Measurements of Diffuse Galactic Emission at 5 GHz with C-BASS

Luke Jew

Oriel College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

August 2017

Acknowledgements

First and foremost, thank you Amanda for always being there to support and encourage me, keeping me company and keeping me sane! I would like to thank my Mom for her wise advice and the rest of my family for generally being an ace group of people.

Thank you to the C-BASS collaboration, you have been a pleasure to work with. In particular, a big thank you goes to Mike and Jamie in Oxford both for sharing their knowledge and guidance but also for making the last four years such great fun! Finally, a massive thank you goes to my supervisor Angela. Thank you for your continuous support, motivation and expertise.

Abstract

The C-Band All-Sky Survey (C-BASS) is a project to produce an all-sky map in intensity and polarization at a central frequency of 5 GHz with 1 GHz bandwidth and $\lesssim 1^\circ$ resolution. The central frequency is low enough for the map to be dominated by synchrotron and free-free emission but high enough so that Faraday rotation and depolarization are small across most of the sky. The C-BASS map will enable a more accurate removal of contaminating foregrounds from measurements of the cosmic microwave background, particularly in polarization where the B-mode signal from inflation is likely to be orders of magnitude weaker than the diffuse Galactic foreground emission.

To produce an all-sky map from the ground requires two telescopes, one in the northern and one in the southern hemisphere. This thesis focuses on analysis of C-BASS North data.

The noise properties of time-ordered data are characterised by fitting a noise model to periodograms. Using simulations, the errors introduced into the C-BASS maps by a destriping mapmaker are quantified and we reduce the signal error by masking the brightest pixels during baseline offset estimation. Jackknife tests are used to test the C-BASS data for systematics and to test the accuracy of the sensitivity maps. In total intensity, the spectral index of diffuse Galactic emission between 5 GHz and 408 MHz is measured using an extended T-T plot method and the results are compared to simulations. The spectral index of polarized diffuse Galactic emission between 5 GHz and 30 GHz is estimated in 55 arcmin pixels, modelling the polarized intensity as a Rician random variable.

Contents

1 The cosmic microwave background and its foregrounds **1**

1.1 The cosmic microwave background 2

1.1.1 The expanding Universe 2

1.1.2 Cosmic inflation 4

1.1.3 The origins of the cosmic microwave background radiation . . 5

1.1.4 The CMB angular power spectrum 9

1.2 CMB foregrounds 13

1.2.1 Diffuse Galactic synchrotron emission 14

1.2.2 Free-free emission 18

1.2.3 Thermal dust emission 19

1.2.4 Anomalous microwave emission (AME) 20

1.2.5 Other foregrounds 21

1.3 Foreground separation techniques 25

1.3.1 Internal linear combination 25

1.3.2 Template fitting 27

1.3.3 Independent component analysis 28

1.3.4 Bayesian parametric model fitting 29

1.4 Observations 30

1.4.1 CMB temperature fluctuations and E-mode power spectrum . 30

1.4.2 CMB B-mode power spectrum 32

1.4.3 Ancillary surveys 32

1.5 C-BASS and thesis layout 35

2 The C-Band All-Sky Survey (C-BASS) **37**

2.1 Overview of the survey and the telescopes 37

2.1.1 Optics 39

2.1.2 Basic radiometers 40

2.1.3 Polarimetry 45

2.1.4 Scan strategy . . . 47
2.1.5 Target sensitivity . . . 49
2.2 Component separation forecasts . . . 52

3 C-BASS North Data Analysis 1: Time Ordered Data 61
3.1 Data pipeline . . . 61
3.1.1 Pointing correction . . . 62
3.1.2 Gain and polarization angle calibration . . . 65
3.1.3 RFI flagging . . . 66
3.1.4 Ground removal . . . 67
3.1.5 1.2 Hz contamination . . . 68
3.1.6 Astronomical calibration . . . 70
3.1.7 Pipeline summary . . . 71
3.2 C-BASS data subsets . . . 72
3.3 TOD noise properties . . . 72
3.3.1 Noise model . . . 73
3.3.2 Statistical properties of the periodogram . . . 74
3.3.3 Least-squares method . . . 77
3.3.3.1 Tests on simulated data . . . 79
3.3.3.2 Recovered noise statistics from NCP stares . . . 79
3.3.4 Bayesian method . . . 81
3.3.4.1 Testing choice of prior with simulated data . . . 81
3.3.4.2 Testing sky subtraction on simulated data . . . 84
3.3.4.3 Recovered noise statistics from survey scans . . . 87
3.3.5 Noise statistics summary . . . 93

4 C-BASS North Data Analysis 2: Maps 95
4.1 Mapmaking and the destriping technique . . . 95
4.2 Destriping errors . . . 98
4.3 DESCART . . . 99
4.3.1 Masking to reduce signal error . . . 100
4.3.2 Residual correlated noise . . . 107
4.4 Jackknife tests . . . 112
4.4.1 Procedure . . . 112
4.4.2 Jackknife results . . . 116
4.4.2.1 Even/odd numbered schedules . . . 118
4.4.2.2 Sun cuts . . . 123

4.4.2.3 First half/second half . . . 129
4.4.2.4 Seasons . . . 132
4.4.2.5 Receiver channels . . . 136
4.4.2.6 Summary of jackknife results . . . 137
4.5 Northern survey maps . . . 139
4.5.1 Total intensity . . . 139
4.5.2 Polarization . . . 145
4.6 Summary of preliminary C-BASS maps . . . 152

5 Constraints on the Total Intensity Emission of Diffuse Galactic Foregrounds 158
5.1 A 3-colour map from Haslam, C-BASS and *WMAP* data . . . 158
5.2 Constraining the synchrotron spectral index with T-T plots . . . 164
5.2.1 Overview of the T-T plot method . . . 165
5.2.2 Voronoi binning . . . 166
5.2.3 Straight line fitting with errors in both axes . . . 167
5.2.4 Extension to standard T-T plot method . . . 169
5.2.5 Tests on simulated data . . . 171
5.2.5.1 Only synchrotron emission . . . 172
5.2.5.2 Synchrotron and free-free components . . . 176
5.2.6 Results from real data . . . 179
5.2.7 Discussion of results from real data . . . 188
5.2.8 Barnard's Loop . . . 188
5.3 Discussion . . . 195

6 Constraints on the Polarized Diffuse Galactic Synchrotron Spectral Index 197
6.1 Polarization datasets and a rudimentary estimate of the spectral index between 5 and 30 GHz . . . 198
6.1.1 Datasets . . . 198
6.1.2 Checking the calibration of the C-BASS P map . . . 199
6.1.3 Rudimentary estimate of spectral index . . . 202
6.2 Bayesian analysis of the spectral index of polarized emission between 5 and 30 GHz . . . 203
6.2.1 Model and choice of priors . . . 203
6.2.2 Tests on simulated datasets . . . 205
6.2.3 Results from real data . . . 209

7 Summary and Future Work **215**
7.1 The CMB, its foregrounds and C-BASS 215
7.2 C-BASS data analysis 216
7.3 Measurements of diffuse Galactic emission 217
7.4 C-BASS South 219

A Temperature definitions **220**

B Bayesian statistics and Markov Chain Monte Carlo methods **225**

Bibliography **228**

Chapter 1

The cosmic microwave background and its foregrounds

The cosmic microwave background (CMB) is radiation released from the surface of last scattering around 380,000 years after the Big Bang. The CMB provides one of the most important routes to understanding cosmology and the early Universe. Since its discovery, increasingly accurate measurements of the CMB angular power spectrum have allowed us to better constrain cosmological parameters and these are now a cornerstone of modern cosmology. As the sensitivity of observations has increased it has become more important to accurately account for the contaminating foreground emission from our own solar system, from the Milky Way Galaxy and from beyond. Foreground emission is now the limiting factor in CMB measurements. In particular the polarized CMB signal from cosmic inflation may be weaker than the foregrounds across the entire sky.

To better understand the foreground emission, sky maps at a range of frequencies are required. The C-Band All-Sky Survey (C-BASS) is a project to produce an all-sky map at a central frequency of 5 GHz that will provide significant insight into the diffuse low frequency foregrounds to the CMB.

In this chapter we introduce cosmic inflation and the CMB (Section 1.1), foreground emission to the CMB (Section 1.2), methods of separating the CMB from those foregrounds (Section 1.3) and the current status of CMB observations (Section 1.4). In the final section of this chapter we explain how C-BASS fits into the context of this exciting field and set out the contents of this thesis (Section 1.5).

1.1 The cosmic microwave background

We begin by introducing the expanding universe and the standard Big Bang model (Section 1.1.1), cosmic inflation (Section 1.1.2), the origins of the CMB (Section 1.1.3) and the CMB angular power spectrum (Section 1.1.4) before introducing C-BASS and setting out the structure of the rest of this thesis (Section 1.5).

1.1.1 The expanding Universe

This section contains a brief introduction to the standard Big Bang model and the expanding Universe.

When Einstein was developing his theory of general relativity he thought that the universe was static, or unchanging with time (Einstein, 1915, 1916). He introduced the cosmological constant, Λ, into his field equations in order to enforce a such a static universe. Friedmann demonstrated that a static universe is unstable and solved Einstein's equations in the case of an expanding or contracting universe (Friedmann, 1922). A few years later Lemaître (1927) showed that the expansion rate of the Universe could be measured from the redshifts of distant galaxies and also realised that if the Universe was observed to be expanding then it must be of finite age and have emerged from "a single quantum" (Lemaître, 1931).

At the end of the 1920s Hubble measured the recession velocities and distances of 46 galaxies. He found that their recession velocities were proportional to their distances from us and therefore demonstrated that the Universe was not static but expanding (Hubble, 1929).

