. Using GPS tracking, this delay can be measured and an estimate of the PWV above the site can be estimated ([Angel et al. \(2011\)](#)).

Direct measurements of a site are not always possible and therefore, ground-based and GPS measurements generated by the Global Navigation Satellite System (GNSS) can be used to create models. These global models can be used to estimate the PWV of a site depending on weather data given at different sites ([Choy et al. \(2015\)](#)).

6.1.3 Estimating Sky Temperature

This subsection describes the two methods used in estimating the sky temperature. One uses the software AATM, and the other uses a method outlined in [ITU \(2012\)](#) in order to generate the specific attenuation of the atmosphere, and then scale it to a fixed sky temperature of 240 K.

The software package AATM is a development of ATM, “Atmospheric Transmission at Microwaves” modelling software, written by [Pardo et al. \(2001\)](#) for ALMA. AATM consists of a library of models for transmission through the Earth’s atmosphere, developed to deal with phase changes in ALMA data caused by atmospheric dispersion. [Curtis et al. \(2009\)](#) define a simple model atmosphere consisting of thin slab of water vapour. The refractive index of such an atmosphere can be computed using the Smith-Weintraub equation:

$$N = (n - 1) \times 10^6 = \kappa_1 \frac{P_d}{T} + \kappa_2 \frac{e}{T} + \kappa_3 \frac{e}{T^2}, \quad (6.7)$$

where n is the refractive index at radio frequencies, P_d is the partial pressure of dry gas, e is the partial pressure of water vapour and T is the absolute temperature. [Smith and Weintraub \(1953\)](#) define these absorption coefficients as $\kappa_1 = 77.6$, $\kappa_2 = 72$ and $\kappa_3 = 3.7 \times 10^{-5}$.

AATM only considers the terms due to the water vapour content, excluding the dry air partial pressure term. It uses the Kramers-Krönig relations to describe the spectral dispersion. AATM takes the values for the zenith precipitable water vapour column density, PWV (mm), the ground-level pressure, P_b (mb), the ground level temperature T_g (K), the tropospheric lapse rate, Γ_T (K km⁻¹) which is kept at -5.6 K km⁻¹ and the water vapour scale height, wvl (km), this is kept fixed at $wvl = 2$ km ([Nikolic \(2005\)](#)).

AATM obtains the dry and wet opacities by Bayesian calculation of the coefficients for the water lines. It integrates this over a change in column density to obtain a frequency dependent sky temperature. To cross-check the output of AATM, I have used the ITU model outlined in the paper ([ITU \(2012\)](#)), to generate the level of atmospheric gain for frequencies below 40 GHz. I have fixed the water vapour density to 7.5g/m³ and scaled it to the appropriate altitude using the relation $\rho_w = \rho_w^h \exp(h/2)$, where ρ_w^h is the density of water vapour at altitude h . The factor 2 comes from the scale height of water vapour which is assumed to be 2 km. I then generate the total zenith attenuation in dB, $A(dB)$, by finding the equivalent heights h_0 , h_w of dry and

wet air respectively along with the moist and dry specific attenuation, γ_0 and γ_w respectively.

$$A(dB) = \gamma_0 h_0 + \gamma_w h_w \quad (6.8)$$


The zenith attenuation is attributable to the sky noise, which determines the background level of gain picked up by the antenna. The sky noise in terms of the optical depth can be obtained by the equation $A(dB) = 10 \log_{10}(\exp(-\tau))$. From this the sky temperature can be estimated by using a typical value for the sky temperature in the stratosphere. This is usually given by $T_{atm} = 240$ K. From this, the sky temperature can be found.

In AATM, the CMB temperature is included in the sky temperature, therefore we must add the CMB temperature to the ITU sky to compare the two models. The total sky temperature for the ITU model is given by:


$$T_{sky}(K) = T_{atm}(1 - \exp(-\tau)) + T_{CMB}, \quad (6.9)$$

where T_{atm} is the temperature of the atmosphere, taken to be 240 K, τ is the total opacity of the atmosphere and $T_{CMB} = 2.73$ K.

Figure 6.3 shows that the ITU and AATM models are in good agreement with each other in terms of generating the sky noise and estimating the sky temperature. I find that AATM generates greater values for the sky temperature at higher frequencies, this is due to the inclusion of precipitable water vapour content and the summation over atmospheric layers. AATM will hence provide a more accurate estimate of sky temperature.



(a) Sky Gain [dB]



(b) Sky Temperature [K]

Figure 6.3: The sky gain in dB, and sky temperature calculated for both the AATM, ITU models (Nikolic (2005), ITU (2016)). We can see that both the Atacama and Klerefontein models are in good agreement with some variation in amplitude of the 23 GHz water line and the tail rise of the O₂ resonance lines.

From these comparisons we can see that the precipitable water vapour content is important for estimating the 23 GHz water peak and the overall background level of gain across all frequencies. The tail is underestimated by the ITU model for the Atacama parameters. The AATM generated system temperature will be the first component in the receiver chain when calculating the system temperature.

6.1.4 Location

A good site must have atmospheric conditions that will not impact the quality of data taken, and the conditions must be consistently stable so the majority of data can be kept. It would not be possible to carry out X-BASS or NextBASS from a site in the UK because of the high levels of precipitable water vapour at any given time. The low altitude combined with the high water vapour content results in much higher sky temperatures. This would also be subject to large changes in cloud cover over short time periods so it would be difficult to maintain a consistent sensitivity over a single scan. The UK is also susceptible to long periods of thick cloud cover from low pressure systems.




Figure 6.4: A basic model of the 23 GHz line excluding PWV data, demonstrating how altitude can affect the total sky temperature as a function of frequency.

The C-BASS site at Klerefontein is an easily accessible site with better weather conditions than any in the UK. It performs well at frequencies below the 23 GHz water line due to its higher altitude and its drier atmosphere. The site has fairly stable in weather conditions apart from the occasional thunder storm, which cause PWV values of up to 22 mm and high winds. However, these storms are short lived, predictable and observing does not take place during such conditions.

There are two potential sites in the Canary islands that offer better observing conditions than those of Klerefontein. These are the 'Roque de los Muchachos' observatory in La Palma and the Teide Observatory in Tenerife. The weather in Canary islands is dominated by the North East trade winds which form an inversion layer over the islands. This causes warmer air to hold above the cooler air, which suppresses convection meaning the cloud cover is unable to rise higher than the 1700 m altitude of the inversion layer. Both sites are situated at approximately 2500 m, well above the inversion layer meaning they have favourable atmospheric conditions all year round.

Location	Altitude (m)	P_g (mb)	pwv (mm)	T_g	wvl (km)	Γ_T (Kkm $^{-1}$)
Klerefontein	1371	1005	7	283	2	-5.6
Atacama	4800	560	1	273	2	-5.6
Tenerife	2500	1000	4	283	2	-5.6

Table 6.1: Site specifications used in generation of the frequency dependent opacity and sky temperature data. These values typically represent these sites and come from weather data for each location.

The most optimum location would Antartica, but this is not practical due to

infrastructure limitations. Therefore, the best choice would be the Chajnantor plateau in the Atacama desert in Chile. The region is extremely dry as wet air from the western coast is forced up onto the main land and then trapped due to the Andes mountain range. This means the weather is always dry and the extremely high altitude makes it excellent for observing. The ALMA telescopes are located here so the infrastructure for a new experiment already exists. Comparisons of the site specifications used in conjunction with the AATM software are detailed in Table 6.1.




Figure 6.5: Sky temperature for the three locations, generated using AATM. This includes the CMB temperature.




Figure 6.6: The dry, wet and total opacities for the three sites Klerefontein, Teide and the Atacama plateau. These have been generated using the AATM software.

The wet and dry opacities along with the total opacity for the three locations are given in Figure 6.6. In this Figure, we can see that the Atacama and Tenerife sites

have very similar dry opacities however due to the increased water vapour content and lower altitude, the Tenerife site has a much sharper water line. Klerefontein has a much larger opacity at frequencies above 15 GHz. This results in a sky temperature of up to 14 K higher than the Atacama site, see Figure 6.5.

X-BASS will be located at the current C-BASS site so I will be comparing the weather conditions in Klerefontein to those of an ideal NextBASS experiment located on the Chajnantor plateau. To assess the atmospheric conditions at these sites, reliable estimates of the PWV must be established. Estimates for the Atacama desert indicate a $PWV = 1.2$ mm, which can fluctuate annually by up to 50% (Giovanelli et al. (2001)). The water vapour content at the Klerefontein site can vary throughout the day depending on amount of PWV present. Radiosonde and GPS measurements have not been carried out directly at the C-BASS site, however global meteorological modelling can indicate the amount of total precipitable water across the globe. Data obtained by the Goddard Earth Observing System Model (GEOS-5), commissioned by NASA has been realised by the open source programme ‘earth’ <https://earth.nullschool.net>, which models atmospheric airflow and ocean currents every three hours (Baccario (2018)). Using this I have estimated the PWV value at Klerefontein to be at 7 mm for fair weather days.

Figures 6.7 and 6.8 show the data obtained by the software, ‘earth’ (Baccario (2018)). In these Figures, we can see that Klerefontein has a higher value of total precipitable water, given by TPW in the Figure in kg/m^2 . This indicates that on average, the Atacama site rarely experiences values of $TPW > 1.5\text{kg/m}^2$ ($PWV = 1.5$ mm). Klerefontein on the other hand generally has a surface level TPW of 7 – 10 kg/m^2 , or 7 – 10 mm (PWV).

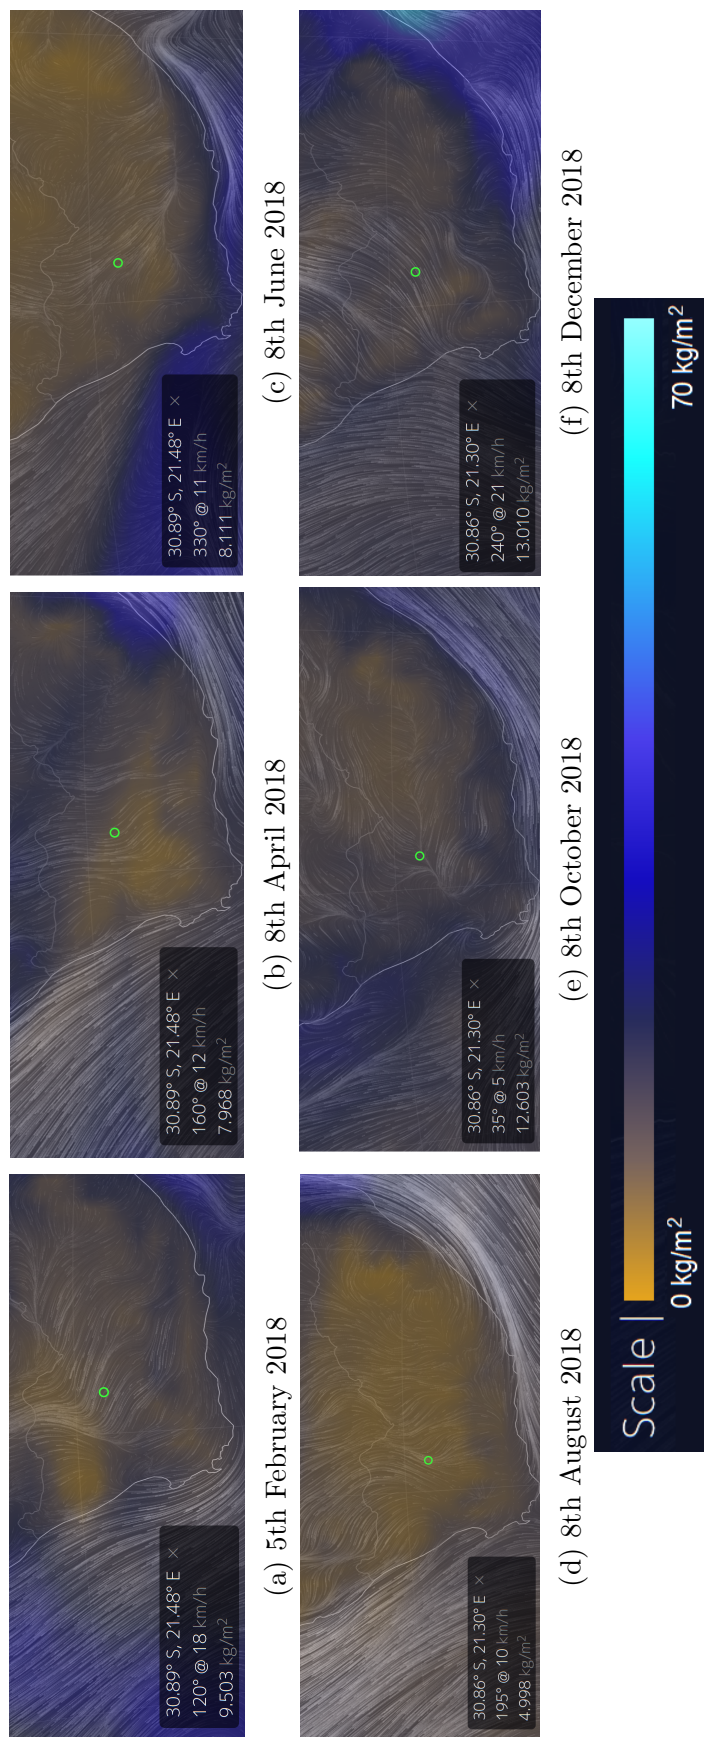


Figure 6.7: Different weather conditions around the Klerefontein, C-BASS site for a selection of dates throughout 2018. These images were taken from <http://earth.nullschool.net>. The contour lines indicate the wind pattern, and the colour scale gives the total precipitable water vapour, *TPW* (*PWV*). The *TPW* at Klerefontein varies between 4 and 15 kg/m² (4 mm to 15 mm). The *TPW* is indicated in the bottom left-hand corner of each plot.