In the standard Big Bang model the Universe started with an initial injection of energy and expanded out, with a decelerating expansion rate due to the mutual gravitational attraction of all the matter in the Universe. Evidence now suggests that the expansion rate of the Universe is actually accelerating as though driven by a cosmological constant (Riess et al., 1998; Perlmutter et al., 1999) and that maybe Einstein was right to include Λ all along, if only for the wrong reasons!

The Friedmann-Robertson-Walker metric is an exact solution to Einstein's field equations for a homogenous, isotropic and expanding (or contracting) universe,

$$ds^2 = -c^2dt^2 + a^2(t)\left(\frac{dr^2}{1-kr^2} + r^2d\theta + r^2\sin^2(\theta)d\phi^2\right), \tag{1.1}$$

where $a(t)$ is the scale factor defined to equal unity today and k is the curvature.

From this metric we get the Friedmann equations,

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

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

where ρ is the energy density, P is the pressure and Λ is the cosmological constant. These equations determine the evolution of the expansion of the universe.

It is useful to rewrite the Friedmann equations in term of the Hubble parameter and the density parameters. The Hubble parameter, $H = \frac{\dot{a}}{a}$, relates the recession velocity of distant objects in the Universe to their distance away from us. It is often measured in km $\mathrm{s}^{-1}\mathrm{Mpc}^{-1}$. The density parameters express the energy density as a fraction of the critical density required for a flat universe with zero curvature. If we subsume Λ into ρ then, for a given H, the critical density needed for a flat universe is $\rho_\mathrm{c} = \frac{3H^2}{8\pi G}$. We can use ρ_c to define the density parameter $\Omega_\mathrm{tot}(t) = \frac{\rho(t)}{\rho_\mathrm{c}}$. The density parameter is the energy density of the Universe expressed as a fraction of the critical density and includes contributions from all sources of energy, $\Omega_\mathrm{tot} = \Omega_\mathrm{matter} + \Omega_\mathrm{radiation} + \Omega_\Lambda$.

The Friedmann equation can then be rewritten,

$$\Omega_\mathrm{tot} - 1 = \frac{k}{a^2H^2}, \tag{1.4}$$

and we can define a density parameter for curvature such that,

$$\Omega_\mathrm{tot} + \Omega_k = 1. \tag{1.5}$$

In an expanding universe the physical distance, d, between two points changes with time. It is therefore useful to define comoving distances that remain fixed, $\chi = d/a$. The particle horizon, χ_p is the maximum comoving distance that light can travel since the beginning of time $t = 0$,

$$\chi_\mathrm{p} = \int_0^t c\frac{dt}{a(t)} = \int_0^a \frac{c}{aH}\frac{da}{a}. \tag{1.6}$$

Objects will only be causally connected if their separation is less than the particle horizon, $\lambda < \chi_\mathrm{p}$. The Hubble radius, $\chi_H = \frac{c}{aH}$, is the distance at which the recession velocity equals the speed of light and sets the size of our observable universe. χ_p is the logarithmic integral of χ_H and in a standard universe $\chi_\mathrm{p} \propto \chi_H$.

The Hubble parameter evaluated today, H_0, is sometimes expressed as a multiple of the dimensionless Hubble parameter, h,

$$H_0 = h \times 100\,\mathrm{kms}^{-1}\mathrm{Mpc}^{-1}. \tag{1.7}$$

1.1.2 Cosmic inflation

The inflationary paradigm, originally proposed by Guth (1981), has revolutionised modern cosmology by solving several major problems of the standard Big Bang model. Cosmic inflation explains the horizon, flatness and magnetic monopole problems, which are discussed further below, as well as providing a mechanism for seeding the perturbations that evolved to become the large scale structure that is seen in the Universe today. For all of inflation's successes it has yet to be confirmed with direct observational evidence. This section contains a brief description of cosmic inflation, how it solves the problems with the standard Big Bang model and the possible direct observational evidence.

Theories of inflation posit that a tiny fraction of a second after the big bang ($\sim 10^{-36}$ s) the Universe underwent a massive expansion in a very short period of time, of order 50 e-foldings in 10^{-34} s.

The first problem with the Big Bang model is called the horizon problem. The Universe is observed to be very isotropic and homogenous on large scales. The standard Big Bang model predicts that the Hubble radius should increase with time and so larger and larger scales should only just be coming into causal contact now. Yet somehow opposite sides of the observable Universe that are separated by the Hubble radius are already in thermodynamic equilibrium, even though they are only just now apparently able to exchange information with each other.

For this to make sense the particle horizon must be much greater than the Hubble radius, $\chi_{\rm p} \gg \chi_H$. This could be the case if the Hubble radius used to be much larger than it is today before decreasing for some period of time. A decreasing Hubble radius, $\dot{\chi_H} < 0$, corresponds to an accelerating expansion, $\ddot{a} > 0$. A period of accelerating expansion means that regions appearing to only come into causal contact now were in fact causally connected in the past. If the current observable Universe was at one point significantly smaller than the particle horizon it could have reached thermodynamic equilibrium before being expanded out.

The second problem is called the flatness problem. From Equation 1.4, if the Universe were perfectly flat with $k = 0$ then $\Omega_{\rm tot}$ would remain unchanged from unity for all time. If, however, k deviated even slightly from 1 then $\Omega_{\rm tot}$ would quickly diverge from unity as the Universe aged. To match the measured value of $\Omega_{\rm tot}$ today would have required incredible fine tuning at early times. Inflation solves this by causing the denominator of the right hand side of Equation 1.4 to increase during the

brief period of rapid expansion,

$$0 < \ddot{a} = \frac{d(\dot{a})}{dt} = \frac{d(aH)}{dt}, \tag{1.8}$$

and therefore driving Ω_{tot} to unity.

Thirdly is the magnetic-monopole problem. Grand unified theories predict that stable magnetic monopoles should have been produced at the very high temperatures of the early Universe and that these monopoles should persist to the present day, yet no magnetic monopole has ever been observed. If the monopoles were produced before inflation, the rapid expansion would dilute their density to unobservably low levels afterwards.

Additionally the rapid expansion during inflation would have created small inhomogeneities that could evolve into the structure in the Universe today. Inflation would have stretched out the quantum fluctuations that existed before it started into over and under densities after it had finished.

Finally cosmic inflation is also predicted to have created primordial gravitational waves by stretching out the primordial graviton field. A direct detection of primordial gravitational waves would be the "smoking gun" evidence that inflation occurred. These gravitational waves may be detectable through their effect on the polarization of the CMB, which is discussed in the following section.

1.1.3 The origins of the cosmic microwave background radiation

Fast-forwarding from inflation and the first fraction of a second after the Big Bang to several minutes and a lot has happened in the Universe. The four fundamental forces have become distinct, most of the matter has been annihilated in collisions with antimatter and nucleosynthesis is complete. A few minutes after the Big Bang the Universe existed as a hot dense plasma of baryons, electrons and photons. In this section we explain how the CMB came to be released from this primordial plasma and why it is polarized.

Thanks to cosmic inflation the primordial plasma of the early Universe was very homogenous and isotropic, but not perfectly so. Where there was a slight excess of matter, gravity would pull in even more. As baryonic matter fell into the gravitational wells of the over dense regions it compressed and heated up. This in turn increased the radiation pressure in these regions. Eventually the pressure was great enough to overcome the gravity and the plasma would rarefy, cool and the radiation pressure

drop back down. All the while non-collisional dark matter would continue to fall into the wells. Eventually gravity would overcome the radiation pressure and once again plasma would collapse back into the wells. This repeated compression and rarefaction of the plasma set up acoustic oscillations, or sound waves, rippling through the primordial plasma and as time went by larger and larger scales began to oscillate.

The primordial plasma was kept ionised by frequent high energy particle collisions. In the primordial plasma the number density of photons, n_γ, was significantly greater than the number density of baryons, n_b, with a ratio $n_{\gamma\mathrm{b}} = \frac{n_\gamma}{n_\mathrm{b}} \sim 10^9$, and initially photons typically had energies much greater than the ionisation potential of neutral Hydrogen (1 Ry=13.6 eV). This meant that any atom that formed was almost instantly ionised in a photon collision. The frequent collisions maintained thermal equilibrium and also ensured that the photon mean free path (MFP) was short (MFP $\sim 1/n_e\sigma_e$ where n_e is the number density of electrons and σ_e is the scattering cross section).

As the Universe expanded, the mean photon energy dropped. Eventually it dropped below 1 Ry but, because $n_{\gamma\mathrm{b}}$ was large, there were still enough photons in the tail of the photon energy distribution to keep the plasma ionised. As the Universe continued to expand, the mean photon energy continued to drop further below 1 Ry. Eventually there were insufficient high energy photons to keep the plasma fully ionised and its ionisation fraction began to fall. Fewer photons had enough energy to excite bound electrons to even the first excited state and the photon MFP grew. Once the MFP reached the Hubble distance, photons decoupled from matter and the Universe became transparent for the first time. This occurred at a redshift of around $z \sim 1100$ (380,000 years after the Big Bang) when the ionisation fraction had dropped to around 10%.

The radiation released during this transition from an opaque to a transparent Universe is called the cosmic microwave background (CMB) and the apparent surface from which it was released is called the surface of last scattering (SLS).

At the SLS the photons had a thermal distribution of energies with a characteristic temperature of ~ 3000 K. The continued adiabatic expansion of the Universe maintained the thermal distribution of photon energies but lowered the characteristic temperature. Today the temperature of the CMB is measured to be $T_0 = 2.72548 \pm 0.00057$ K (Fixsen, 2009). The emission of a blackbody at 3000 K peaks in the visible part of the spectrum. The emission of a blackbody at ~ 3 K peaks at 217 GHz, which is where the CMB is observed today.

The CMB is isotropic to a very high accuracy but there are fluctuations of a few parts in 10^5 (Wright, 2004). In a direction $\hat{\mathbf{n}}$ the temperature fluctuation $\delta T/T$ is given by the sum of a number of physical effects,

$$\frac{\delta T}{T} = \frac{\delta \rho_\gamma}{4\rho_\gamma} - \mathbf{v}_\mathrm{B} \cdot \hat{\mathbf{n}} - \Phi - 2\int_{\eta_\mathrm{SLS}}^{\eta_0} \dot{\Phi}\, \mathrm{d}\eta. \tag{1.9}$$

The first term, $\frac{\delta \rho_\gamma}{4\rho_\gamma}$, is called the intrinsic term where ρ_γ is the density of radiation. The temperature of the radiation was higher where it was more compressed. The second term, $-\mathbf{v}_\mathrm{B} \cdot \hat{\mathbf{n}}$, is the Doppler term where $\mathbf{v}_\mathrm{B}$ is the baryon velocity at the SLS. Each CMB photon received a Doppler shift in its final collision with a baryon. The third term, $-\Phi$, is the Sachs-Wolfe term where Φ is the local gravitational potential. The Sachs-Wolfe effect is the decrease in temperature caused by a photon escaping its local gravitational well (Sachs & Wolfe, 1967). The final term, $-2\int_{\eta_\mathrm{SLS}}^{\eta_0} \dot{\Phi}\, \mathrm{d}\eta$, is the integrated Sachs Wolfe effect, which is the net gravitational blue/redshifting of the CMB photon as it travelled from the SLS to now.

The CMB is linearly polarized and Figure 1.1 shows how this is caused by quadrupole temperature anisotropies at the SLS. Thomson scattering introduces polarization aligned with the cold axis of a local quadrupole. Quadrupole anisotropies arise from scalar, vector and tensor perturbations caused by density perturbations, velocity perturbations and primordial gravitational waves respectively. A useful basis to decompose the polarization vectors into are E and B-modes, shown in Figure 1.2. Whilst the E-modes can be produced by scalar (density), vector (velocity) and tensor (gravitational waves) perturbations, B-modes can only be produced by tensor perturbations and as such are a direct probe of primordial gravitational waves and the inflationary epoch (Dodelson, 2003). The ratio of power in tensor perturbations to those in scalar perturbations is called the tensor-to-scalar ratio, r, and is determined by the energy scale of inflation. Different inflationary models predict different non-zero values of r.

The CMB is not expected to be significantly circularly polarized. Circular polarization could be introduced by the first stars, primordial magnetic fields and beyond the Standard model particle physics. See King & Lubin (2016) and the references therein for details.

There are three major difficulties in measuring CMB B-modes and hence determining r. Firstly, r has no theoretical lower limit and it could be vanishingly small. Secondly, on small angular scales CMB E-modes can be gravitationally lensed into B-modes. Unless the E-modes can be accurately de-lensed, r can only be estimated from the largest scales. Thirdly, across the whole sky and at all frequencies, the

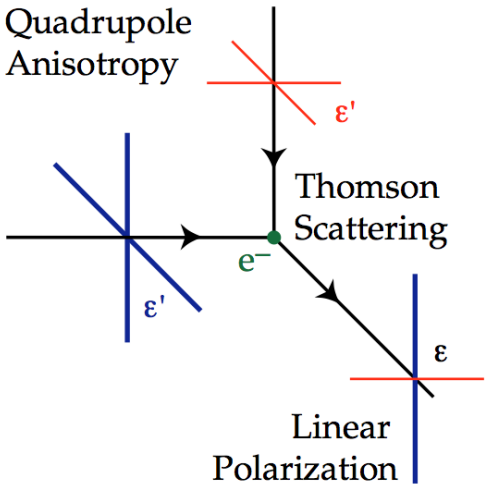

Figure 1.1: Thomson scattering of photons with a quadrupolar distribution of intensities results in polarized emission aligned with the cold (red) axis (Hu & White, 1997). Figure reprinted from New Astronomy, Volume 2, Hu & White, A CMB polarization primer, 323-344, Copyright 1997, with permission from Elsevier.

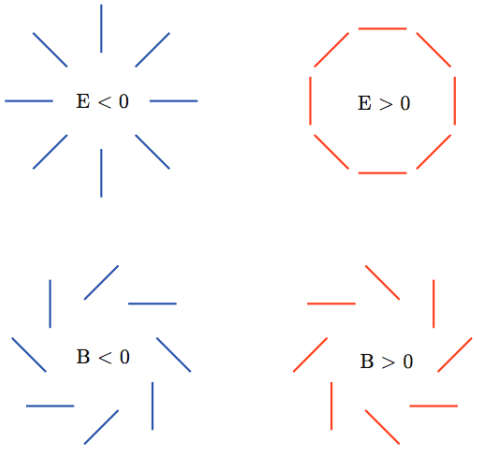

Figure 1.2: E and B-mode polarization patterns. Figure reprinted from Baumann et al. (2009), with the permission of AIP Publishing.

CMB B-mode signal is weaker than the diffuse Galactic foregrounds, which necessitates highly accurate foreground removal. Figure 1.3 shows the predicted frequency spectra of the CMB and of diffuse Galactic foregrounds on 1° scales at high Galactic latitudes (see Section 1.2 for a discussion of foregrounds to the CMB). In total intensity, at around 100 GHz the CMB signal is at least equal to, if not greater than, the contaminating foregrounds across large areas of the sky. The E-mode CMB signal is subdominant to foregrounds at all frequencies and the CMB B-mode signal is orders of magnitude smaller still.

1.1.4 The CMB angular power spectrum

The CMB temperature anisotropies, $\Theta(\hat{\mathbf{n}})$, are a scalar field defined on the surface of a sphere and as such can be decomposed into a sum of spherical harmonics,

$$\Theta(\hat{\mathbf{n}}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^{T} Y_{\ell m}(\hat{\mathbf{n}}), \tag{1.10}$$

where $Y_{\ell m}$ are the spherical harmonics, ℓ is the multipole moment (related to the angular scale) and m encodes the orientation of the spherical harmonics.

The polarization anisotropies can be expressed as Stokes Q and U parameters. Stokes parameters are more useful than polarization angles and amplitudes because angles cannot be sensibly averaged, summed or differenced whereas Stokes parameters are powers and therefore can be due to conservation of energy.

The Stokes Q and U parameters transform as a spin-2 quantity and so the polarization field can be decomposed into spin-2 spherical harmonics;

$$(Q \pm iU)(\hat{\mathbf{n}}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} {}_{(\pm 2)}a_{\ell m (\pm 2)} Y_{\ell m}. \tag{1.11}$$

The E and B coefficients are then

$$a_{\ell m}^{E} = -\frac{1}{2}\left({}_{(+2)}a_{\ell m} + {}_{(-2)}a_{\ell m}\right) \tag{1.12}$$

$$a_{\ell m}^{B} = -\frac{i}{2}\left({}_{(+2)}a_{\ell m} - {}_{(-2)}a_{\ell m}\right). \tag{1.13}$$

If we assume that the Universe is isotropic and has no preferred direction then the statistical average properties of the $a_{\ell m}^{T,E,B}$ can only depend on ℓ and not on m. The CMB anisotropies are measured to be highly Gaussian with a mean of zero and variance $C_{\ell}^{TT,EE,BB}$,

$$\langle a_{\ell m} \rangle = 0 \tag{1.14}$$

$$\langle a_{\ell m} a_{\ell' m'}^{*} \rangle = \delta_{\ell \ell'} \delta m m' C_{\ell}. \tag{1.15}$$