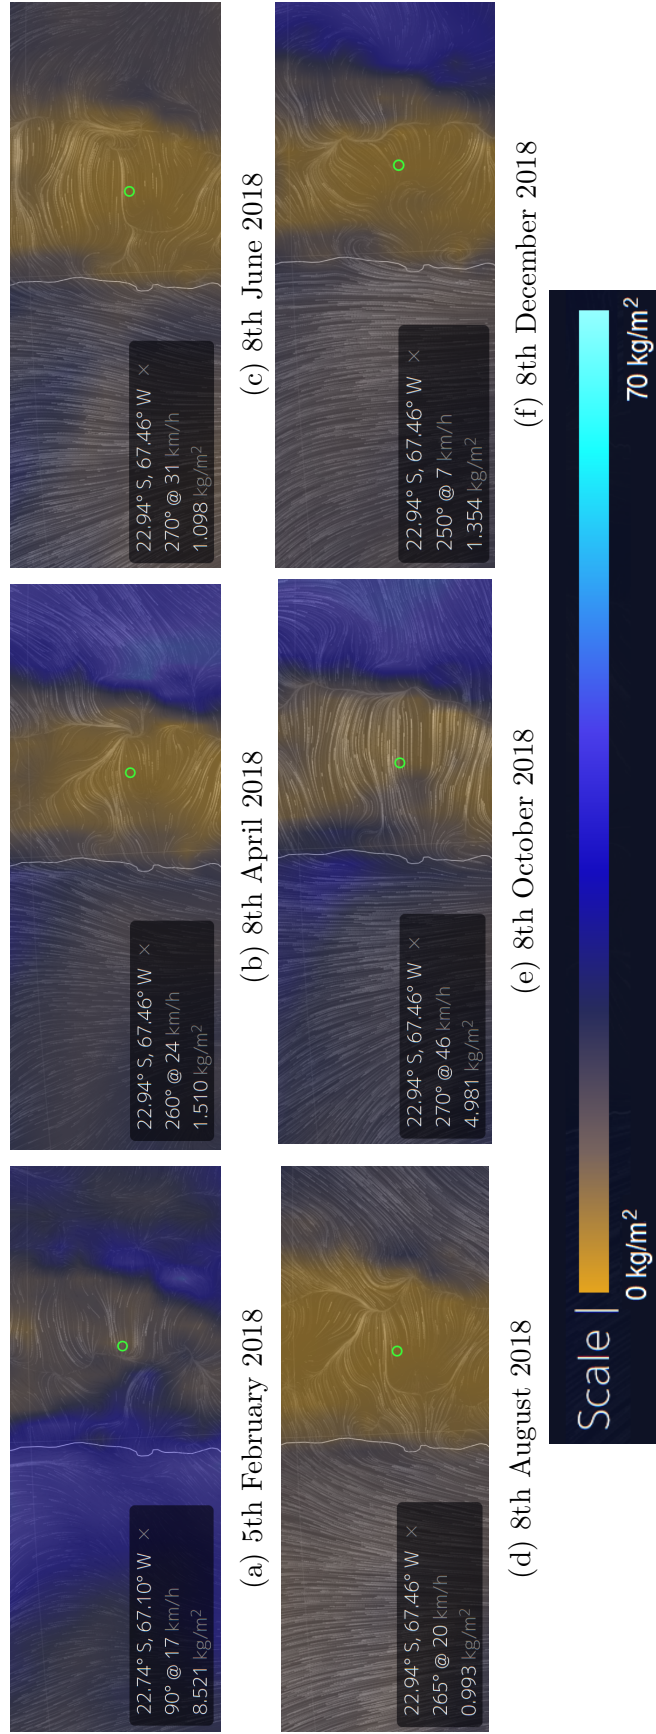


Figure 6.8: Different weather conditions around the Llano de Chajnantor for a selection of dates over the year 2018. These images were taken from <http://earth.nu1school.net>. The contour lines show the wind and their direction, and the colour scale gives the total precipitable water vapour. Due to the Andes mountain range, the *TPW* (or *PWV*) in the Atacama rarely rises above 1.5 kg/m² (1.5 mm). The *TPW* is indicated in the bottom left-hand corner of each plot.

6.2 Receiver System Temperature

In this section, I use the receiver architecture specified in Chapter 5 to develop a frequency dependent model of the receiver chain to generate an accurate estimate of the system temperature across the band. This is based on the stacking of components in the NextBASS receiver architecture, the schematic for which is given in Figure 5.13.

6.2.1 Receiver chain

To establish a cascading receiver model, we have to consider the components in the order at which they appear. The first component is the CMB and the second component would be the contribution from the sky. This propagates into the receiver system via the antennas, the feed horn and the OMT, each contributing to the gain. The first electric component in the receiver chain is the coaxial cable that couples the OMT to the LNA. Following this is the RF and IF filtering and amplification stages. In the case of NextBASS, attenuators will also be required to reduce the output power so as to not saturate the 3.5 bit ADC. The physical temperature of each of these stages is part of the Friis equation given by the equations in Section 5.5.1. The physical temperature, gain and corresponding gain temperature of each component in the chain is given in Table 6.2.

The gain is the figure of merit for noise for each component in the receiver. A component that has a high (positive) gain such as the LNA's, RF and IF amplifiers are described by a noise figure that only considers their electrical contribution, as a result the noise figure need not be described by the physical temperature. Components that are lossy such as the bandpass filters and interlinking cables are described by their physical temperature and their loss (negative gain). Figure 6.9, lists each cascading component the full receiver, with their corresponding gain in dB. In the receiver model, the only frequency dependent components are the sky, of which the gain is given by the atmospheric modelling outlined in the previous section and the LNA, see Figure 5.9.

Component	Gain [dB]	Gain Temperature [K]	T_{phys}
LNA	38	10	5
RF amplifier	14.00	3.0	290
IF amplifier	24.00	3.0	290
Primary mirror	-0.01	0.0	290
Secondary mirror	-0.01	0.0	290
Feed horn (stage 1)	-0.02	0.0	290
Feed horn (stage 2)	-0.02	0.0	70
OMT	-1.00	0.0	5
SMA connector	-0.10	0.0	-
Connector	-0.20	0.0	-
Coax cable	-0.10	0.0	70
1 st Bandpass filter	-3.50	0.0	70
RF Bandpass filter	-3.50	0.0	290
Anti-aliasing filter	-3.50	0.0	290
Mixer	-10.00	0.0	290
Attenuator	-10.00	0.0	290

Table 6.2: Table of gain and gain temperatures used for components in the receiver system temperature modelling. The lossy components have a negative gain contribution and no gain temperature. The physical temperature of the connectors and SMA cables depends on their positioning in the chain. The temperatures are separated into the ambient temperature at 290K, an intermediate cryo-stage temperature of 70K and the cryogenic temperature of 5K. The cryogenic temperature can also be between 10 – 15 K, however the 5 K as stated here has been based on the GM-cooled cryostat for C-BASS, which can be re-used for X-BASS ([Jones et al. \(2018\)](#)).

As shown in Table 6.2, each component has a physical temperature. Following the receiver chain, each component is at a different temperature. The CMB is at 2.731 K, followed by an average sky temperature of 240 K, this however varies band. The horn is split into two stages, an ambient temperature stage of 290 K and a cryogenic stage, assuming an ideal cryostat temperature of 5 K. This cryostat temperature is based on the GM-cooler cryostat for C-BASS, but temperatures between 10 – 15 K are also common cryostat temperatures. In practice, the coax cable, connector and LNA are at the cryogenic temperature. However when calculating the Friis equation, amplifiers only contribute gain temperature. I have also included an intermediate stage at 70 K where the first bandpass filter is located. After the first RF amplifier, the receiver is maintained at the ambient temperature of 290 K. Figure 6.10 shows the component temperatures throughout the chain.




Figure 6.9: The gain of each component in the receiver chain. The lossy components such as the filters give a decrease in gain, whereas the amplifiers provide positive gain. This plot takes the maximum value over the frequency range for the NextBASS¹ frequencies.




Figure 6.10: The physical temperature of components in the gain chain. This starts at the CMB at 2.71 K, then increases to ≈ 5 K at the sky physical temperature. The jump to 290 K represents the ambient temperature of the optics. The components in the cryostat are at 5 K whilst the intermediate stage temperature is 70 K. All amplifiers are set to a physical temperature of zero as only their electrical temperature is relevant for noise calculations.




Figure 6.12: Power at the NextBASS receiver chain. The increase in power across NextBASS¹ is due to a gradual increase in system temperature from a combination of the atmospheric model and the LNA gain. The rise between 20 GHz and 23 GHz is due to the 22 GHz water line.

6.2.3 System Temperature

In this subsection, I will demonstrate how the cryostat temperature can affect the total system temperature. I will also show that the noise figure of the first coaxial cable before the LNA can be a large contributor to increases in system noise. This cable is known as the ‘Golden Cable’, and it is essential that this component is not lossy. The cumulative system temperature can tell us which components contribute most significantly to the system temperature. Figure 6.13 shows that as we work through the sequential components in the receiver, the system temperature changes. The system temperature is obtained by equation 5.24. The system temperature begins at 2.731 K and gradually increases due to the sky signal, which raises the system temperature to a maximum of 11 K for the 22.5 GHz channel. Even though the largest sky contribution occurs at 22.5 GHz, the frequency-dependent LNA gain will have a more considerable contribution, all successive components have a negligible impact. The final output system temperature varies from around 14 K at 7.4 GHz to 26 K at 28.4 GHz.




Figure 6.13: This Figure shows the cumulative system temperature as a function of component in the chain across all the NextBASS frequencies. The behaviour at the start shows that the two greatest contributors to the system temperature are the atmospheric contributions and the LNA. Any component that comes after the LNA has negligible contribution to the system temperature.

To generate the NextBASS system temperature used for the final sensitivity calculations, I have considered cryogenic stages between 5 – 10 K. Cryostat temperatures below 10 K are not essential for radio observations, so PT-coolers (Pulse-Tube) are often used ([Janis Research Company \(2019\)](#)). The C-BASS cryostat has a Gifford-McMahon (GM) cooler, this means it is capable of achieving cooling down to 4 K ([Jones et al. \(2018\)](#)). Since NextBASS is operating at higher frequencies, a GM-cooler may be necessary to reduce the overall system temperature. Figure, [6.14a](#) shows how the system temperature as a function of frequency for three cryostat temperature of 5 K, 7 K and 15 K. The difference between 5 K and 15 K can introduce a rise in the final system temperature of ≈ 4 K.

As explained in Chapter 5 when introducing the Friis equation, any lossy component before the LNA will drastically increase the noise power of the system, hence increasing the system temperature. The coaxial cable before the LNA, known as ‘The Golden Cable’, is a lossy component. If the loss is anything more than 2 dB, then the system temperature increases to a level where the sensitivity of the instrument is no longer useful. Figure [6.14b](#) shows different insertion losses for the ‘Golden Cable’

can change the system temperature. By changing the coaxial isolation from -0.1 dB to -2.0 dB, there is an increase in system temperature of ≈ 15 dB.

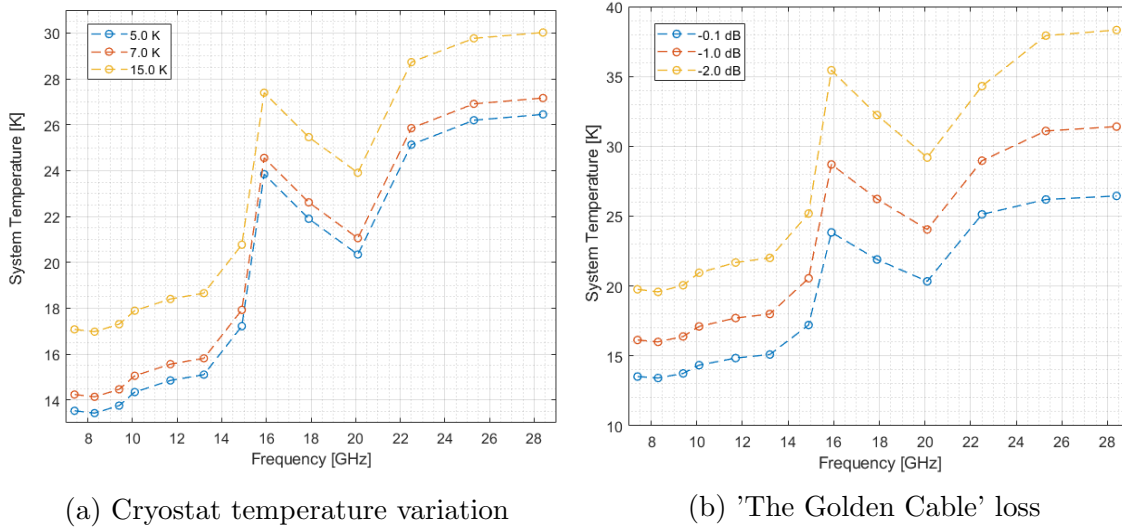


Figure 6.14: Differences in system temperature due to a) temperature changes in the cryostat and b) loss at the coaxial cable positioned before the LNA. These plots demonstrate that the decisions made on the type of cryostat and the coaxial cables can play an important part in the overall system temperature.

6.3 Realistic System Modelling

The receiver system temperature has been generated using a fairly precise method that takes into consideration frequency dependent effects from the amplifiers and the atmospheric effects. However, there are more factors to consider that could degrade performance. One of these factors is the observing elevation. When observing at an elevation that is not the zenith, the beam is looking through a thicker airmass, which would increase the sky temperature.

Another factor that affects noise is the modelling of a realistic telescope beam. The optics have a specific beam pattern that varies in polar coordinates. Across the θ -angle, the beam will see different temperatures, the main beam will see the sky temperature at the elevation it is observing, and the spillover lobes will see a function of both the sky and ground temperatures. I will now go into some discussion about each scenario.

6.3.1 Scanning Strategy

In reality, a telescope that makes a map of the sky does not observe at the zenith as this would not achieve a large enough sky area. Using the scanning strategy of C-BASS as an example, the experiment would slew at a constant elevation while it takes data. This is to minimise atmospheric variation. Then it repeats the same observation at a different elevation to get a full sky coverage. These observations suffer from considerable ground interference when the telescope observes at elevations below 30 degrees. The best angles of observation telescope would be between 45 and 60 degrees (see page 48 (Jew (2017))).

In practice, if observing at a single elevation of θ_{el} then the limits of the integral that determines the sky fraction must change. The upper limit would stay as $2|\delta + \phi|$ and the lower limit would become $|\phi + \theta_{el}|$. The overall effect of this is a reduction in the sky fraction, which is why two observation elevations are usually used.

6.3.2 Observation Elevation

By factoring the observation elevation into the ITU model, it is possible to predict the increase in system temperature. This is the same as scaling the total attenuation by a factor of $1/\sin(e)$ where e is the elevation in degrees. In Figure 6.15, we can see that at an elevation of 60° , there is an increase in sky temperature of 2 K.




Figure 6.15: The sky temperature at Klerefontein showing the 22 GHz water line for different observation elevations. This plot was obtained using the ITU atmospheric model and have been scales by a factor of $1/\sin e$ where e is the observation elevation.

6.3.3 Accurate Brightness Temperature

In the receiver chain, the antenna loss has been assumed to be -0.01 dB since this accommodates for the small electrical losses coming from unwanted currents generated on the surface in the dish. These current have a negligible effect on the electrical loss of the antenna and therefore this value holds valid. However, the telescope optics have an efficiency which determines the brightness temperature of the telescope. This means that we can write the overall system temperature as a combination of CMB, T_{CMB} , atmospheric T_{ATM} , spillover, T_{SL} and receiver, T_R temperatures, given by (Burke and Graham-Smith (2010)):

$$T_{SYS} = T_{CMB} + T_{ATM} + T_{SL} + T_R. \quad (6.10)$$

The spillover noise term can be thought of the as the increase in system temperature as a result of power in the spillover lobes creating an increase in beam power away from the main beam. In order to add this spillover temperature into the model, the overlap integral must be used alongside a ground-sky model that describes the gradient between the ground and sky temperatures as a function of angle θ . This is more accurate than merely multiplying the antenna temperature by the beam efficiency. The resulting spillover temperature could be added into the receiver chain as a component with no electrical properties but having a gain temperature equal to the calculated spillover temperature.

6.3.4 Temperature Integral

While the telescope is observing at some elevation, it is looking through a greater airmass than at the zenith. As its elevation changes the airmass changes, then at 90° it will start looking at the ground. A very simplified model of the sky-ground discontinuity would be to have a step function, where at 90° the brightness temperature becomes the ground. This is too simplified a model and would underestimate the brightness temperature at angles near the main beam, then vastly over estimate the antenna temperature at larger angles.

Instead, the sky-ground model can be split into two parts, a sky component that considers the airmass, and a ground component that is calculated from the reflection coefficients that describe sky reflections off the ground. To find the total brightness temperature, the sky-ground function is multiplied by the polar beam pattern and integrated with respect to the polar coordinates, θ and ϕ .

$$T_A(\nu|\hat{\mathbf{r}}_0) = \frac{\int \int_{4\pi} T_b(\nu, \theta, \phi) P_n(\nu, \theta, \phi|\hat{\mathbf{r}}_0) \sin \theta d\theta d\phi}{\int \int_{4\pi} P_n(\nu, \theta, \phi) \sin \theta d\theta d\phi} \quad (6.11)$$

where $\hat{\mathbf{r}}_0$ denotes the direction that the telescope is pointing in. P_n is the total power (cross-polar + co-polar), T_b is the brightness temperature distribution from the area surrounding the telescope, at a given frequency, ν .

Assuming that galactic emissions are going to cause negligible contributions to noise temperature at these frequencies, and that the CMB has already been incorporated into the sky model as calculated by AATM, then the sky brightness model will be a function of the T_{sky} AATM output over angle plus a ground component.

In the SKA Memo on noise temperature calculations, [Mendellín \(2007\)](#), the ground reflection coefficients and the sky brightness temperature function are derived. I have used these equations to develop a ground function. This memo applies the sky brightness function to a slab atmosphere. The slab atmosphere means that at the horizon, the airmass tends to infinity. However, in reality, the Earth's atmosphere has curvature that results in a finite airmass as the horizon. I have used a method outlined by [Han and Westwater \(2000\)](#), to correct the airmass to account for the Earth's curvature. In the appendix B and C, I have detailed these methods and how they have been used to generate a sky-ground function.




Figure 6.16: Visualisation of the calculation of airmass for curved and slab atmospheric models.

Figure 6.17 shows the ground-sky function developed using the papers outlined above. This function shows the $1/\sin(e)$ function up to 90° , followed by a reflection of the sky temperature on the ground, then the function gradually tends towards the ground brightness temperature. The accuracy of the intensity of this reflection could

be improved upon, however it is a good approximation for this stage of instrument modelling, and is more accurate than using a step-model.




Figure 6.17: Sky Brightness Temperature generated using the reflection coefficients of the ground and the zenith opacities obtained through the AATM model. The sky brightness temperature increases as a sec function of one over the zenith angle. At 90°, there is a sudden drop which is a result of the ground reflecting the sky signal. The function then rolls off to a constant ground temperature.

6.3.5 Antenna Temperature: X-BASS and NextBASS

The integral defined in equation 6.11, can be used with the real optical beam simulated in GRASP (see Chapter 3). If we consider the simple model of antenna temperature (equation 6.12), then we know that the beam efficiency is a good indication of whether the antenna temperature will be a large contribution to overall system temperature. As shown in Chapter 3, both X-BASS and NextBASS achieve beam efficiencies of over 95% across the band, so we expect the realistic antenna temperature to not contribute considerably to the overall system temperature.

$$T_A = \eta_{null} T_{sky} + (1 - \eta_{beam}) * T_{gnd} \quad (6.12)$$

In Table 6.3, I list the T_A values for X-BASS and NextBASS, calculated with the method as outlined above. In this Table, I have also listed the corresponding beam efficiencies and true beam full-width half maxima. This Table demonstrates how both atmospheric contributions and the amount of power in the outer spillover lobes can change the antenna temperature. We know that X-BASS would be subject to poorer atmospheric conditions at the Klerefontein site than NextBASS at a site similar to

ν [GHz]	$T_A(z_{90})$ [K]		$T_A(z_{60})$ [K]		η_{null} [%]	
	X-BASS	NextBASS	X-BASS	NextBASS	X-BASS	NextBASS
7.0	5.6	11.9	7.7	10.8	97.2	95.5
8.0	5.7	8.4	8.4	8.4	96.4	97.4
9.0	5.9	7.6	9.0	7.7	95.5	97.9
10.0	5.8	6.8	8.4	7.1	95.6	98.3
11.0	5.9	6.2	7.8	6.7	97.2	98.7
12.0	6.0	5.3	8.2	6.0	97.4	99.2
13.0	6.2	5.3	8.5	6.1	97.7	99.2
14.0	6.4	5.1	8.6	5.9	97.9	99.3
15.0	6.7	4.9	9.0	5.8	97.8	99.5
15.0	-	12.3	-	11.1	-	95.5
18.0	-	11.4	-	10.8	-	95.9
21.0	-	11.6	-	11.6	-	96.8
24.0	-	10.5	-	10.9	-	97.4
27.0	-	9.9	-	10.3	-	97.4
30.0	-	9.2	-	9.9	-	98.0

Table 6.3: The calculated Antenna temperature for X-BASS and NextBASS for the zenith angle (90.0°) and an elevation of $e = 60.0^\circ$. The beam efficiency has been calculated between the first two nulls of the GRASP beam patterns outlined in Chapter 3. It can be seen that the system temperature follows the trend that the higher the beam efficiency, the less power is in the sidelobes and hence the system temperature is lower.

ALMA, however at the lowest frequency, the NextBASS beam does not achieve a higher beam efficiency and hence suffers from an increase in antenna temperature. At the upper end of the first band however, we see that as beam efficiency increases, the antenna temperature becomes closer to the Zenith sky temperature. The wider beam means there is more power between the first nulls of the beam, making NextBASS achieve a higher beam efficiency.

When comparing the antenna temperature in Table ?? at the zenith to that at an elevation of $e = 60^\circ$, X-BASS shows the trend we expect to see, such that there is a decrease in system temperature of ≈ 2 K at the zenith. However, at the low frequency channels for NextBASS, the zenith temperature has actually increased. This is because the spike that appears in the ground-sky model at approximately 90° is aligned with the large secondary and primary spillover sidelobes of the cross-Dragone optics. At an elevation of $e = 60^\circ$, these spillover lobes miss this reflection feature. As the spillover reduces with increasing frequency, the zenith temperature behaves as we expect it to. The conclusions from this are that the ground-sky model at the point of intersection between the sky and the ground is unrealistic, and that the NextBASS sidelobes at low frequencies are severely problematic to the system

temperature, without sufficient ground-shielding of the optics. An improved model would consider the stability provided by the baffle, and to set the baffle temperature to a fixed 300 K, with only the central beam looking at the sky.

We can see from this sky-ground model how the optical shielding would provide a consistent system temperature at different scanning elevations. Following on from the discussion in Section 3.4.5, we can estimate the system temperature for the shielded system. If we approximate the ground temperature to 280 K, the ambient temperature to 300 K and the sky temperature to 5 K, we can demonstrate the difference in system temperature for NextBASS with and without the optical shielding. For a beam efficiency of 94%, an unshielded NextBASS has a system temperature equivalent to $94\% \times 5\text{K} + 6\% \times 280\text{K} = 16.4\text{K}$. In the case of the shielded optics, we know from Section 3.4.5 that 4% of the power was sent to the optical shield. Therefore we can calculate the system temperature to be $94\% \times 5\text{K} + 4\% \times 300\text{K} + 2\% \times 280\text{K} = 22.4\text{K}$. We can see that there has been an increase in system temperature for the shielded optics, however, since the ground temperature is susceptible to change with frequency, the time of observation and the elevation, this slight increase is a fair trade-off for a consistent system temperature for all observations.

Finally, in order to compare the true performance of both instruments, I used the antenna temperature as a replacement for the AATM sky gain temperature in the receiver model. As this propagates through the receiver chain, the overall system temperature and hence the instrument noise will be higher than the original model. Figure 6.18 shows the system temperatures and polarisation sensitivities for both X-BASS and NextBASS, with and without the spillover modelling. The resulting polarisation sensitivity has been generated using the methods described in Section 6.4, the polarisation sensitivity is scaled to suit a 3° beam.

In this Figure, we can see that for X-BASS, the increase in antenna temperature is fairly constant at $\approx 6\text{K}$ across the band. For NextBASS, there is an increase in system temperature at the start of the band, the result is a decrease in sensitivity to less than that of X-BASS, corresponding to a drop in sensitivity by approximately $2\mu\text{Kdeg}$. This is because the NextBASS optics are susceptible to larger sidelobes than the C-BASS south antenna, and this is worse at lower frequencies. In Table 6.3, the beam efficiency between the two first nulls indicates that there is a 2% difference between 7.0 GHz data point of X-BASS and NextBASS. For band-2 of NextBASS, there is a sudden jump in temperature due to the change in channel and the spillover for the high frequency beam at 15.0 GHz being greater due to the smaller aperture size. This starts to improve until 21 GHz, where the 22 GHz water line causes an