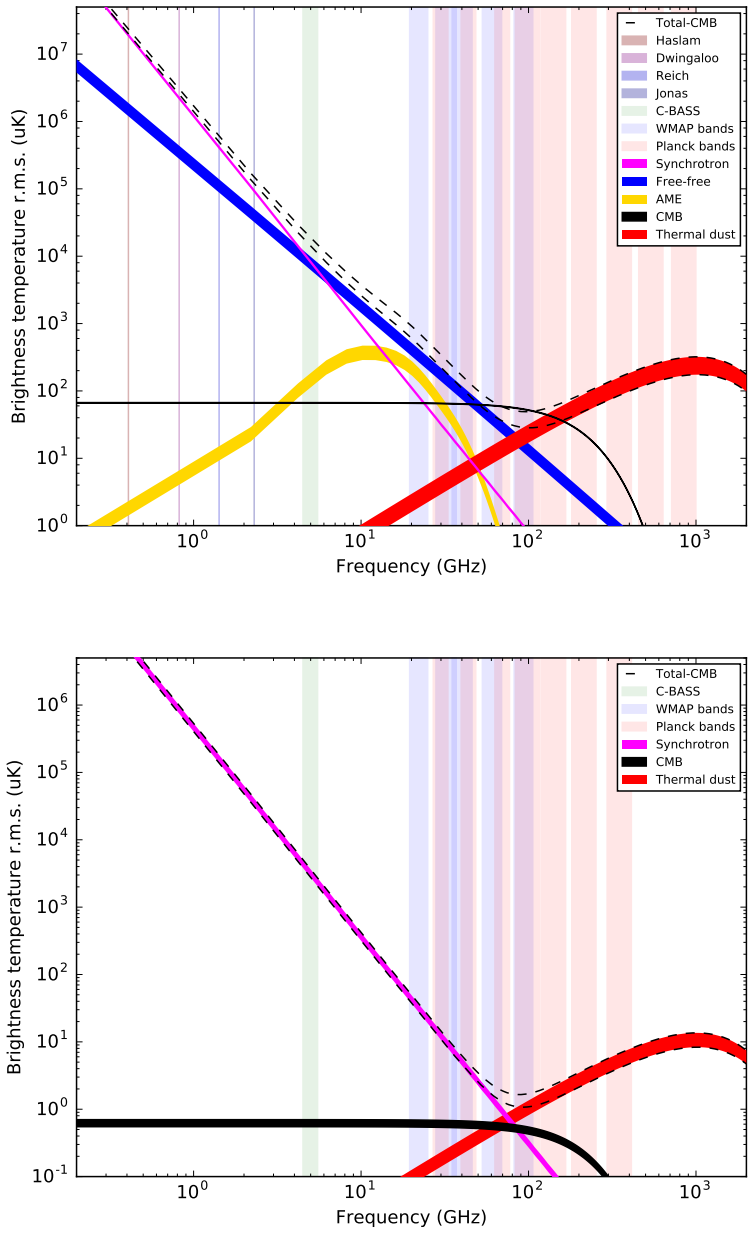

Figure 1.3: Frequency spectra of the CMB and diffuse foregrounds in temperature (*top*) and polarization (*bottom*) at 1° FWHM resolution. The lower bound of each component (except for the E-mode polarization) is the RMS fluctuation of that component in the Planck Collaboration et al. (2016b) model for the region outside the *Planck* 2015 HFI Galactic plane masks that include 80% of the sky. The upper bound of each line is the r.m.s. fluctuation outside the mask that includes 90% of the sky. The E-mode polarization amplitude has been calculated from the Planck Collaboration et al. (2016b) best-fit power spectrum. Underlaid are the bands of *Planck*, *WMAP*, C-BASS and other lower frequency radio surveys. Figure taken from Jones et al. (2017, in prep.).

For T, E and B, each C_ℓ can be estimated from the $2\ell+1$ independent $a_{\ell m}$ values,

$$C_\ell^{TT} = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell}\left(a_{\ell m}^{T*}a_{\ell m}^{T}\right)^2, \tag{1.16}$$

$$C_\ell^{EE} = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell}\left(a_{\ell m}^{E*}a_{\ell m}^{E}\right)^2, \tag{1.17}$$

$$C_\ell^{BB} = \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell}\left(a_{\ell m}^{B*}a_{\ell m}^{B}\right)^2. \tag{1.18}$$

The finite number of spherical harmonic modes at a given ℓ means there is a fundamental limit on how accurately C_ℓ can be measured. This limit is called cosmic variance,

$$\left(\frac{\Delta C_\ell}{C_\ell}\right)_{\text{cosmic variance}} = \sqrt{\frac{2}{2\ell+1}}. \tag{1.19}$$

Cosmological models cannot predict the individual $a_{\ell m}$s but instead predict the distribution from which they are drawn and therefore C_ℓ values. The models can be tested by comparing them to the observed power spectrum.

We now briefly describe the main features in the CMB power spectrum, focusing on the TT, EE and BB spectra in turn. Theoretical CMB power spectra for a range of tensor-to-scalar ratios along with the diffuse Galactic foregrounds are shown in Figure 1.4.

Inflation set up scale-invariant fluctuations in the early Universe and the TT spectrum of scale invariant fluctuations is proportional to ℓ^{-2}. The acoustic oscillations in the primordial plasma left peaks in the TT spectrum. The Hubble distance, $\chi_{\rm H}$, at the SLS corresponds to features around 1° ($\ell \sim 180$) in the CMB. The first acoustic peak should therefore be at around this scale and corresponds to a single compression of the primordial plasma. The second peak corresponds to smaller scales that had time for one compression and one rarefaction. The third peak corresponds to even smaller scales with enough time until the SLS for two compressions, and so on. The precise location of the first peak is a probe of the curvature k. The ratio of power in odd numbered acoustic peaks to the power in even numbered peaks is determined by the baryon density. The more baryons in the Universe the stronger the compressions of the primordial plasma in comparison to the rarefactions. The TT spectrum at low ℓ is a direct measure of the fluctuations generated by inflation and at high ℓ is determined by damping effects (Silk, 1968). The overall tilt of the spectrum is determined by the properties of the cosmic inflation that seeded the initial perturbations.

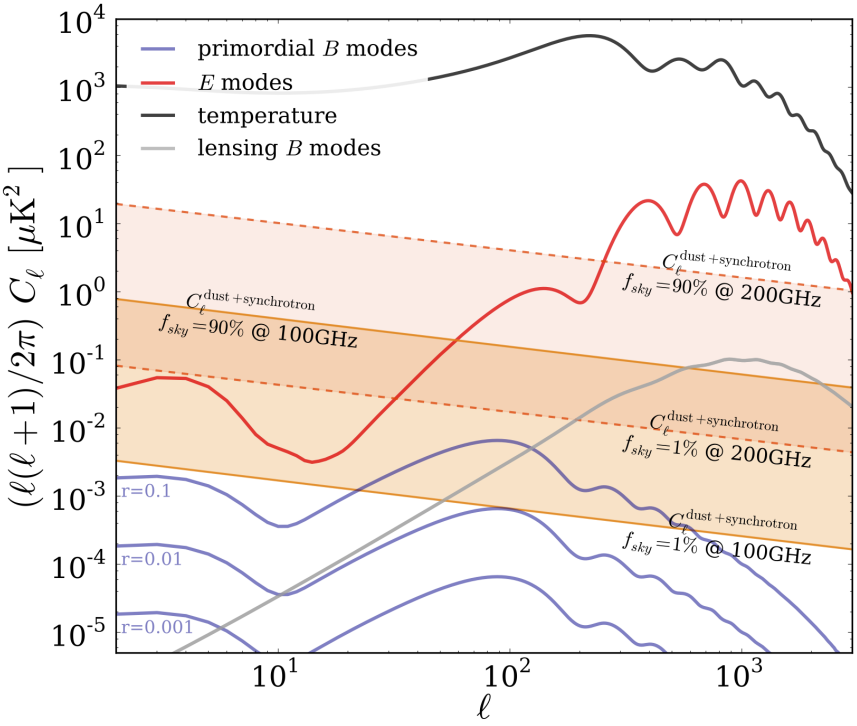

Figure 1.4: CMB TT, EE and BB angular power spectra. The BB spectra has been separated into its lensed component and the component from primordial gravitational waves. The thermal dust and synchrotron foregrounds expected in the cleanest 1% and 90% are also shown. Figure reprinted from Errard et al. (2016), ©SISSA Medialab Srl. Reproduced by permission of IOP Publishing. All rights reserved.

The acoustic peaks are measurable in the EE spectrum too but are shifted relative to the TT peaks. The quadrupole anisotropies needed to polarize the CMB were created by velocity gradients in the primordial plasma and so the polarized signal is greatest when the velocities are greatest. The acoustic peaks in the EE spectrum are therefore out of phase with the TT peaks. The later reionization of the Universe introduces extra power on large scales and so the EE spectrum contains a wealth of information on this later period in the history of the Universe. For a review of observable consequences of reionization see, for example, Reichardt (2016) and the references therein.

The BB spectrum is the linear sum of the signal from primordial gravitational waves and gravitationally lensed E-modes. Like the first acoustic peak, the primordial BB spectrum peaks at the angular scale corresponding to the Hubble length at the SLS and decays quickly at smaller angular scales. At small angular scales the lensed E-modes dominate and so to determine r from small angular scales requires the signal to be de-lensed.

Table 1.1: Base ΛCDM parameter constraints from the full likelihood analysis of Planck Collaboration (2015).

Parameter	Value	Description
$\Omega_b h^2$	0.02225 ± 0.00016	Density of baryonic matter today
$\Omega_c h^2$	0.1198 ± 0.0015	Density of cold dark matter today
$100\theta_{MC}$	1.04077 ± 0.00032	Approximation to the angular scale of the sound horizon at SLS
τ	0.079 ± 0.017	Optical depth due to reionization
$\ln(10^{10} A_s)$	3.094 ± 0.034	Log power of primordial curvature perturbations at $0.05\,\mathrm{Mpc}^{-1}$
n_s	0.9645 ± 0.0049	Spectral index of primordial scalar perturbations

The Λ - Cold Dark Matter (ΛCDM) model is now well established as the "Standard Cosmological Model". With six free parameters, summarised in Table 1.1, ΛCDM can accurately fit not only the CMB power spectra but other cosmological observations including: the relative abundances of hydrogen, helium and lithium; the Type Ia supernovae magnitude-distance relation (Betoule et al., 2014); measurements of baryon acoustic oscillations (Anderson et al., 2014; Alam et al., 2016); the clustering of galaxies on large scales (Jenkins et al., 1998); and cosmic shear (Kitching et al., 2014; Troxel et al., 2017).

1.2 CMB foregrounds

There are many sources of microwave radiation between the SLS and a radio telescope on the Earth today. These sources can be extragalactic, Galactic and even from within our solar system. This subsection briefly outlines the major components of the contaminating foreground radiation to the CMB.

On large angular scales the diffuse emission from our Galaxy dominates much of the radio and microwave sky. The frequency spectra of the CMB and components of the diffuse Galactic emission in a typical 1° pixel were shown in Figure 1.3. The Galactic emission includes:

- Diffuse Galactic synchrotron radiation, described in Section 1.2.1.
- Free-free emission, described in Section 1.2.2.
- Thermal dust emission, described in Section 1.2.3.
- Anomalous microwave emission, described in Section 1.2.4.

Three further foregrounds to the CMB are described in Section 1.2.5. These include

- Line emission, which dominates on large angular scales and at narrow frequency ranges.
- Extragalactic point sources, which are important on small angular scales through both their emission and the Sunyaev-Zel'dovich effect.
- Zodiacal emission.

1.2.1 Diffuse Galactic synchrotron emission

The spectra in Figure 1.3 show that synchrotron emission is the dominant component of Galactic emission in total intensity up to $\sim 10\,\mathrm{GHz}$ and up to $\sim 70\,\mathrm{GHz}$ in polarization. Synchrotron radiation is emitted when highly relativistic charged particles are accelerated along a curved path. Diffuse Galactic synchrotron emission is radiated by cosmic ray electrons and positrons spiralling in the Galactic magnetic field. The high energy electrons are injected via supernova explosions in the Galactic plane before they diffuse into the Galactic halo. If the Galactic magnetic field is ordered, diffuse synchrotron radiation can in principle be up to 75% polarized, at high Galactic latitudes it is typically measured to be up to $\sim 40\%$ polarized (Vidal et al., 2015; Planck Collaboration et al., 2015d). Where the magnetic field is more tangled the polarization fraction will be lower.

The propagation of cosmic rays through the Galaxy and the radiation that they then emit can be modelled with codes such as GALPROP (Moskalenko & Strong, 1998)[1]. Figure 1.5 (taken from Orlando & Strong (2013)) shows the best fit GALPROP synchrotron and free-free model to data from radio surveys for the inner Galaxy (*left*) and high latitudes (*right*). The synchrotron spectrum obeys a power law over many decades in frequency with a low frequency turn over. In the inner Galaxy the model is in agreement with the data. At high latitude the model underestimates the synchrotron emission in polarization, which suggests that the magnetic field is not accurately modelled over the whole sky, and overestimates the total intensity emission between 10–100 GHz, which suggests that the free-free component is not well modelled here. The remainder of this section describes the physical origins of the synchrotron spectra plotted in Figure 1.5, first focusing on the power law behaviour before explaining the low frequency turn over and then also discussing a higher frequency cut-off.

[1] http://galprop.stanford.edu/

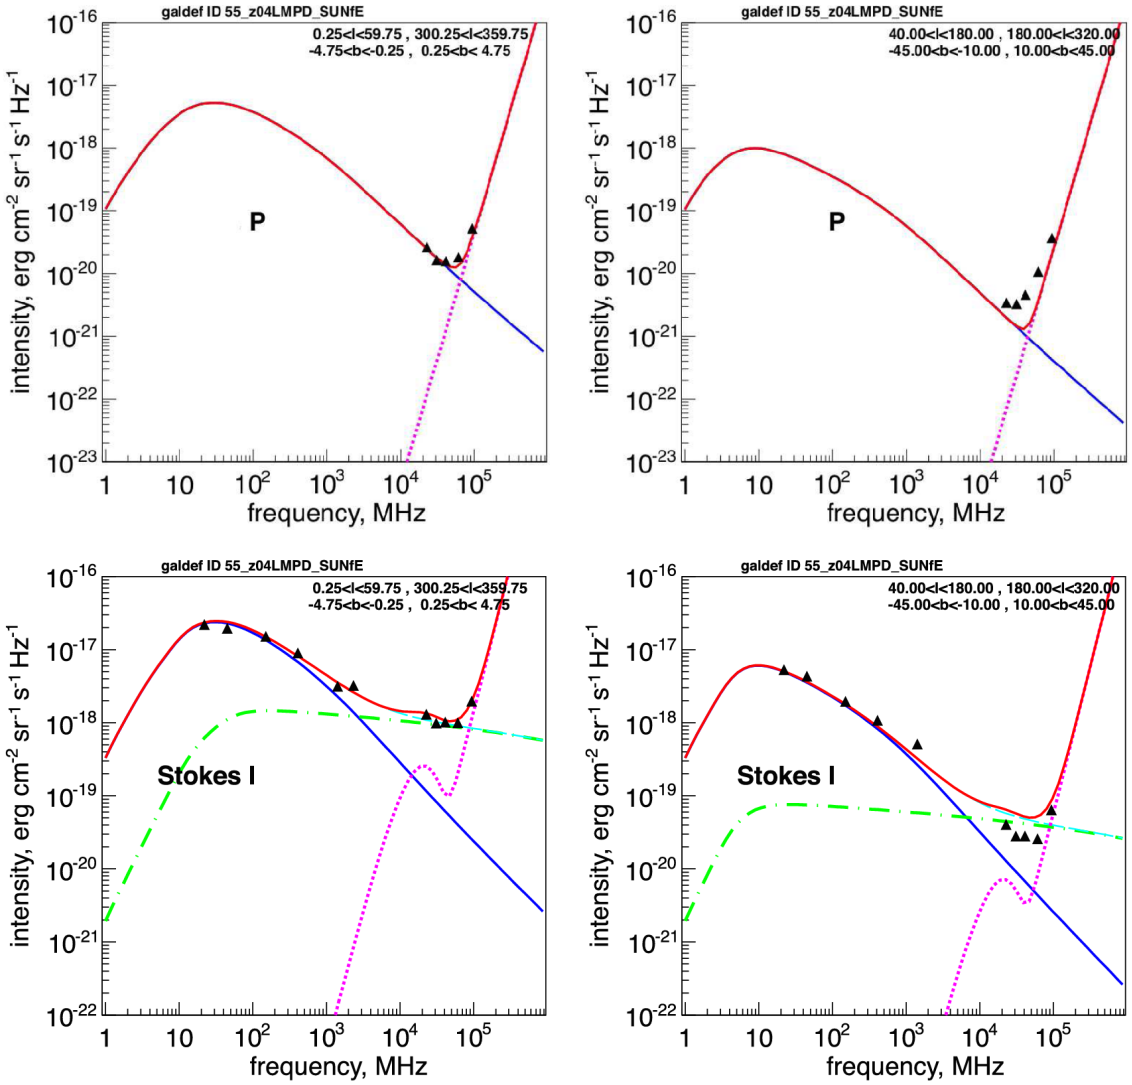

Figure 1.5: The best fit polarization, P, (*top*) and I (*bottom*) frequency spectra to radio surveys both in the inner Galaxy (*left*) and at high latitudes (*right*). The various components are a GALPROP synchrotron model (*blue solid line*), dust model (*pink dotted line*), free-free (*green dash-dotted line*) and free-free plus synchrotron (*cyan dashed line*). The total spectrum is the *red line* and the data are *dark triangles*. Figure reprinted from Orlando & Strong (2013), with permission from OUP. For permission to reuse, please contact the rights holder.

Consider a single synchrotron emitting electron with energy $E = \gamma m_e c^2$, where γ is the Lorentz factor and m_e is the electron rest mass. Its frequency spectrum can be approximated as a narrow spike at a frequency

$$\nu_{\rm crit} \simeq \gamma^2 \nu_{\rm cyc}, \tag{1.20}$$

where $\nu_{\rm cyc}$ is the cyclotron frequency. On average, an observer in an inertial reference frame would see the electron radiate power

$$-\frac{dE}{dt} = \frac{4}{3}\gamma^2 c \sigma_T U_B, \tag{1.21}$$

where σ_T is the Thomson cross-section and U_B is the magnetic field energy density. Galactic synchrotron emission comes from clouds of electrons with a range of energies. The electron energy distribution, $N(E)$, in the interstellar medium (ISM) is plotted in Figure 1.6 and obeys a power law over at least ~ 2 orders of magnitude. We can parametrize this distribution with an energy spectral index, p, like

$$N(E)dE \propto E^{-p}dE. \tag{1.22}$$

In our galaxy $p \simeq 2.4$ (Casadei & Bindi, 2004).

The specific brightness[2] of synchrotron emission from a cloud of electrons, $S_\nu(\nu)$, is given by

$$S_\nu(\nu)d\nu = -\frac{dE}{dt}N(E)dE. \tag{1.23}$$

From our knowledge of $-\frac{dE}{dt}$, $N(E)$ and $\frac{dE}{d\nu}$ we find that

$$S_\nu \propto \nu^{(1-p)/2} \simeq \nu^{-0.7}. \tag{1.24}$$

Expressing this as a brightness temperature ($\rm K_{RJ}$) we have $T_{\rm RJ} \propto \nu^{-2.7}$. Because the power radiated by an electron is proportional to E^2 and $N(E)$ obeys a power law over ~ 2 orders of magnitude in frequency, Galactic synchrotron emission should itself obey a power law over ~ 4 orders of magnitude in frequency (Condon & Ransom, 2016). Radio surveys of the sky between 38–1420 MHz and at *WMAP* frequencies show that the temperature spectral index of synchrotron emission varies between 2.5–3 across the sky, for example Lawson et al. (1987); Reich & Reich (1988); Platania et al. (2003); Davies et al. (2006); Guzmán et al. (2011).

The lower frequency cut off seen in Figure 1.5 is caused by synchrotron self-absorption. The brightness temperature of a cloud of electrons cannot be greater

[2]Specific brightness and commonly encountered unit conversions are explained in Appendix A.