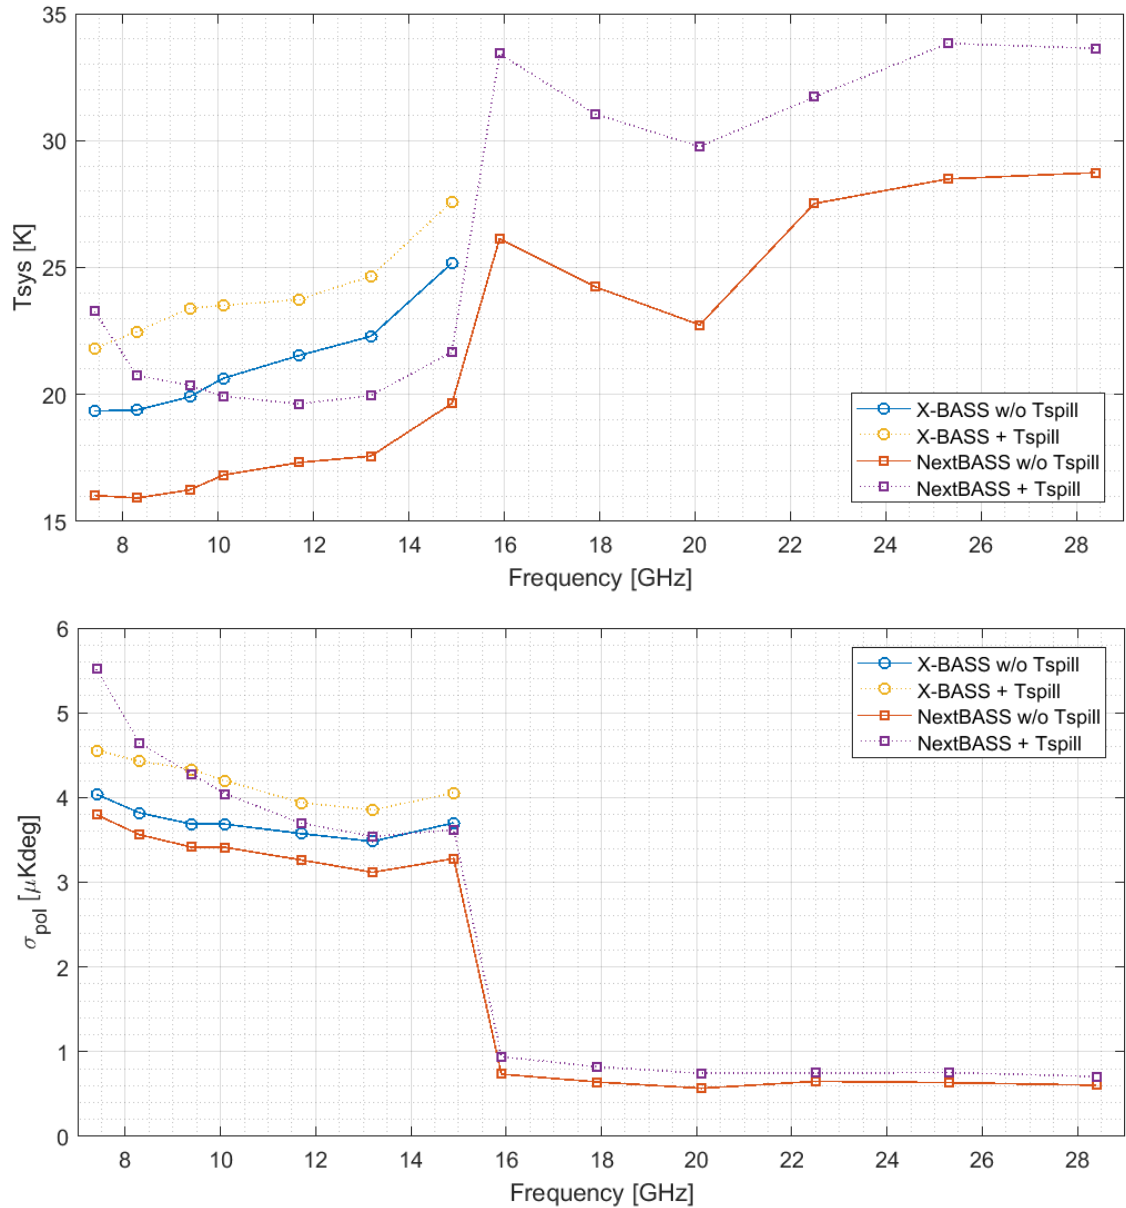


Figure 6.18: A comparison of the system temperature and polarisation sensitivity of both X-BASS and NextBASS, with and without spillover contributions from the main beam. We can see that the addition of the antenna temperature causes an increase of system temperature at 7.0 GHz for the NextBASS optics of ~ 4 K, corresponding to an loss in sensitivity of $\sim 2\mu\text{Kdeg}$. The system temperature increase is more severe at higher frequencies, however since it affects the overall sensitivity as a percentage, for the sensitive band-2 of NextBASS, the sensitivity is less affected. These results indicate the important of ground-shielding and spillover reduction when designing the optics.

increase in the antenna temperature. As the spillover becomes less at the higher frequencies, the antenna temperature improves.

Figure 6.18 shows that there is a decrease in sensitivity for X-BASS of $\approx 0.5\mu\text{Kdeg}$. However, for the lower channel of NextBASS there is a decrease in sensitivity of $\approx 2\mu\text{Kdeg}$ at 7.0 GHz. This model has assumed that the cross-Dragone optics will not be shielded from ground spillover, therefore the ground-sky function will give an overestimate on the decrease in performance. For NextBASS, it would be better to use the shielded beam pattern and apply a step function such that the main beam sees the sky and the spillover lobes only see the blackbody temperature of the absorbing material surrounding the optics.

6.4 Calculating Survey Sensitivity

In this section, I describe the equations that have been used to calculate the raw sensitivities of X-BASS and NextBASS. These calculations depend on the total observing time, the angular size of the beam, the total number of years observing and the fraction of the sky visible at the site that the instrument is located. In this section, I also describe how the total number of feeds for the NextBASS higher frequency channel was obtained by using the atmospheric data for different sites to match NextBASS to an equivalent sensitivity to synchrotron of $1\mu\text{Karcmin}$ at a frequency of 100 GHz.

6.4.1 Survey Parameters

Before calculating the sensitivity of an experiment, the number of years observing along with the fractional area of the sky need to be calculated. The longer the total length of observing time, the more sensitive the instrument, as the total integration time on the sky has increased. Realistically, an experiment cannot observe continuously throughout the day or through the year, as events such as RFI, or the sun passing through the maps can cause loss of data. To fairly represent these events, I have assumed that the observations will take place for night-time only and that only 80% of the data is usable. This amounts to 40% of the total observing time.

Although ideally, NextBASS would be a full-sky survey, in the generation of the sensitivities, I have used the location to estimate the total sky fraction observable at each latitude. The limit at which the telescope can observe depends on its latitude, given by $\Delta = \pm| - 90^\circ + 2|\phi||$ where ϕ is the latitude. The sky fraction is then given as the ratio between the solid angle that is observable at this latitude and 4π (Jew (2017)).

$$f_{sky} = \frac{1}{4\pi} \int_0^{2\pi} \int_{\Delta}^{2(\Delta+\phi)} \sin\theta d\theta d\phi \quad (6.13)$$

By matching the elevation to the latitude, all declinations below the pole will be observed. Table 6.4 shows the latitude for each of the potential NextBASS sites and the corresponding sky fraction. The Atacama site is capable of observing up to 78% of the sky, whereas Klerefontein can only view $\approx 66\%$.

Location	Atacama	Klerefontein	Tenerife
Latitude	24.63	30.97	28.30
f_{sky}	78.09%	66.38 %	71.55%

Table 6.4: The Latitudes of the three locations considered for NextBASS, and their corresponding visible sky fractions.

6.4.2 Beam width estimation

The beam width of each frequency can be extracted from the corresponding beam pattern generated in GRASP. By finding the angle at which the power has decreased by 3 dB, twice this value gives the full-width half-maximum (θ_{fwhm}) of the beam.

To correctly get estimate the number of beams that are on the sky, one must also take into account the overlap of the beams required to correctly sample the sky. For a Gaussian beam, there must be an overlap of at least one half of the full width half maximum of the beam. This means that to get good sky coverage, twice as many beams must be used, making the effective solid angle on the sky half as large. The solid angle of a Gaussian beam can be expressed as:

$$\Omega_{beam} = \frac{1}{2} \frac{\pi \theta_{fwhm}^2}{4 \ln 2}. \quad (6.14)$$

where the factor of $\frac{1}{2}$ comes from the doubling up of beams on the sky. The FWHM has influences the raw sensitivity, such that a narrower beam gives a higher sensitivity. However, since we are scaling each beam to a given pixel size in order to estimate a baseline sensitivity, (1° in intensity, and 3° in polarisation), θ_{fwhm} cancels out, so the beam width does not affect the sensitivities used in the pixel simulations.

6.4.3 Calculating Sensitivity

The number of pixels on the sky is required to find the effective integration time needed for each pixel. This scales with the angular beam size on the sky. The total

observing time t_{obs} and the sky fraction, f_{sky} , as calculated in the survey parameters sub-section are then scaled with frequency to obtain an integration time that matches the pixel size on the sky, for each frequency in the band. N_{pixels} will be much greater for the higher frequencies, resulting in shorter integration times. The longest integration time required for a 5 year NextBASS survey, would occur at 7.4 GHz, with an integration time of 247s, or ≈ 4 minutes of integration on each pixel. The number of pixels on the sky is given by:

$$N_{pixels} = \frac{4\pi f_{sky}}{\Omega_{eff}} \times \left(\frac{180}{\pi}\right)^2. \quad (6.15)$$

Therefore, the total integration time for each pixel is given by:

$$t_{int} = \frac{t_{obs}}{N_{pixels}}. \quad (6.16)$$

A good indicator of the raw sensitivity of the instrument without considering the survey parameters is the Noise Equivalent Temperature (NET). This is a measure of a detectors sensitivity to thermal radiation, measured in $\text{K}\sqrt{\text{Hz}^{-1}}$. The system temperature, T_{sys} is generated in the receiver chain and is dependent on frequency. This also scales with bandwidth. NET is defined as:

$$NET = \frac{T_{sys}}{\sqrt{\Delta\nu}}. \quad (6.17)$$

Using the NET, it is possible to then find the sensitivity to thermal noise for each individual channel. By introducing more feeds looking at that one pixel on the sky ie. increasing N_f , the sensitivity can improve by a factor of $1/\sqrt{N_f}$. This gives the raw sensitivity per beam, as we have scaled the integration time to suitably match each beam.

$$\sigma_{pp} = \frac{NET}{\sqrt{t_{int}N_f}} \quad (6.18)$$

When comparing the raw sensitivity to other experiments in order to perform the Bayesian analysis as carried out in Chapter 2, it is necessary to scale the beam to the pixel size used in the simulations. The scaling can be done linearly on the FWHM to some reference FWHM.

$$\sigma_s = \sigma_{pp} \frac{\theta_{fwhm}}{\theta_{ref}} \quad (6.19)$$

In the case of generating a sensitivity for the single pixel simulations, the intensity and polarisation reference FWHM are $\theta_{ref}^I = 1.0^\circ$ and $\theta_{ref}^P = 3.0^\circ$, respectively.

6.4.4 Number of Feeds

In order to demonstrate the influence of the 20 GHz water line at different locations, along with how the other survey parameters such as location and observing time affect the sensitivity of the instrument, I generated a script to compare the minimum number of feeds required to achieve an arbitrary sensitivity. For a given spectral index of synchrotron, assuming there are no spatial or spectral variations, it is possible to scale the instrument noise to a reference frequency in the synchrotron spectrum.

$$\sigma_{synch}^P = \sigma_s^P \left(\frac{\nu}{\nu_{ref}} \right)^\beta \quad (6.20)$$

The reference frequency I have chosen is $\nu_{ref} = 100$ GHz. I set the limit of the minimum number to occur when the sensitivity to synchrotron met $1\mu\text{Karcmin}$ at 100 GHz to match the sensitivity of *Planck* at the 100 GHz channel ([Planck Collaboration et al. \(2016b\)](#)). This value is a good representation of the sensitivity future CMB missions are aiming to achieve, PICO for example expects to obtain a sensitivity of $2.9\mu\text{Karcmin}$ at frequency of 107 GHz. I performed this calculation for the three potential sites outlined in the subsection on Location. For each location I was careful to update the sky fraction to match the corresponding latitude. I performed each calculation for total survey times 1 year and 5 years.

Figure 6.19 tells us that as we increase the frequency response, we require an exponential growth in the number in feeds to maintain the same sensitivity. As these calculations include atmospheric data, we can see that the presence of the water line affects the amount of feeds required to improve sensitivity at frequencies between 20 and 23 GHz. This plot also shows us that if we observe for a total of 5 years, location is no longer a limiting factor for frequencies below 15 GHz. This plot demonstrates that the single feed proposal for X-BASS at the Klerefontein site can match this sensitivity at the lower part of the band. The plot also indicates that at 30 GHz, at least 33 feeds are required in to meet this sensitivity. Therefore, this is the number of feeds selected for the NextBASS higher frequency channel.




Figure 6.19: The number of feeds required to achieve an equivalent sensitivity to synchrotron in intensity of $1\mu\text{Karcmin}$ at a reference frequency of 100 GHz, based on the sensitivity of Planck at this frequency (Planck Collaboration et al. (2016b)). The calculations include data from the atmospheric model using location specific attenuations obtained using the AATM software. At frequencies below 15, only a single feed is required to achieve this sensitivity, therefore the Klerefontein site is a good location for the X-BASS receiver. As the frequency increases, the required number of feeds increases exponentially. Factoring in the atmospheric effects also means that at 22 GHz, more feeds are required to compensate for the water line. At 30 GHz, ~ 33 feeds are required, hence this is the minimum number of feeds chosen for band-2 of NextBASS.

6.5 X-BASS and NextBASS Sensitivities

Figure 6.19 in the previous section allows us to come to the conclusion that for a full-sky, ground-based experiment it will be difficult to match the sensitivity of future space-based missions at 30.0 GHz. Although, this should not be considered as a downfall. As demonstrated in chapter 2, NextBASS offers a considerable improvement on foreground parameter constraints through its broad-band response and its high sensitivity at lower frequencies., It will measure the bright foregrounds to obtain a higher accuracy model and does not make measurement of the B-mode polarisation. To end this section on sensitivity calculation, I will now specify the finalised parameters used to obtain the X-BASS and NextBASS sensitivities used in the single pixel simulations in Chapter 2.

In Figure 6.20 I have generated the E and expected B -mode spectra along with a thermal dust component and three synchrotron spectral components with spectral indices $\beta_s = -2.8$, $\beta_s = -3.0$ and $\beta_s = -3.2$. We can see that foregrounds are always brighter than the B -mode CMB spectrum at $r = 10^{-3}$. This demonstrates the importance sampling the foreground spectra across all frequencies where the CMB B -mode signal is present.