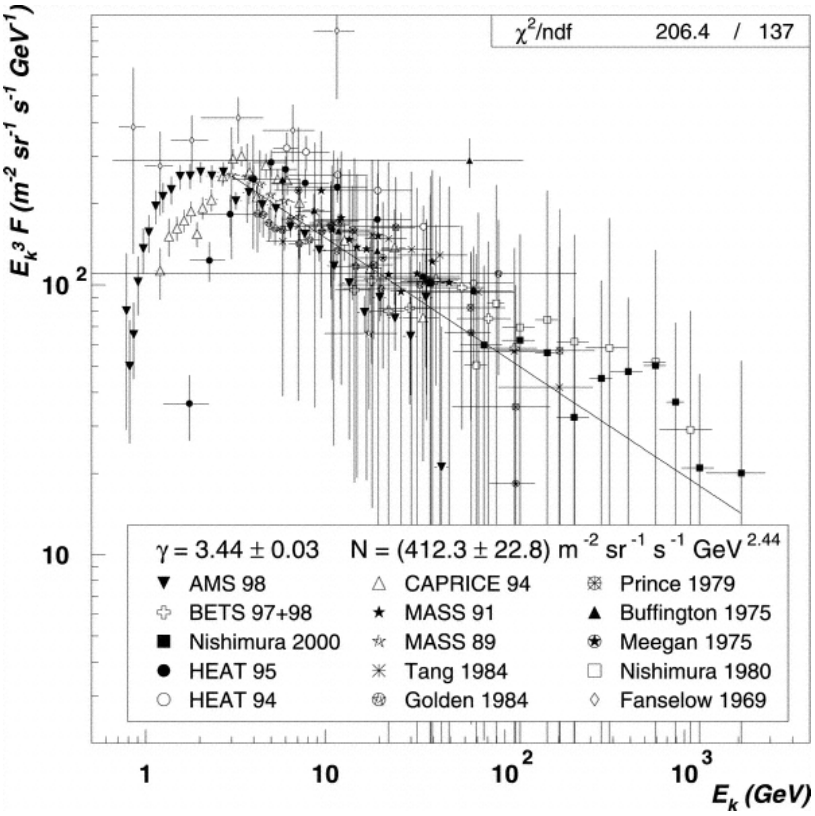

Figure 1.6: The distribution of local interstellar medium electron energies and the best fit power law. Figure reprinted from Casadei & Bindi (2004), reproduced by permission of the original authors.

than that cloud's thermodynamic temperature. From the power law spectrum, $T_{\rm RJ}$ increases with decreasing frequency. At low enough frequencies $T_{\rm RJ}$ approaches the mean kinetic temperature of the emitting electrons, $T_e = \frac{\gamma m_e c^2}{3k_b}$, and at this point self absorption becomes important. At these high temperatures the electrons can re-absorb the synchrotron radiation released by other electrons, and emission and absorption balance along the line of sight. This causes the spectrum to flatten off and eventually turn over. Below the turnover frequency the spectrum once again obeys a power law with $T_{\rm RJ} \sim T_e \propto \gamma \propto \nu^{0.5}$

At high frequencies, the spectrum of synchrotron emission from clouds of relativistic electrons will eventually cut-off. The high frequency cut-off is caused by the cloud of electrons ageing. From Equation 1.20, synchrotron emission at the highest frequencies is emitted by electrons with the greatest energies and, from Equation 1.21, the greater the energy of an electron then the faster it will radiate that energy away. Therefore, as a population of synchrotron emitting electrons ages, the number of high energy electrons will drop faster than the number of lower energy electrons. This

causes an exponential cut-off in the spectrum and the cut-off frequency drops as the population of electrons gets older.

1.2.2 Free-free emission

Free-free radiation (or bremsstrahlung) is released by electrons scattering off ions in warm interstellar plasma. Free-free emission can be the dominant component of diffuse Galactic emission in Stokes I between a few GHz and 100 GHz. Because of the uniform distribution of scattering angles between the electrons and ions in the plasma, free-free radiation is not highly polarized. Upper limits on the diffuse free-free polarization fraction are $< 3\%$ (Macellari et al., 2011).

The two-parameter model of Draine (2011) accurately describes the diffuse Galactic free-free emission at typical CMB experiment frequencies. The brightness temperature, $T_{\rm RJ}^{\rm ff}$, is given by

$$g_{\rm ff} = \log\left(\exp\left[5.960 - \sqrt{3}/\pi \log\left(\nu_9 T_4^{-3/2}\right)\right] + e\right) \tag{1.25}$$

$$\tau = 0.05468 T_{\rm e}^{-3/2} \nu_9^{-2} EM g_{\rm ff} \tag{1.26}$$

$$T_{\rm RJ}^{\rm ff} = 10^6 T_{\rm e}\left(1 - \exp^{-\tau}\right) \tag{1.27}$$

where T_e is the kinetic temperature of the electrons, EM is the effective emission measure, $T_4 = T_e/(10^4\,{\rm K})$, $\nu_9 = \nu/(10^9\,{\rm Hz})$, $g_{\rm ff}$ is the Gaunt factor and accounts for quantum effects (assuming ions have a charge of +1). The effective emission measure is the integrated squared electron density along the line of sight measured in cm^{-6}pc. In this parametrization and at the frequencies of typical CMB experiments, EM can be interpreted as the amplitude of the free-free emission and T_e weakly changes the spectral index.

The free-free emission frequency spectrum obeys a power law over a large range of frequencies with a typical spectrum of $T_{\rm RJ}(\nu) \sim \nu^{-2.1}$ that does not vary significantly across the sky (Dickinson, Davies & Davis, 2003; Gold et al., 2009; Draine, 2011). As the frequency increases, $g_{\rm ff}$ causes the spectrum to slowly steepen. Like synchrotron radiation, self absorption becomes important at lower frequencies. This causes the spectrum to flatten off below frequencies around 1–10 GHz.

Most diffuse Galactic free-free radiation is from warm hydrogen plasma ($\sim 10^4$ K) and Hα emission can be used as a tracer for both the plasma and of free-free emission. In the plasma, free electrons and protons can recombine to form atomic hydrogen with the electron in an excited state. As the electron cascades down to the ground state it radiates away its energy. Hα emission is the first line in the Balmer series and

is released when an electron transitions from the $n = 3$ to the $n = 2$ energy level. The Hα line is at a wavelength 656.28 nm. Like free-free emission, the intensity of Hα emission is proportional to EM. In dusty regions this optical emission will be absorbed, but if this absorption is accounted for then the corrected Hα maps can be used as a tracer for free-free emission (VallsGabaud, 1998). For an example see Finkbeiner (2003).

1.2.3 Thermal dust emission

In both total intensity and polarization the dominant foreground to the CMB above ~ 100 GHz is thermal dust emission from vibrating dust grains that have been heated by interstellar radiation.

Interstellar dust is created in stellar winds and the circumstellar ejecta of dying stars where the heavy elements in the gas can condense to solid grains. The dust grains can grow in the ISM through accretion and coagulation.[3] There are small grains of polycyclic aromatic hydrocarbons (PAHs) with diameters up to a nanometer, slightly larger grains are made mostly of carbon and the largest grains are made of silicates coated in carbon and can reach a micron in diameter (Partridge et al., 2010). Interstellar dust grains are typically heated to $T_\mathrm{d} = 10 - 30$ K but can reach up to 40 K and above in the most extreme environments (Draine, 2011).

Dust grains are not spherically symmetric and emit more efficiently along their long axis, furthermore the grains can be aligned by the Galactic magnetic field. This causes thermal dust emission to have a net polarization and Planck Collaboration et al. (2015c) have found that thermal dust emission can be up to 20% polarized.

Vibrating dust grains radiate as modified blackbodies, also called grey-bodies. One way of parametrizing a grey-body frequency spectrum is

$$T_\mathrm{RJ} = A_\mathrm{d} \left(\frac{\nu}{\nu_0}\right)^{\beta_\mathrm{d}+1} \frac{\exp(\frac{h\nu_0}{k_\mathrm{B}T_\mathrm{d}}) - 1}{\exp(\frac{h\nu}{k_\mathrm{B}T_\mathrm{d}}) - 1}, \tag{1.28}$$

where A_d is the dust brightness temperature at a reference frequency ν_0, β_d is the emissivity index and T_d is the physical temperature of the dust. It is important to note that this is a phenomenological model of the thermal dust emission.

The frequency spectrum of diffuse thermal dust emission between 10–1000 GHz can be well modelled by a single modified blackbody component (Planck Collaboration et al., 2014, 2016b), although there is some evidence to suggest that two independent components of grey-body thermal dust may fit the data better (Finkbeiner,

[3] Dust grains can be destroyed by radiation, grain-grain collisions and gas-grain collisions.

Davis & Schlegel, 1999; Meisner & Finkbeiner, 2014). At frequencies higher than 1000 GHz the grey-body model is too simplistic to match the measured spectrum (Planck Collaboration et al., 2016b).

1.2.4 Anomalous microwave emission (AME)

Since its initial discovery in the late 1990s, AME has been observed in total intensity across a frequency range $\sim 10 - 60$ GHz by many CMB experiments including: Python V (Mukherjee et al., 2003), ACME/SP94 (Hamilton & Ganga, 2001), Tenerife (Mukherjee et al., 2001) and QMAP (de Oliveira-Costa et al., 2000). Polarized diffuse AME is yet to be firmly detected and currently there are only upper limits of a polarization fraction $\sim 1\%$ over a frequency range 20–30 GHz (see for example, Rubiño-Martín et al., 2012).

AME was discovered by Kogut et al. (1996). They measured excess emission in low frequency COBE satellite data. The excess emission was spatially correlated with thermal dust but the signal was significantly stronger than the expected thermal dust emission at these frequencies. The excess emission was initially interpreted as free-free emission. Leitch et al. (1997) made deep observations of the north celestial pole (NCP) region and confirmed the presence of dust correlated emission at low frequencies but saw no corresponding signal in Hα maps. This meant that if free-free were the responsible emission mechanism then the plasma would have to be very hot ($T_e > 10^6$ K). Draine & Lazarian (1998) realised that this was unphysical and proposed that the AME was caused by electric dipole emission from small spinning dust grains. The following year Draine & Lazarian (1999) proposed a second emission mechanism, magnetic dust, which was also consistent with the observations.