Figure 6.20: The spectral components in the sky along with existing and future CMB B -mode satellite missions along with present foreground experiments C-BASS and QUIJOTE. The X-BASS and NextBASS data points show an improvement in sensitivity to QUIJOTE along with essential additional frequency resolution. The blue line gives the rms E -mode at $\ell = 100$, the orange line is the predicted B -mode frequency spectrum at $r = 10^{-3}$ for $\ell = 100$, the three red lines correspond to varying spectral indices of synchrotron $\beta_s = -3.2, -3.0, -2.8$ and the yellow line represents thermal dust. The dotted line is the extrapolated C-BASS sensitivity scaled across all frequencies. The QUIJOTE sensitivities come from Table 1 in [Rubiño Marín et al. \(2010\)](#), where the sensitivity per pixel has been scaled to a 3 degree pixel. C-BASS: [Jones et al. \(2018\)](#), QUIJOTE: [Rubiño Marín et al. \(2010\)](#), *Planck*: The *Planck* Collaboration ([Planck Collaboration et al. \(2016b\)](#)), *LiteBIRD*: Table 2 of [Remazeilles and Chluba \(2018\)](#), PICO: [PICO Collaboration et al. \(2017\)](#).

By extrapolating the C-BASS polarisation sensitivity across frequency, it is possible to see how future experiments compare, and set a baseline for NextBASS. The lower channels of PICO join with this line well. For frequencies 7 – 15 GHz, both X-BASS and NextBASS match C-BASS in sensitivity, offering improvement on the current sensitivity proposed by QUIJOTE, when scaled to 3°. NextBASS² with 33 feeds at a total observing time of 5 years meets this sensitivity for the first three frequencies of the 15 – 30 GHz channel. Improvements to sensitivity can be made by increasing the total number of feeds, the total integration time or by supplementing full-sky survey data with deeper integrations, obtaining a higher spatial resolution.

Experiment	Years	Site	Channels
X-BASS	5.0	Klerefontein	1 feed \times [7.4, 8.3, 9.4, 10.1, 11.7, 13.2, 14.9] GHz
NextBASS	5.0	Atacama	1 feed \times [7.4, 8.3, 9.4, 10.1, 11.7, 13.2, 14.9] GHz 33 feeds \times [15.9, 17.9, 20.1, 22.5, 25.3, 28.4] GHz

Table 6.5: X-BASS and NextBASS surveys and band separation. The NextBASS² feed array design is given in Chapter 4, Figure 4.24. The bandwidths are based on 20% bandwidths, equally spaced in frequency space and hence generating the central frequencies given in this Table.

ν [GHz]	$\Delta\nu$ [GHz]	θ_{fwhm}°	NET $\mu\text{K}/\sqrt{\text{Hz}}$	T_{SYS} [K]	$\sigma_{pp}^{\mu\text{K}}$	$\sigma_I^{\mu\text{Kdeg}}$	$\sigma_P^{\mu\text{K3deg}}$	$\sigma_{synch}^{\mu\text{Karcmin}}$
7.4	1.48	0.46	503	19.35	23.52	12.11	4.04	0.08
8.3	1.66	0.42	476	19.39	24.96	11.46	3.82	0.10
9.4	1.88	0.38	459	19.91	27.28	11.06	3.69	0.15
10.1	2.02	0.38	459	20.64	29.31	11.06	3.69	0.18
11.7	2.34	0.34	445	21.54	32.92	10.72	3.57	0.28
13.2	2.64	0.30	434	22.29	36.19	10.45	3.48	0.3
14.9	2.98	0.26	461	25.17	43.42	11.11	3.70	0.61

Table 6.6: Table of X-BASS sensitivities derived with the methods described in this chapter. The intensity (σ_I) and polarisation (σ_P) noise values were used in the single pixel forecasting simulations. The FWHM has been calculated from the 3 dB beam width of the GRASP simulations. The atmospheric contribution comes from the AATM model of conditions based on those at the C-BASS site. These sensitivities exclude the additional antenna temperature discussed in Section 6.3.5 since the model required further development.

ν [GHz]	$\Delta\nu$ [GHz]	θ_{fwhm}°	NET $\mu\text{K}/\sqrt{\text{Hz}}$	T_{SYS} [K]	$\sigma_{pp}^{\mu\text{K}}$	$\sigma_I^{\mu\text{Kdeg}}$	$\sigma_P^{\mu\text{K3deg}}$	$\sigma_{\text{synch}}^{\mu\text{Karcmin}}$
7.4	1.48	0.54	416	16.02	22.11	11.39	3.80	0.07
8.3	1.66	0.50	391	15.92	23.27	10.68	3.56	0.09
9.4	1.88	0.46	375	16.24	25.27	10.24	3.41	0.13
10.1	2.02	0.46	374	16.82	27.12	10.23	3.41	0.17
11.7	2.34	0.42	358	17.32	30.06	9.79	3.26	0.25
13.2	2.64	0.38	342	17.57	32.39	9.35	3.11	0.35
14.9	2.98	0.38	360	19.64	38.47	9.84	3.27	0.54
15.9	3.18	0.26	463	26.14	9.21	2.21	0.73	0.15
17.9	3.58	0.22	405	24.25	9.06	1.93	0.64	0.18
20.1	4.02	0.22	359	22.74	9.00	1.71	0.57	0.24
22.5	4.5	0.18	410	27.52	11.53	1.95	0.65	0.38
25.3	5.06	0.18	400	28.49	12.66	1.91	0.63	0.54
28.4	5.68	0.14	381	28.74	13.53	1.81	0.60	0.73

Table 6.7: Table of NextBASS sensitivities derived with the methods described in this chapter. The intensity (σ_I) and polarisation (σ_P) noise values were used in the single pixel forecasting simulations. The FWHM has been calculated from the 3 dB beam width of the GRASP simulations. The atmospheric contribution comes from the AATM model of conditions similar to an ideal observing day on the Chajnantor Plateau, Chile. The antenna temperature estimation has been excluded from these calculations as the ideal NextBASS instrument will be placed in a ground-shield and system temperature increases due to ground-spillover will be minimised.

6.6 Summary

This chapter has addressed a large amount of variables that can contribute to the system temperature of the receiver and hence the overall sensitivity of X-BASS and NextBASS. The final sensitivities used in the simulation have taken into consideration the atmospheric conditions of the two locations, but have not included the antenna spillover temperature in the overall system temperature calculation. This decision was made because further work for the NextBASS optics is to develop a better model of the ground-shield and use an improved ground-to-sky temperature gradient for the shielded optics. For X-BASS, there was not a considerable loss in sensitivity that could not be overcome by using slightly longer integration times. The main improvement on the current generation of low frequency foreground instruments is

that X-BASS and NextBASS have much simpler receivers to implement, meaning considerably lower system temperatures.

NextBASS also gains more sensitivity from the feed array for band-2, which I have shown is necessary to achieve similar sensitivity to C-BASS and the future CMB satellite instruments. As shown by the pixel simulations in Chapter 2, these are the sensitivities that are required to achieve meaningful parameter constraints that can help remove bias on the CMB B -mode amplitude.

The work in this chapter is open to continued development such as obtaining more location-specific atmospheric and ground-spillover models, modelling the full feed array with a suitable cryostat at an attainable temperature, investigating noise sources from potential leakage and developing the survey parameters to combine information from two potential scans of different elevations. The survey strategies are important and will tie into the map making process when removing residual scanning features. Future modelling could include determining sensitivities for different scanning strategies to create a simulated sky map.

Chapter 7

A Summary: The Next Generation of Galactic Foreground Instruments

“ When you know and respect your inner nature, you know where you belong. You also know where you don’t belong.”

Benjamin Hoff, *The Tao of Pooh*

This thesis has attempted to cover the breadth of investigations required to justify a new Galactic Foregrounds instrument. As a result, the summary will address the proposed X-BASS and NextBASS instruments, along with a more general approach to designing instruments for this purpose.

To begin the summary, I would like to re-address each chapter and briefly comment on the key factors, and the further work required for the development of ideas outlined in this thesis. Firstly, and most importantly, the single pixel simulations are the robust conclusion that an experiment similar to NextBASS is required to make a definitive constraint on the synchrotron curvature, and to any additional polarised, low frequency foregrounds. I have demonstrated that a new instrument can constrain foregrounds further than any other current instruments that generate large-area sky maps. I’d like to stress the importance of the polarised AME simulations. Statements from the literature I would like to remark on, include that if future observations find the AME polarised to greater than 0.01% at frequencies greater than 10 GHz, this will provide sufficient evidence that the AME is not entirely rotational emission, and that starlight is a contributing factor to the background polarisation ([Draine and Hensley \(2016\)](#)). Furthermore, if a detection of 1% AME polarisation is made, then

it may be possible to distinguish between spinning dust and magnetic dust mechanisms. NextBASS can detect a potential 1% polarised AME component in pixels with strong foreground emission, and determine whether or not the polarised sky model used to fit the data, is incorrect. This has been demonstrated by the obvious bias that the NextBASS frequency data introduced on foreground parameters and the CMB amplitude when the sky model is incorrect. If deep observations were considered then NextBASS could reign in on different regions of the ISM and perhaps get down to sensitivity levels capable of detecting much lower AME polarisation fractions. Therefore, we can conclude that the frequencies and sensitivities of NextBASS are essential to identifying the physical mechanisms behind the anomalous microwave emission.

NextBASS has shown to improve constraints and remove bias on the CMB B -mode amplitude. It has proven to do this for a variety of pixels with different strengths of foregrounds, and for any level of foreground modelling complexity. The error bars on the CMB amplitude for both the intensity and polarisation frequency spectra show a considerable improvement with the addition of both the X-BASS and NextBASS data points. Therefore, we can conclude that frequencies between 7 – 30 GHz, are required to better constrain the CMB amplitude. We can also conclude that X-BASS is capable of improving constraints on the CMB amplitude in pixels that have strong synchrotron emission, however in dustier regions with AME, NextBASS will achieve better constraints than X-BASS. This justifies the 15 – 30 GHz channel of NextBASS. Therefore, we can conclude that to constrain synchrotron, only X-BASS is necessary, however, to constrain and de-tangle other foreground components, the full NextBASS instrument is essential.

I have also shown that more data is needed when dealing with complex sky models with multiple free-parameters, such as in intensity. Future work includes trying a more efficient sampling methods and looking into the use of reference priors. It would also be beneficial to perform full-sky simulations and generate simulated data that includes spatially varying spectral indices. These simulations have shown that the dust parameters are still poorly constrained. In the paper ‘Component Separations on future CMB instruments’ by [Remazeilles \(2018\)](#), it is shown that in order to effectively constrain the dust parameters, frequencies up to 799 GHz (as provided by the PICO experiment). A stronger case may be made if these simulations included contributions from PICO at higher frequencies. Finally, no work has been done on investigating the effects of NextBASS and X-BASS on deep sky surveys to compliment experiments such as CMB-S4. Deeper observations mean much higher sensitivities and a chance

to see how NextBASS and X-BASS could be used in efforts to obtain an improved CMB B -mode power spectrum.

I have compared optical configurations for NextBASS and used an analysis of the focal plane for the two potential antenna candidates to confirm the possibility of making a feed array between 15 – 30 GHz. I have demonstrated ways of controlling ground-spillover for the NextBASS optics, by using a ground-shield and introducing an absorbing baffle. I have also shown that the surface distortions present on the C-BASS South main reflector are small enough that it can be used up to 15 GHz, without distorting the beam. I have designed hyperbolic feed horns to match the NextBASS optics and I have created an optimised feed for X-BASS to match the C-BASS optics and obtain a low cross-polarisation. I have performed full optical analyses for X-BASS and NextBASS indicating that they are both capable of obtaining high beam efficiencies and low cross-polarisation. The final combined feed and optical simulations have shown that the feed horns achieve similar edge tapers to those used in the design stage.

Future work includes calculating the intrinsic cross-polarisation (IXR) as an indication of the polarisation performance of each telescope ([Carozzi and Woan \(2011\)](#)). An important piece of future work that must be done in order to proceed with the NextBASS experiment, is to develop an optical model of the full feed array in GRASP, and obtain beam patterns for each feed. It would also be important to optimise the NextBASS feed horns to achieve lower cross-polarisation and have shorter lengths so that the phase centres are less variable. The main systematic that caused problems in the C-BASS data was the noise contribution from ground-spillover. I would recommend to investigate ways of reducing X-BASS ground-spillover, such as building a ground shield around C-BASS.

In this thesis I have explored the options for both the X-BASS and NextBASS receivers, and made a decision on the architecture based on the suitability for a broadband receiver. I have shown that many components are readily available and have been used in similar experiments, meaning that the architectures have a high technology readiness level. I have designed the appropriate microwave filters and shown that the proto-type anti-aliasing filter is well-matched to the design cut-off frequency. However, the tests carried out in this performance indicate that further development of the band-pass filters is required to ensure their reliability. Such developments include using alternative manufacturing techniques or using a different style of filter.

I have demonstrated that the receiver chain will not cause degradation in sensitivity. I have suggested a potential FPGA architecture to compliment the choice

of DBE. Future work will require testing a full receiver chain to compare with the receiver model and investigate the effects of impedance mismatching between the components. An important investigation for the NextBASS receiver would be the quantisation noise loss in conjunction with the low bit ADC and the CW architecture. It may be such that having an ADC with a small dynamic range for a broadband architecture is not ideal. The next stage is to build a full receiver proto-type and make measurements of the $1/f$ noise and investigate polarisation leakage.