Spinning dust models are currently the most well developed theories of AME. The rest of this section contains a brief overview of the spinning dust spectrum and an even more brief description of magnetic dust.

Small PAH dust grains in the ISM can develop and maintain electric dipole moments. A classically rotating dust grain with electric dipole moment $\mu(\mathbf{t})$ radiates power at its rotation frequency according to the Larmor formula,

$$P = \frac{2}{3c^3}\ddot{\mu}. \tag{1.29}$$

The frequency spectrum of spinning dust emission from a cloud of dust grains therefore depends on the distribution of angular velocities and dipole moments in the cloud. The theoretical modelling of the physical processes that affect the dipole moments

and angular velocities of small dust grains is well developed but more detailed observations are needed to help distinguish between the many models that are consistent with current data (Ali-Haïmoud, 2013).

Emission from a single rotating dipole is 100% polarized and, therefore, diffuse Galactic spinning dust emission could be polarized if there is a mechanism maintaining the alignment of the grains. Serkowski's law (Serkowski, Mathewson & Ford, 1975), however, suggests that the smallest grains will not be aligned and so spinning dust emission is only expected to be weakly polarized ($< 1\%$). This is consistent with observations.

SPDUST2 [4] is an IDL program that calculates the frequency spectrum of spinning dust given a range of environmental parameters. It implements the models described in Silsbee, Ali-Haïmoud & Hirata (2011) and Ali-Haïmoud, Hirata & Dickinson (2009). The emissivity spectra calculated for a range of environments are shown in the top panel of Figure 1.7 (taken from Bennett et al. (2013)). The bottom panel shows the same spectra all scaled to the same peak frequency and amplitude. The scaled plots are fairly consistent with one another and demonstrate that the spinning dust emission can be empirically modelled by a template scaled in amplitude and frequency (Bennett et al., 2013; Planck Collaboration et al., 2016b).

Draine & Lazarian (1999) have proposed a second mechanism that could contribute up to half the observed AME, emission from magnetic dust. Thermal fluctuations of magnetic dust grains could excite the magnetisation of a grain away from its lowest energy direction. As the grain relaxes back to its ground state it will emit radiation. Free-flying magnetised grains could produce 20% polarized emission, however if the grains are embedded in larger non-metallic grains then the polarization properties are more complex (Draine & Hensley, 2013).

1.2.5 Other foregrounds

Line emission Rotational transitions in small molecules such as CO and HCN are significant at very narrow frequency ranges, particularly in star forming regions.

The rotational energy eigenvalues of the Schrödinger equation for a diatomic molecule are $E = \frac{J(J+1)\hbar^2}{2I}$, where I is the moment of inertia. Allowed rotational transitions have a change in quantum number $\Delta J = \pm 1$ and therefore release a photon with energy $\Delta E = \frac{\hbar^2 J}{I}$. The frequencies of the first few rotational transitions of HCN and CO molecules are shown in Table 1.2.

[4] http://pages.jh.edu/~yalihai1/spdust/spdust.html

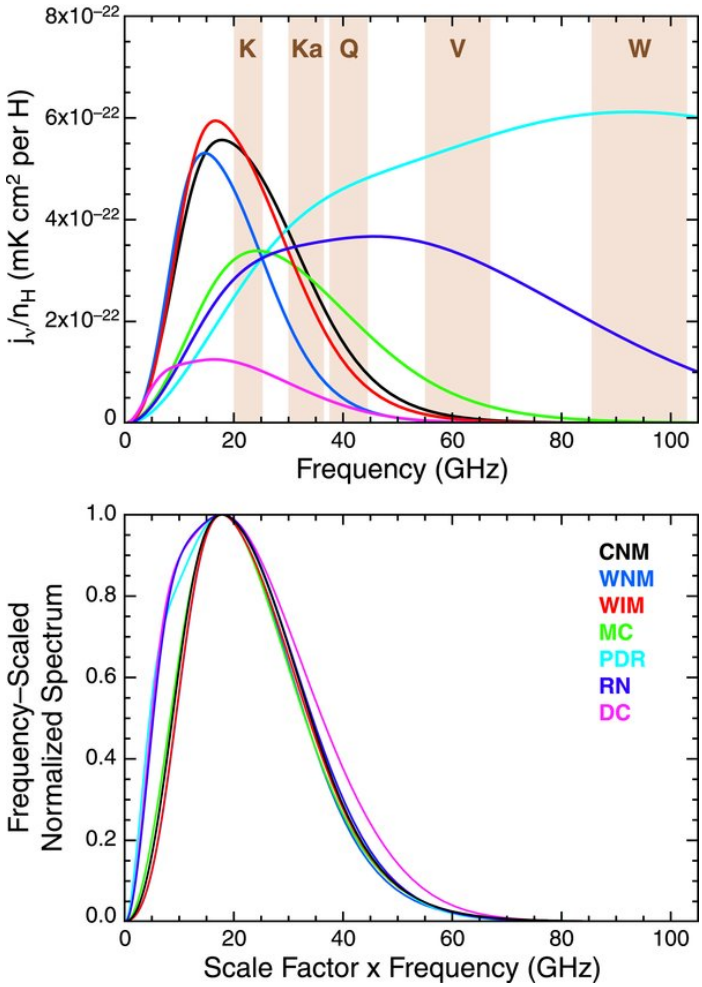

Figure 1.7: The emissivity spectra produced by SPDUST2 and presented in Silsbee, Ali-Haïmoud & Hirata (2011) for a range of ISM environments (*top plot*) and then those spectra scaled to the same peak frequency and amplitude (*bottom plot*). Figure reprinted from Bennett et al. (2013), reproduced by permission of the original authors.

Table 1.2: Rest frame frequencies of the first few rotational transitions of HCN and CO molecules. Data taken from Splatalogue (`http://www.cv.nrao.edu/php/splat/`).

Species	Frequency [GHz]	Transition
HCN	88.6	$J = 1 - 0$
CO	115.3	$J = 1 - 0$
HCN	177.3	$J = 2 - 1$
CO	230.5	$J = 2 - 1$
HCN	265.9	$J = 3 - 2$
CO	345.8	$J = 3 - 2$
HCN	354.5	$J = 4 - 3$
HCN	443.1	$J = 5 - 4$

The *Planck* maps at 100, 217 and 353 GHz and the *WMAP* W-band maps are strongly contaminated by line emission (Planck Collaboration et al., 2013a). To quote Planck Collaboration et al. (2011):

> After launch it became apparent that the contribution of CO rotational transitions to the HFI [High frequency instrument] measurements was greater than anticipated, especially for the 100 GHz band.

Line emission can be fitted using knowledge of the different bandpasses of individual *WMAP* and *Planck* detectors (Planck Collaboration et al., 2013a).

Extragalactic foregrounds Other galaxies also radiate synchrotron, free-free, thermal dust and anomalous microwave emission. Depending on whether or not a galaxy is active it can also emit synchrotron radiation with a range of spectral indices from jets powered by its central black hole.

The brightest extragalactic objects can be individually identified in maps. The fainter sea of unresolved point sources adds an extra level of temperature fluctuations to a map called confusion noise. The standard deviation of the confusion noise is called the confusion limit. The confusion limit is much lower in polarization than in total intensity because most extragalactic radio sources are not significantly polarized. Battye et al. (2011) analysed the *WMAP* point source catalogue to find extragalactic radio sources between 8.4–43 GHz have a frequency independent mean polarization fraction of 3.5%. Extrapolating their results to higher frequencies they show that to detect a tensor-to-scalar ratio of 0.01 will require a method of mitigating against polarized radio sources.

Extragalactic radio sources become an increasingly important foreground to the CMB on smaller angular scales - or equivalently at large ℓ. For a given frequency it can be shown that the angular power spectrum of unresolved point sources is given by

$$C_\ell = \int_0^{\phi_c} \frac{\partial \bar{n}}{\partial \phi} \phi^2 d\phi, \tag{1.30}$$

where ϕ is the flux, $\bar{n}(\phi)$ is the number density per steradian of point sources with flux less than ϕ and $\frac{\partial \bar{n}}{\partial \phi}$ is the differential source count (Tegmark & Efstathiou, 1996). The angular power spectrum of point source emission has no dependence on ℓ and C_ℓ is flat. For comparison the spectrum of the CMB does depend on ℓ and the perturbations seeded by inflation are scale invariant and $C_\ell \propto \ell^{-2}$.

Galaxy clusters provide a further foreground through the Sunyaev-Zel'dovich (SZ) effect, caused by CMB photons inverse Compton scattering off hot intergalactic electrons (Sunyaev & Zeldovich, 1972). The scattering increases the CMB photon energy and shifts the blackbody spectrum of scattered CMB photons to a higher frequency such that

$$\frac{\delta I_\nu^{\rm CMB}}{I_\nu^{\rm CMB}} = -2 \int n_e \sigma_T \frac{k_B T_e}{m_e c^2} dl, \tag{1.31}$$

where n_e is the number density of electrons and the integral is along the line of sight. The SZ effect decreases the observed CMB temperature at frequencies below 217 GHz and increases it at frequencies above 217 GHz. Due to their proximity, the Coma and Virgo clusters are the most dominant SZ objects in the sky.

Zodiacal emission Our solar-system is full of inter-planetary meteoroids (IPM) with diameters $\sim 10\,\mu$m that are heated by the sun to temperatures around ~ 270 K (Peucker-Ehrenbrink & Schmitz, 2001). These grains are much larger and warmer than the dust typically found in the ISM. The IPMs are concentrated in the ecliptic plane between the Sun and Jupiter. They reflect light from the sun and from dark locations can be seen with the naked eye at dawn and dusk as a triangular, diffuse, white glow that extends along the ecliptic plane.