The receiver design of NextBASS is heavily reliant on the use of the FPGA to provide active level control of the continuous wave signal and to perform gain stabilisation in real-time. Circular polarisation will also be generated in the FPGA; an investigation needs to be done into whether this stage will be carried out before, or after channelisation, as the Stokes parameters are typically obtained in post-processing. Further more, the development of an FPGA to simultaneously perform these functions over a 2:1 bandwidth will require a lot of technology development into FPGA programming. As a result, new DSP algorithms could be developed to optimally achieve each of these steps at fast processing speeds. This DSP technology is applicable to all future Astronomy projects that acquire large sets, as real-time data processing reduces the problem of ‘Big-Data’, a commonly discussed issue arising with the SKA ([Adam \(2016\)](#)).

In Chapter 6, I estimated the sensitivities of X-BASS and NextBASS from a fair representation of the receiver and atmospheric conditions at each potential site. This took a lot of time and consideration into the frequency dependent effects from the atmosphere, and the resulting propagation through a realistic receiver chain. The trial runs allowed me to determine the minimum number of feeds required at 30 GHz for a 5 year survey at a site with atmospheric conditions similar to those of the Llano de Chajnantor observatory in the Atacama, Chile. This was in order to match an equivalent sensitivity to synchrotron at 100 GHz of $1\mu\text{K arcmin}$, which is the sensitivity of the *Planck* satellite at this frequency ([Planck Collaboration et al. \(2016b\)](#)). This number gave $N_f = 33$, which is a realistic number of corrugated feed horns to put in the focal plane of the 6m, cross-Dragone NextBASS optics. I have also shown that these attainable sensitivities are comparable to future CMB missions such as *LiteBIRD* and *PICO* (See [Figure 6.20](#)). There is a lot more potential for this chapter to develop into a full survey proposal. For example, I have based these on the assumption that the surveys will scan at a constant elevation of 45° . It would be good to incorporate this scanning strategy and the corresponding sensitivities into generating full-sky simulated data by similar methods outlined in Chapter 2.

Furthermore, I also demonstrated that further work needs to go into the development of an accurate ground-spillover model that can use the properties of the telescope beam to obtain an estimate of the total antenna temperature.

The NextBASS project could coincide with launch dates as early as 2030, and due to the availability of infrastructure and technology, X-BASS could be integrated on the C-BASS South antenna, and taking data by 2025. By treating X-BASS as a ‘path-finder’, it could pave the way as an example of what can be done for the CMB foreground community. X-BASS would compliment data obtained by S-PASS (Carretti et al. (2019)) and C-BASS (Jones et al. (2018)) as these are all of similar resolution. The result would be a full coverage of frequency space where synchrotron is the dominant foreground component. This would obtain a measurement of the synchrotron spectral index and curvature, across the full sky to an incredible accuracy. This data would be invaluable to those researching the Galaxy’s magnetic field.

There is not currently a ‘NextBASS’ collaboration. The concept has been presented at various CMB conferences and has always received an interested and motivated response. Perhaps it would be better to re-brand the experiment and open it to an international collaboration. The name ‘NextBASS’ implies a full-sky survey, however the work in this thesis has only considered a southern-sky candidate. A site of equivalent atmospheric conditions for the Northern Hemisphere would be La Palma, in the Canary Islands. Perhaps European efforts could mean this work would be developed into a ‘European Low-Frequency Survey’, which has the much more enjoyable and memorable acronym, ‘*ELFS*’. The CMB community agrees that more frequency data is needed to improve the likelihood of detecting a tensor-to-scalar ratio as low as $r = 10^{-3}$.

I conclude this thesis with the statement that it should be imperative that an instrument with continuous frequency coverage from 7 to 30 GHz, at a sensitivity paralleled to C-BASS and future CMB satellite missions, will play a vital role in the planning of future efforts made by the CMB community.

Appendix A

Markhov Chain Monte Carlo

It can be very difficult to analytically find the parameter values that maximise multi-parameter posterior distributions and to determine the shape of the posteriors near the maximum. To carry out an unbiased exploration of the aforementioned foreground parameters in Section 2.2, Markhov Chain Monte Carlo (MCMC) methods offer a way of assessing the highly correlated and multidimensional parameter space. If we consider the probability distribution of a and b , but we are uninterested in b , then the marginalized distribution of a $p(a)$ can be found using:

$$p(a) = \int_{-\infty}^{\infty} p(a, b) db = \int_{-\infty}^{\infty} p(a|b)p(b) db \quad (\text{A.1})$$

By sampling the posterior distribution using Markhov chains, Monte Carlo methods estimate a probability distribution by taking random samples from a chain and testing whether one step in the chain increases the posterior probability than another step. If the $(n + 1)^{th}$ step is determined only by the n^{th} step's state, the the chain is Markovian. The cycle begins by choosing a starting value for the parameters and then the sampling begins. The initial chains can be ignored (burned) as they are not likely to be representative of the posterior distribution.

There are many MCMC algorithms, which I will not go into too much detail. The simulations outlined in this thesis apply the Metropolis-Hastings (M-H) algorithm, which works by taking a random walk through parameter space. The M-H algorithm, developed by (Metropolis et al., 1953) and (Hastings, 1970), is a commonly used MCMC step method. It works such that at the n^{th} step of the chain, the $(n + 1)^{th}$ is found by deciding on the next step in the random walk by defining a proposal distribution, the acceptance ratio is then calculated and finally, if the acceptance ratio is greater than one, the step is accepted, if this ratio is less than one, the

proposed step is accepted with a probability of acceptance ratio. The chain then repeats. We can write the algorithm out in steps:

1. Draw a step from the proposal distribution: $g(\theta_{n+1}|\theta_n)$
2. Calculate the acceptance ratio: $\alpha = \frac{p(\theta_{n+1}|\mathbf{d})}{p(\theta_n|\mathbf{d})}$
3. If $\alpha \geq 1$ then the proposed step is accepted. If $\alpha < 1$ the proposed step is accepted with a probability α
4. Repeat.

The main disadvantage of this method it is lengthy in computational time, especially when increasing the number of steps. An extension of the M-H algorithm is the Gibbs Sampling algorithm, this particular version has been used in the COMMANDER-RULER code which generated the component separated maps for the Planck Collaboration ([Planck Collaboration et al. \(2013a\)](#)).

Appendix B

Sky Model

Starting at the zenith, the airmass can be written as a function of theta, such that the zenith is at 90° to the ground. This model holds for a slab-style atmospheric model however it is not realistic at angles approaching 0° (see Figure 6.16). The simple airmass is by the equation:

$$a_0 = \csc(\theta) \quad (\text{B.1})$$

where θ is the angle deviating from the zenith. This is equal to $|90 - \theta_{el}|$, where θ_{el} is the elevation angle. If we compensate for curvature, we can write the corrected airmass as:

$$\tau(\theta) = \frac{\tau(0)}{\sin \theta} - \frac{H}{r_e \sin \theta} \left(\frac{1}{\sin^2} - 1 \right) \tau(0). \quad (\text{B.2})$$

where H describes the effective height of the atmosphere, also known as the absorption scale height. It can be given by the following equation:

$$H = -\log_{10}\left(\frac{\exp(-\tau(\nu))}{r_e}\right). \quad (\text{B.3})$$

The corrected airmass, a_c can now be given by the relationship:

$$a_c - a_0 = a_0 \left(\frac{H}{r_e} \right) \left(\frac{1}{\sin^2(\theta)} - 1 \right) \quad (\text{B.4})$$

The sky brightness becomes :

$$Tb = T_{sky} \csc(\theta) (1 - \exp(-\tau_\nu \csc \theta)) + a_c T_{sky} \exp(-\tau_\nu a_c) \quad (\text{B.5})$$

Appendix C

Ground Model

The ground component needs to include emission and scattering from the ground. We can consider the reflection coefficients at the interface of ground and the air, $\Gamma_{\parallel}(\theta_1)$ and $\Gamma_{\perp}(\theta_1)$ where θ_1 is $\pi - \theta$. These are described in SKA: Memo 95 ([Mendellín \(2007\)](#)) originally from ([Ulaby et al. \(1981\)](#)) and represent the polarised reflection coefficients.

$$\Gamma_{\parallel}(\theta_1) = \left| \frac{\cos \theta_1 - \sqrt{\epsilon_2 - \sin^2 \theta_1}}{\cos \theta_1 + \sqrt{\epsilon_2 - \sin^2 \theta_1}} \right|. \quad (\text{C.1})$$

$$\Gamma_{\perp}(\theta_1) = \left| \frac{\epsilon_2 \cos \theta_1 - \sqrt{\epsilon_2 - \sin^2 \theta_1}}{\epsilon_2 \cos \theta_1 + \sqrt{\epsilon_2 - \sin^2 \theta_1}} \right|. \quad (\text{C.2})$$

For these antenna noise calculations we can assume that the ground emission and bright sky noise sources are un-polarised. We can also assume that for dry land $\epsilon_2 \approx 3.5$ meaning that $\Gamma_{\parallel}(\theta_1)$ and $\Gamma_{\perp}(\theta_1)$ are similar. We can hence define an average reflection coefficient.

$$\tilde{\Gamma}(\theta_1) = \frac{\Gamma_{\parallel}(\theta_1) + \Gamma_{\perp}(\theta_1)}{2}. \quad (\text{C.3})$$

We can now write the full integral as:

$$T_b^{sky}(\theta, \nu, \phi) P_n(\nu, \theta, \phi) = \begin{cases} P_n(\nu, \theta, \phi | \hat{\mathbf{r}}_0) T_b^{sky}(\nu, \theta), & 0 \leq \theta < \pi/2. \\ P_n(\nu, \theta, \phi | \hat{\mathbf{r}}_0) \left[(1 - \tilde{\Gamma}(\theta_1)) T_{gnd} + \tilde{\Gamma}(\theta_1) T_b^{sky}(\theta)_1 \right], & \pi/2 \leq \theta \leq \pi. \end{cases} \quad (\text{C.4})$$

Bibliography

- K. Abazajian et al. CMB–S4 Science Book, First Edition. *arXiv eprint*, 2016. URL <https://arxiv.org/abs/1610.02743>.
- R. Adam. The SKA and Big Data. *NACI Business Symposium on Science, Technology and Innovation*, June 2016.
- P. Ade et al. A measurement of the cosmic microwave background B-mode polarization power spectrum at sub-degree scales with Polarbear. *ApJ*, 794:171, 2014. doi: 10.1088/0004-637X/794/2/171.
- Y. Ali-Haimoud et al. Search for polycyclic aromatic hydrocarbons in the Perseus molecular cloud with the Green Bank Telescope. *MNRAS*, 447(1):315–324, 2015. doi: 10.1093/mnras/stu2476.
- R. Alpher and C. Herman, R. On the evolution of the universe. *Nature*, 162:772–775, 1948.
- Analog Devices. *HMC462LP5462LP5E, GaAs pHEMT MMIC LOW NOISE AMPLIFIER, 2 to 20 GHz*, 2015. URL www.analog.com.
- Analog Devices. *GaAs MMIC Fundamental Mixer, 6 GHz to 26 GHz, HMC773A Data Sheet*, 2015. URL www.analog.com.
- Analog Devices. *2.25 to 18 GHz, MMIC, Double Balanced Downconverter*, 2018. URL www.analog.com.
- C. Angel, R. Querel, and F. Kerber. Precipitable Water Vapor: Considerations on the water vapor scale height, dry bias of the radiosonde humidity sensors, and spatial and temporal variability of the humidity field. *arXiv eprint*, 2011. URL <https://arxiv.org/abs/1103.3025>.
- ANSYS Inc. Designer version 8.0, 2012. URL <https://www.ansys.com/en-gb/products/electronics/ansys-electronics-desktop>.

- AtlanTecRF. Miniature Synthesisers ASY Series, 2019a. URL http://www.atlantecrf.com/products/oscillators/miniature_synthesisers_asy_series.htm.
- AtlanTecRF. 90 Degree Hybrid 3dB Couplers, 2019b. URL http://www.atlantecrf.com/products/passive_components/couplers/90_deg_hybrid_couplers.htm.
- K. Aylor et al. A Comparison of Cosmological Parameters Determined from CMB Temperature Power Spectra from the South Pole Telescope and the Planck Satellite. *ApJ*, 850:101, 2017. doi: 10.3847/1538-4357/aa947b.
- C. Baccario. A visualization of global weather conditions forecast by supercomputers updated every three hours, 2018. URL <https://earth.nullschool.net/>.
- C. A. Balanis. *Antenna Theory*, chapter 2. Wiley, third edition, 2005.
- C. A Balanis. *Antenna Theory: Analysis and Design*. Wiley, third edition, 2009.
- D. Baumann et al. CMBPol Mission Concept Study: Probing Inflation with CMB Polarization. *AIP Conference Proceedings*, 1141:10–120, 2008. doi: 10.1063/1.3160885.
- C. Bennett et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Final maps and results. *ApJ*, 208:20, 2013. doi: 10.1088/0067-0049/208/2/20.
- L. Bennett, C. et al. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Preliminary Maps and Basic Results. *ApJ*, 148:1–27, 2003. doi: 10.1086/377253.
- B. Bhattacharyya et al. A Long-term study of three rotating radio transients. *MNRAS*, (3):4090–4103, 2018. doi: 10.1093/mnras/sty923.
- BICEP Collaboration et al. BICEP2. II. Experiment and three-year data set. *ApJ*, 792:62, 2014. doi: 10.1088/0004-637X/792/1/62.
- BICEP2 Collaboration. Detection of *B*-Mode Polarization at Degree Angular Scales by BICEP2. *Physical Review Letters*, 112(24):241101, 2014. doi: 10.1103/PhysRevLett.112.241101.

- BICEP2/Keck, Planck Collaboration et al. A Joint Analysis of BICEP2/Keck Array and Planck Data. *Physical Review Letters*, 114:101301, 2015. doi: 10.1103/PhysRevLett.114.101301.
- C. Bischoff. The QUIET Instrument. *ApJ*, 768:9, 2013. doi: 10.1088/0004-637X/768/1/9.
- Bernard Burke and F. Graham-Smith. *An Introduction to Radio Astronomy*. Cambridge Press, third edition, 2010.
- D. Carozzi, T. and G. Woan. A Fundamental Figure of Merit for Radio Polarimeters. *IEEE*, 59(6):2058–2065, 2011. doi: 10.1109/TAP.2011.2123862.
- E. Carretti et al. S-band Polarization All Sky Survey (S-PASS): survey description and maps. *arXiv eprints*, 2019. URL <https://arxiv.org/abs/1903.09420>.
- J. Chluba, C. Hill, J., and H. Abitbol, M. Evolution of CMB spectral distortion anisotropies and tests of primordial non-Gaussianity. *MNRAS*, 466(2):2390–2401, 2017. doi: 10.1093/mnras/stw3230.
- S. Choy et al. Application of GPS Atmospheric Sounding for Climate Studies in the Australian Region. *Proc. FIG*, 7574:1–13, 2015.
- J. B. Clarricoats, P. and D. Olver, A. *Corrugated horns for microwave antennas*. IEEE, 18 edition, 1984.
- COESA. U.S. Committee on Extension to the Standard Atmosphere. 1976. URL https://ccmc.gsfc.nasa.gov/modelweb/atmos/us_standard.html.
- Collaboration for Astronomy Signal Processing and Electronics Research, 2009. URL <https://casper.ssl.berkeley.edu/wiki/Projects>.
- Charles Copley. *D.Phil Thesis: The C-Band All Sky Survey*. University of Oxford, 2014.
- CORE Collaboration et al. Exploring cosmic origins with core: Inflation. *Journal of Cosmology and Astroparticle Physics*, 04:016, 2016. doi: 10.1088/1475-7516/2018/04/016.
- R. I. Curtis, E., B. Nikolic, S. Richer, J., and R. Pardo, J. ALMA Memo 590: Atmospheric dispersion and the implications for phase calibration. *ALMA Memo*, 590, 2009. URL <http://www.mrao.cam.ac.uk/~eic22/ALMAmemos/memos>.

- P. de Bernardis et al. Mapping the CMB sky: The BOOMERanG experiment. *New Astronomy Reviews*, 43:289–296, 1999. doi: 10.1016/S1387-6473(99)00022-6.
- P. de Bernardis et al. First results from the BOOMERanG experiment. 555:85–94, 2000. doi: 10.1063/1.1363510.
- G. de Gasperis et al. Optimal cosmic microwave background map-making in the presence of cross-correlated noise. *A&A*, 593:A15, 2016. doi: 10.1051/0004-6361/201628143.
- D. A. Del Rizzo and M. A. B. Garstin. Monitor and control for the SKA1 CSP Mid.CBF utilizing the Stratix10 FPGA equipped with HPS. *SPIE*, 10707:107071S, 2018. doi: 10.1117/12.2309986.
- J. Delabrouille, F. M. Cardoso, J., M. Le Jeune, M. Betoule, G. Fay, and F. Guillox. A full sky, low foreground, high resolution CMB map from WMAP. *A&A*, 493(3): 835–857, 2009. doi: 10.1051/0004-6361:200810514.
- C. Dickinson et al. Studies of Anomalous Microwave Emission AME with the SKA. *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 2014. URL <https://arxiv.org/abs/1412.5054>.
- C. Dickinson et al. The State-of-Play of Anomalous Microwave Emission (AME) Research. *A&A*, 80:1–28, 2018a. doi: 10.1016/j.newar.2018.02.001.
- C. Dickinson et al. The C-Band All-Sky Survey (C-BASS): Constraining diffuse Galactic radio emission in the North Celestial Pole region. *MNRAS*, 485:2844–2860, 2018b. doi: 10.1093/mnras/stz522.
- R. D’Inverno. *Introducing Einstein’s Relativity*. Oxford University Press, first edition, 1992.
- F. Dominguez. Lectures, 2011. URL http://www.atmo.arizona.edu/students/courselinks/spring11/hwr523/HWR423_523_files/Module1b.pdf.
- B. Draine and A. Lazarian. Electric Dipole Radiation from Spinning Dust Grains. *ApJ*, 508:157–179, 1998. doi: 10.1086/306387.
- B. T. Draine and Brandon S. Hensley. Quantum Suppression of Alignment in Ultra-small Grains: Microwave Emission from Spinning Dust will be Negligibly Polarized. *ApJ*, 831:915311, 2016. doi: 10.1117/12.2056267.

- T. Draine, B. *Physics of the interstellar and intergalactic medium*. Princeton University Press, 2011.
- T. Draine, B. and B. Hensley. Magnetic Nanoparticles in the Interstellar Medium: Emission Spectrum and Polarization. *ApJ*, 765:159, 2013. doi: 10.1088/0004-637X/765/2/159.
- A. Einstein. *Sitzungsberichte der Königlich Preußischen Akad. der Wissenschaften*, 84:831–839, 1915.
- J. Ekin, W. *Experimental Techniques for Low-Temperature Measurements: Cryostat Design, Material Properties, and Superconductor Critical-Current Testing*. Oxford University Press, 2007.
- K. Eriksen, H., J. Banday, A., M. Górski, K., and B. Lilje, P. On Foreground Removal from the Wilkinson Microwave Anisotropy Probe Data. *ApJ*, 155:633–646, 2004. doi: 10.1086/422807.
- K. Eriksen, H., B. Jewell, J., C. Dickinson, J. Banday, A., K. M. Górski, and M. Lawrence. Joint Bayesian Component Separation and CMB Power Spectrum Estimation. *ApJ*, 676:10–32, 2008. doi: 10.1086/525277.
- J. Errard et al. Robust forecasts on fundamental physics from the foreground-obscured, gravitationally lensed CMB polarization. *Journal of Cosmology and Astroparticle Physics*, page 052, 2016. doi: 10.1088/1475-7516/2016/03/052.
- O.G. King et al. Astronomical Modelling Using Scattering Matrices. *RAS*, 446(2): 1252–1267, 2014. doi: 10.1093/mnras/stu2172.
- J. Fixsen, D. The Temperature of the Cosmic Microwave Background. *ApJ*, 707: 916–920, 2009. doi: 10.1088/0004-637X/707/2/916.
- A. Friedmann. *Zeitschrift für Phys.*, 10:377–386, 1922. doi: 10.1007/BF01332580.
- U. Fuskeland et al. Spatial Variations in the Spectral index of polarised Synchrotron emission in the 9 yr WMAP sky maps. *ApJ*, 790:104, 2014. doi: 10.1088/0004-637X/790/2/104.
- R. Génova-Santos et al. Measurements of the Intensity and Polarization of the Anomalous Microwave Emission in the Perseus molecular complex with QUIJOTE. *MNRAS*, 452(4):4169–4182, 2015. doi: 10.1093/mnras/stv1405.

- R. Giovanelli et al. The Optical/Infrared Astronomical Quality of High Atacama Sites. II. Infrared Characteristics. *The Astronomical Society of the Pacific*, 113: 803–813, 2001. doi: 10.1086/322136.
- M. Girolami and B. Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Royal Statistical Society*, 73:1–37, 2011.
- B. Gold et al. Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Galactic Foreground Emission. *ApJ*, 180:265–282, 2008. doi: 10.1088/0067-0049/180/2/265.
- F. Goldsmith, P. *Quasioptical Systems*. Wiley-IEEE Press, 1993.
- C. Granet and L. Graeme, J. Design of Corrugated Horns: A Primer. *IEEE*, 47, 2005.
- J. Grayson et al. BICEP3 performance overview and planned Keck Array upgrade. *Proc. SPIE*, 9914:99140S, 2016. doi: 10.1117/12.2233894.
- R. Grumitt et al. NUTS Component Separation. *MNRAS*, in prep.
- R. Gualtieri et al. SPIDER: CMB polarimetry from the edge of space. *Journal of Low Temperature Physics*, 193:1112–1121, 2017. doi: 10.1007/s10909-018-2078-x.
- Y. Han and R. Westwater, E. Analysis and Improvement of Tipping Calibration for Ground-Based Microwave Radiometers. *IEEE*, 38:1260 – 1276, 2000. doi: 10.1109/36.843018.
- S. Hanany and P. Marrone, D. Comparison of designs of off-axis Gregorian telescopes for millimeter-wave large focal-plane arrays. *OSA*, 41:4666–4670, 2002. doi: 10.1364/AO.41.004666.
- D. Hanson et al. Detection of B-mode polarization in the Cosmic Microwave Background with data from the South Pole Telescope. *Physical Review Letters*, 111: 141301, 2013. doi: 10.1103/PhysRevLett.111.141301.
- K. Hastings, W. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:97–109, 1970. doi: 10.1093/biomet/57.1.97.
- E. Hecht. *HECHT OPTICS*. Addison Wesley, third edition, 1997.

- B. Hensley. Modeling the Anomalous Microwave Emission with Spinning Nanoparticles: No PAHs Required. *ApJ*, 836:179, 2017. doi: 10.3847/1538-4357/aa5c37.
- B. Hensley and P. Bull. Mitigating Complex Dust Foregrounds in Future Cosmic Microwave Background Polarization Experiments. *ApJ*, 853(2):127, 2018. doi: 10.3847/1538-4357/aaa489.
- Hittite Microwave Corporation. *HMCAD5831LP9BE*, 2016. URL www.hittite.com.
- T. Hoang and A Lazarian. A Unified Model of Grain Alignment: Radiative Alignment of Interstellar Grains With Magnetic Inclusions. *A&A*, 831:159, 2016. doi: 10.3847/0004-637X/831/2/159.
- D. Hoffman, M. and . Gelman, A. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *arXiv eprints*, 2011. URL <https://arxiv.org/abs/1111.4246>.
- C. Holler et al. A Circularly Symmetric Antenna Design With High Polarization Purity and Low Spillover. *IEEE*, 61:117–124, 2011. doi: 10.1109/TAP.2012.2219843.
- C. M Holler. C-BASS Memo: Reduction of Sidelobes of the JPL 6.1m Antenna. private communication, 2007.
- C. M Holler. C-BASS Memo: Optical design of the 6.1m and 7.6m antennas for C-BASS. private communication, 2009.
- A. Hyvarinen. Fast and Robust Fixed-Point Algorithms for Independent Component Analysis. *IEEE*, 6:626–634, 1999.
- Y. Inoue et al. POLARBEAR-2: an instrument for CMB polarization measurements. *Proc. SPIE*, 9914:99141I, 2016. doi: 10.1117/12.2231961.
- INTEL. *INTEL STRATIX GX/SX Device Overview*, 2018. URL <https://www.intel.com/content/www/us/en/programmable/documentation/joc1442261161666.html>.
- ITU. Attenuation by atmospheric gases. *ITU*, ITU-R P.676-11, 2012. URL <https://www.itu.int/rec/R-REC-P.676/en>.
- ITU. Attenuation by atmospheric gases. *ITU*, ITU-R P.676-11, 2016. URL <https://www.itu.int/rec/R-REC-P.676/en>.

- Janis Research Company. Cryostat Research Equipment, 2019. URL <https://www.janis.com/Products/Cryocoolers/PulseTubeVsGifford-McMahonCryocoolers.aspx>.
- H. Jeffreys. An invariant form for the prior probability in estimation problems. *Proc. Royal Soc. of London*, 186:453–461, 1946. doi: 10.1098/rspa.1946.0056.
- L. Jew. *D.Phil Thesis: Measurements of Diffuse Galactic Emission at 5GHz with C-BASS*. University of Oxford, 2017.
- L. Jew et al. The C-Band All-Sky Survey (C-BASS): Simulated parametric fitting in single pixels in total intensity and polarisation. *MNRAS*, in prep.
- B. R. Johnson, G. Yassin, and P. Grimes. CLOVER CDR2 Document: The Optics Design for CLOVER. private communication, 2007.
- M. Jones et al. The C-Band All-Sky Survey (C-BASS): constraining diffuse Galactic radio emission in the North Celestial Pole region. *MNRAS*, 485:2844–2860, 2018. doi: 10.1093/mnras/stz522.
- M. Kamionkowski and D. Kovetz, E. The Quest for B Modes from Inflationary Gravitational Waves. *Annual Review of Astronomy and Astrophysics*, 54:227–269, 2016. doi: 10.1146/annurev-astro-081915-023433.
- W. Kester. Taking the Mystery out of the Infamous Formula, $\text{SNR}=6.02N + 1.76\text{dB}$, and Why You Should Care. 2008a. URL <https://www.analog.com/media/en/training-seminars/tutorials/MT-001.pdf>.
- Walt Kester. ADC Noise Figure An Often Misunderstood and Misinterpreted Specification. 2008b. URL <https://www.analog.com/media/en/training-seminars/tutorials/MT-006.pdf>.
- G. King, O. et al. The C-Band All-Sky Survey: Instrument design, status and first-look data. *SPIE*, 7741:77411I, 2010. doi: 10.1117/12.858011.
- G. King, O. et al. The C-Band All-Sky Survey (C-BASS): design and implementation of the northern receiver. *MNRAS*, 438:2426–2439, 2014. doi: 10.1093/mnras/stt2359.
- I. Kiyotomo. CMB Foreground: A Concise Review. *PTEP*, 2014:06B109, 2014. doi: 10.1093/ptep/ptu065.

- A. Kogut. Synchrotron Spectral Curvature from 22 MHz to 23 GHz. *ApJ*, 753:110, 2012. doi: 10.1088/0004-637X/753/2/110.
- E. Komatsu et al. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Cosmological Interpretation. *ApJ*, 180:330–376, 2009. doi: 10.1088/0067-0049/180/2/330.
- M. Kovac, J. et al. Detection of polarization in the cosmic microwave background using DASI. *Nature*, 420:772–787, 2002. doi: 10.1038/nature01269.
- M. Leach, S. et al. Component separation methods for the PLANCK mission. *A&A*, 491(2):597–615, 2009. doi: 10.1051/0004-6361:200810116.
- Low Noise and Factory. *LNFLNC620C, 620 GHz Cryogenic Low Noise Amplifier*, 2016. URL www.lownoisefactory.com.
- Low Noise Factory. *LNFLNC1529B, 1529 GHz Cryogenic Low Noise Amplifier*, 2017. URL www.lownoisefactory.com.
- T. Luhmann. Close range photogrammetry for industrial applications. *ISPRS*, 65: 558–569, 2010. doi: 10.1016/j.isprsjprs.2010.06.003.
- B. Maffei et al. CLOVER: The CMB polarization observer. *EAS Publication Series*, 14:251–256, 2005. doi: 10.1051/eas:2005039.
- R. Marta et al. A derivation of the free-free emission on the Galactic plane between 20Å and 44Å. *MNRAS*, 422:2429–2443, 2016. doi: 10.1111/j.1365-2966.2012.20796.x.
- J. C. Mather et al. A preliminary measurement of the Cosmic Microwave Background Spectrum by the Cosmic Background Explorer (COBE) satellite. *ApJ*, 354:L37, 1990. doi: 10.1086/185717.
- T. Matsumura et al. Mission design of LiteBIRD. *Journal of Low Temperature Physics*, 176:733–740, 2014. doi: 10.1007/s10909-013-0996-1.
- C. Mendellín, G. SKA Memo 95, Antenna Noise Temperature Calculation. *SKA Collaboration*, 2007. URL <https://www.skatelescope.org/publications/>.
- N. Metropolis, W. Rosenbluth, A., N. Rosenbluth, M., A.H Teller, and E. Teller. *The Journal of Chemical Physics*, 21, 1953. doi: 10.1063/1.1699114.

- D. Michilli et al. Low-frequency pulse profile variation in PSR B2217+47: evidence for echoes from the interstellar medium. *MNRAS*, 476:2704–2716, 2018. doi: 10.1093/mnras/sty368.
- J.L Monteith and M. H. Unsworth. *Principles of Environmental Physics*. Academic Press, 2007.
- K. Morris and M. Luce. DBE Implementation for VLBA. *CASPER2010*, 2010. URL https://casper.ssl.berkeley.edu/wiki/images/3/3a/Workshop_casper2010_vlba_dbe_keith_morris.pdf.
- National Radio Astronomy Observatory (NRAO). Radiometers, 2017. URL <https://www.cv.nrao.edu/course/astr534/Radiometers.html>.
- B. Nikolic. Atmospheric Modelling for ALMA. 2005. URL <https://www.mrao.cam.ac.uk/~bn204/alma/atmomodel.html>.
- N.Macellari et al. Galactic foreground contributions to the 5–year Wilkinson Microwave Anisotropy Probe maps. *MNRAS*, 418:888–905, 2011. doi: 10.1111/j.1365-2966.2011.19542.x.
- D. Olver, A. et al. *Microwave Horns and Feeds*. IEEE, 39 edition, 1994.
- S. Padin et al. First Intrinsic Anisotropy Observations with the Cosmic Background Imager. *ApJ*, 549:L1–L5, 2001. doi: 10.1086/319142.
- A. Page, L. The Wilkinson Microwave Anisotropy Probe. *arXiv eprint*, 2004. URL <https://arxiv.org/abs/astro-ph/0306381>.
- R. Pardo, J., J. Cernicharo, and E. Serabyn. Atmospheric Transmission at Microwaves (ATM): An Improved Model for Millimeter and Submillimeter Applications. *IEEE*, 49:1683 – 1694, 2001. doi: 10.1109/8.982447.
- Patil, Huard, and Fonnesbeck. PyMC: Bayesian Stochastic Modelling in Python. *Journal of Statistical Software*, 35, 2010. doi: 10.18637/jss.v035.i04.
- A. Penzias, A. and W. Wilson, R. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *ApJ*, 142:419–421, 1965. doi: 10.1086/148307.
- PICO Collaboration et al. Scientific Goals and Objectives of PICO – Probe of Inflation and Cosmic Origins. *AAS*, 231:140.17, 2017.

- Planck Collaboration et al. Planck 2013 results. XII. Component Separation. *A&A*, 571:A12, 2013a. doi: 10.1051/0004-6361/201321580.
- Planck Collaboration et al. Planck 2013 results. xv. cmb power spectra and likelihood. *A&A*, 571:A15, 2013b. doi: 10.1051/0004-6361/201321573.
- Planck Collaboration et al. Planck intermediate results. XIX. An overview of the polarized thermal emission from Galactic dust. *A&A*, 576:A104, 2015. doi: 10.1051/0004-6361/201424082.
- Planck Collaboration et al. Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical depth. *A&A*, 596:A107, 2016a. doi: 10.1051/0004-6361/201628890.
- Planck Collaboration et al. Diffuse component separation: CMB maps. *A&A*, 594:A9, 2016b. doi: 10.1051/0004-6361/201525936.
- Planck Collaboration et al. Planck 2018 results. XI. Polarized dust foregrounds. *arXiv eprint*, 2018. URL <https://arxiv.org/abs/1801.04945>.
- F. Poidevin et al. The Quijote Experiment: Prospects for CMB B-mode polarization detection and foregrounds characterization. *arxiv eprint*, 2018. URL <https://arxiv.org/abs/1802.04594v1>.
- A. Pollak. *D.Phil Thesis: Receiver Technology for Radio Astronomy and Deep Space Communications*. University of Oxford, 2018.
- A. Pollak and C. Holler. Continuous Wave Stabilised Receiver. *IEEE*, in prep.
- D. M Pozar. *Microwave Engineering*, chapter 2. Wiley, third edition, 2005a.
- D. M Pozar. *Microwave Engineering*, chapter 8. Wiley, third edition, 2005b.
- C. D Price. Spectrometers and Polyphase Filterbanks in Radio Astronomy. *arXiv eprints*, 2016. URL <https://arxiv.org/pdf/1607.03579v1.pdf>.
- H. Qiu, Y. A novel design for a giant Arecibo-type spherical radio telescope with an active main reflector. *MNRAS*, 301:827–830, 1998. doi: 10.1046/j.1365-8711.1998.02067.x.
- A. Rajantie. Magnetic monopoles in field theory and cosmology. *Royal Society*, 370:5705–5717, 2012. doi: 10.1098/rsta.2011.0394.

- W. Reich et al. The Effelsberg 1.4 GHz Medium Galactic Latitude Survey (EMLS). *The Magnetized Interstellar Medium, Conference Proceedings*, pages 45–50, 2004. URL <https://ui.adsabs.harvard.edu/abs/2004mim..proc...45R>.
- M. Remazeilles. *arXiv eprint*, 2018. URL <https://export.arxiv.org/abs/1806.01026>.
- M. Remazeilles and J. Chluba. Extracting foreground-obscured μ -distortion anisotropies to constrain primordial non-Gaussianity. *MNRAS*, 478:807–824, 2018. doi: 10.1093/mnras/sty1034.
- M. Remazeilles et al. Sensitivity and foreground modelling for large-scale CMB and B-mode polarization satellite missions. *MNRAS*, 458(2):2032–2050, 2016. doi: 10.1093/mnras/stw441.
- M. Remazeilles et al. Exploring Cosmic Origins with CORE: B-mode Component Separation. *Journal of Cosmology and Astroparticle Physics*, 2018(4):023, 2018. doi: 10.1088/1475-7516/2018/04/023.
- G. Riess, A. et al. Cepheid Calibrations of Modern Type Ia Supernovae: Implications for the Hubble Constant. *ApJ*, 183:109–141, 2009. doi: 10.1088/0067-0049/183/1/109.
- ROGERS corporation. RT/duroid Laminates, 2018. URL <https://www.rogerscorp.com/acs/producttypes/6/RT-duroid-Laminates.aspx>.
- A. Rubiño Marín, J. et al. The Quijote CMB Experiment. *Astrophysics and Space Science Proceedings*, 14:127, 2010. doi: 10.1007/978-3-642-11250-8_12.
- J. C. Samuel Chieh et al. Development of a Ku-Band Corrugated Conical Horn Using 3-D Print Technology. *IEEE*, 13:201–204, 2014. doi: 10.1109/LAWP.2014.2301169.
- K. Smith, E. and S. Weintraub. The Constants in the Equation for Atmospheric Refractive Index at Radio Frequencies. *IRE Proc.*, 41:1035 – 1037, 1953. doi: 10.1109/JRPROC.1953.274297.
- M. Stevenson. Derivation of an analytical approximation of the spectrum of spinning dust emission. *ApJ*, 781:113, 2014. doi: 10.1088/0004-637X/781/2/113.
- B. Sutin, M. et al. PICO-Probe of Inflation and Cosmic Origins. *arXiv e-prints*, 2019. URL <https://arxiv.org/abs/1902.10541>.

- A. Suzuki et al. The LiteBIRD Satellite Mission: Sub-Kelvin Instrument. *Journal of Low Temperature Physics*, 193:1048–1056, 2018.
- A. Tauber, J. et al. Planck pre-launch status: The optical system. *A&A*, 520:A2, 2010. doi: 10.1051/0004-6361/200912911.
- A. Taylor. C-Band All-Sky Survey (C-BASS), 2018. URL http://www.iac.es/congreso/cmbforegrounds18/media/talks/day1/Taylor_CBASS_Tenerife_2018_2web.pdf.
- The EBEX Collaboration et al. EBEX: A balloon-born CMB polarization experiment. *SPIE Proc.*, 7741:77411C, 2010. doi: 10.1117/12.857138.
- The LSPE Collaboration et al. The Large-Scale Polarization Explorer (LSPE). *SPIE Proc.*, 8446:84467A, 2012. doi: 10.1117/12.926095.
- The Planck Collaboration et al. Planck 2015 results. XX. Constraints on inflation. *A&A*, 594:A20, 2016. doi: 10.1051/0004-6361/201525898.
- The Polarbear Collaboration et al. The new generation CMB B-mode polarization experiment: POLARBEAR. *arXiv e-prints*, 2010. URL <https://arxiv.org/abs/1011.0763>.
- J. Thornton, R. et al. The Atacama Cosmology Telescope: The polarization-sensitive ACTPol instrument. *ApJ*, 227:21, 2016. doi: 10.3847/1538-4365/227/2/21.
- TICRA. *CHAMP Manual*, 2015a. URL <https://www.ticra.com/>.
- TICRA. *GRASP User's Manual*, 2015b. URL <https://www.ticra.com/software/grasp/>.
- TICRA. *Analysis Methods*, 2015c. URL <https://www.ticra.com/software/grasp/>.
- TICRA. *Useful Geometric Equations for Reflector Antennas*, 2015d. URL <https://www.ticra.com/software/grasp/>.
- R Ulaby, F., K Moore, R., and K. Fung, A. Microwave Remote Sensing. 1, 1981.
- M. Vidal, C. Dickinson, D. Davies, R., and P. Leahy, J. Polarised radio filaments outside the Galactic plane. *MNRAS*, 452(1):656–675, 2015. doi: 10.1093/mnras/stv1328.

- A. Watson, R. et al. Detection of Anomalous Microwave Emission in the Perseus Molecular Cloud with the COSMOSOMAS Experiment. *ApJ*, 624:L89–L92, 2005. doi: 10.1086/430519.
- R. Watson et al. QUIJOTE telescope design and fabrications. *SPIE*, 7733:77330Z, 2010. doi: 10.1117/12.857286.
- L. Wilson, T., K. Rohlfs, and Hüttemeister. *Tools of Radio Astronomy*. Springer, fifth edition, 2008.
- R. Wilson. Radio Receiver Architectures and Analysis. 2012. URL <https://pdfs.semanticscholar.org/cc65/651a756049eb23acaa7ec03269877c5f6644.pdf>.
- XILINX. *Zynq UltraScale + RFSoc Data Sheet: Overview*, 2018. URL <https://www.xilinx.com/products/silicon-devices/soc/rfsoc.html>.
- XILINX. Field Programmable Gate Array (FPGA). 2019. URL <https://www.xilinx.com/products/silicon-devices/fpga/what-is-an-fpga.html>.
- J. Yardly. Timeline of the Universe, 2015. URL planck.cf.ac.uk/science/timeline/universe/.
- G. Yassin, B. Sørensen, S., and K. Grimes, P. Compact Optical Assemblies for Large-Format Imaging arrays. *ISSTT*, pages 89–92, 2005. URL <https://ui.adsabs.harvard.edu/abs/2005stt..conf...89Y>.
- M. Zaldarriaga. Cosmic Microwave Background Polarization Experiments. *ApJ*, 503: 1–15, 1998.
- B. Zhang et al. Metallic, 3D-Printed, K-Band-Stepped, Double-Ridged Square Horn Antennas. *Applied Science*, 8, 2018. doi: 10.3390/app8010033.
- Zotefoam, 2010. URL <http://zotefoams.com/pages/EN/datasheets/LD45.HTM>.