As well as reflecting light from the sun, IPMs emit thermal radiation at infra-red frequencies and are a significant contributor to the diffuse sky at wavelengths between $\sim 1 - 100\,\mu$m. At mid infra-red frequencies, zodiacal emission dominates most of the sky (Kelsall et al., 1998).

Modelling the zodiacal emission is complex because its contribution to the brightness at a given RA/DEC depends on the time of year at which the observation was made. The COBE/DIRBE team constructed a full sky parametric model of this emission (Kelsall et al., 1998). They integrate a source function and 3D density distribution along the line of sight and include scattering terms to calculate the brightness. Many others have created models of the Zodiacal light, including Vrtilek & Hauser (1995); Wright (1998); Rowan-Robinson & May (2013).

The *Planck* Collaboration attempt to remove this emission from their high frequency instrument maps during the mapmaking stage using the Kelsall et al. (1998) model (Planck Collaboration et al., 2013c, 2015b). Planck Collaboration et al. (2016b) extrapolate the 100 GHz template with a power law to correct the LFI maps. This emission is negligible at the lowest frequencies relevant to CMB research.

1.3 Foreground separation techniques

This section outlines four of the techniques most commonly used to separate diffuse foregrounds from the CMB using maps of the sky at different frequencies. The four techniques are the internal linear combination (Section 1.3.1), template fitting (Section 1.3.2), independent component analysis (Section 1.3.3) and bayesian parametric model fitting (Section 1.3.4).

The internal linear combination (ILC) method, template fitting and independent component analysis (ICA) estimate the CMB from linear combinations of maps but they make different assumptions about the maps used. For example, the maps used in ILC analysis must all be accurately calibrated to $\mathrm{K_{CMB}}$ units (Apendix A), whereas in template fitting the maps can have arbitrary gain factors applied to them. The methods are treated separately in the literature, for examples see Dunkley et al. (2009); Planck Collaboration et al. (2014); Ichiki (2014).

1.3.1 Internal linear combination

The internal linear combination (ILC) technique is a blind component separation method that assumes that the CMB frequency spectrum is flat and that we are uninterested in all other foreground components. The ILC method can be used in pixel space, spherical harmonic space or needlet space.

A sky map can be represented as a 1D vector where each index is a pixel in the sky. Assume that a sky map vector $\mathbf{T}^i$ at frequency i is made of a linear sum of a CMB map $\mathbf{T}^i_{\mathrm{CMB}}$, a foreground map $\mathbf{T}^i_{\mathrm{f}}$ and a noise map $\mathbf{n}^i$ such that

$$\mathbf{T}^i = \mathbf{T}^i_{\mathrm{CMB}} + \mathbf{T}^i_{\mathrm{f}} + \mathbf{n}^i. \tag{1.32}$$

An estimator of the CMB can be constructed of the form

$$\hat{\mathbf{T}}_{\mathrm{CMB}} = \sum_i w^i \mathbf{T}^i \tag{1.33}$$

$$= \sum_i \left(w^i \mathbf{T}^i_{\mathrm{CMB}} + w^i \mathbf{T}^i_{\mathrm{f}} + w^i \mathbf{n}^i \right), \tag{1.34}$$

where w^i are weights and the sum is over the frequencies. Constraining the weights such that $\sum_i w^i = 1$ simplifies Equation 1.34 to

$$\hat{\mathbf{T}}_{\mathrm{CMB}} = \mathbf{T}_{\mathrm{CMB}} + \sum_i \left(w^i \mathbf{T}^i_{\mathrm{f}} + w^i \mathbf{n}^i \right). \tag{1.35}$$

The weights should be chosen to minimise the impact of the foregrounds and noise on the estimate of the CMB, $\hat{\mathbf{T}}_{\mathrm{CMB}}$. The simplest measure of this impact is the

variance of $\hat{\mathbf{T}}_{\rm CMB}$. The minimisation can be performed globally or separately over smaller regions of the sky to account for spatial changes in the foreground spectra. For example Bennett et al. (2003) break the sky into 12 regions.

The method can be improved by allowing different weights for different angular scales and performing the minimisation in spherical harmonic space instead of pixel space. On large angular scales the sky is dominated by foreground emission and on small angular scales maps are often noise dominated. The ILC was first used in spherical harmonic space by Tegmark & Efstathiou (1996). The CMB $a_{\ell m}$ can be estimated from the observed $a^i_{\ell m}$ at multiple frequencies by choosing weights that minimise the power spectrum,

$$\langle |a_{\ell m}|^2 \rangle = {\mathbf{w}_\ell}^T \boldsymbol{C_\ell} \mathbf{w}_\ell, \tag{1.36}$$

under the condition that for each ℓ the weights sum to unity, $\mathbf{e} \cdot \mathbf{w}_\ell = 1$, where $\boldsymbol{C}^{ij}_\ell = \langle (a^i_{\ell m})^\dagger a^j_{\ell m} \rangle$ and $\mathbf{e}$ is the all ones column vector. This gives

$$\mathbf{w}_\ell = \frac{\boldsymbol{C}_\ell^{-1}\mathbf{e}}{\mathbf{e}^T \boldsymbol{C}_\ell^{-1}\mathbf{e}}. \tag{1.37}$$

A further improvement can be achieved by using the ILC method on a basis of needlets. Needlets are a special type of spherical wavelet that are quasi-exponentially localised in pixel space and exactly localised in a finite number of multipoles in spherical harmonic space (Narcowich, Petrushev & Ward, 2006). This localisation allows the weights to change smoothly across the sky and with angular scale. The technique was first used by Delabrouille et al. (2009) to produce a CMB map from the *WMAP* 5-year temperature data. The Planck Collaboration use the NILC pipeline (Planck Collaboration et al., 2016b).

The key advantage of a blind separation technique such as ILC is that no assumptions have been made about the frequency spectra of the foreground components. The ILC method has a second advantage of being very fast to run. The main disadvantages of the ILC method are that it does not recover any information about any other components of the sky, the underlying uncertainty of the estimated CMB map is difficult to assess and care has to be taken to break the sky into separate regions across which the spectral properties of the foregrounds are believed to be the same (Vio & Andreani, 2008). Recently this method has been used to produce a CMB map by Planck Collaboration et al. (2014).

1.3.2 Template fitting

The template fitting technique, or cross-correlation analysis, assumes that a sky map $\mathbf{T}$ at some frequency is made of a linear superposition of spatial templates $\mathbf{t}_j$ and noise $\mathbf{n}$. The noise can include both instrumental contributions and the CMB. $\mathbf{T}$ is modelled using

$$\mathbf{T} = \sum_j \alpha_j \mathbf{t}_j + \mathbf{n}, \tag{1.38}$$

where α_j is the cross-correlation coefficient that scales the j^{th} template to the frequency of the map and the sum is over the templates.

An estimator of the CMB map can then be constructed from the sky map and the templates,

$$\hat{\mathbf{T}}_{\text{CMB}} = \mathbf{T} - \sum_j \alpha_j \mathbf{t}_j \tag{1.39}$$

with the values of α_j found by minimising a χ^2,

$$\chi^2 = (\mathbf{T} - \sum_j \alpha_j \mathbf{t}_j)^T \boldsymbol{C}^{-1} (\mathbf{T} - \sum_j \alpha_j \mathbf{t}_j). \tag{1.40}$$

The covariance matrix, $\boldsymbol{C}$, encapsulates the noise properties of both the instrument and the CMB, where $\boldsymbol{C}_{\text{CMB}}$ can be estimated from a CMB power spectrum.

Because the templates are fitted across the whole sky, components can be constrained even if their signal-to-noise ratio in individual pixels is low. This in turn leads to the major disadvantage of this method. It assumes that the spectral properties of each component are uniform across the sky, which may not be the case. Template fitting requires a template for each foreground component and noise from these templates can be aliased into the estimated CMB map. This is problematic if the noise on the template map is large. Template fitting can be used in either pixel space or spherical harmonic space, although more care has to be taken in spherical harmonic space if there is incomplete sky coverage (Mortlock, Challinor & Hobson, 2002; Jaffe et al., 2006).

Template fitting is well known in the literature and widely used, for example it was used to remove foregrounds from the COBE-DMR 4-year maps (Banday et al., 1996) and *WMAP* 5-year maps (Komatsu et al., 2009). SEVEM is an implementation of template fitting (Martinez-Gonzalez et al., 2003) that has been used on *WMAP* polarization data (Fernández-Cobos et al., 2012) and tested on simulated *Planck* data (Leach et al., 2008).

1.3.3 Independent component analysis

Independent component analysis (ICA) techniques assume that a set of sky maps at a range of frequencies are each linear sums of the CMB and the foregrounds. The CMB has a Gaussian distribution but the foregrounds are highly non-Gaussian. There are many implementations of ICA including ALTICA and SMICA.

ALTICA (Maino et al., 2002b) is an implementation of the FASTICA algorithm (Hyvarinen, 1999). Due to the central limit theorem, a linear sum of random variables is more Gaussian than the un-summed variables. The components can be separated by making linear combinations of the input maps that maximise the non-Gaussianity. The non-Gaussianity is measured by the neg-entropy statistic. The FASTICA algorithm has been used on data from *COBE*, (Maino et al., 2003), BEAST (Donzelli et al., 2006), *WMAP* (Maino et al., 2007) and simulated *Planck* data (Ichiki et al., 2013).

Spectral Matching ICA (SMICA) is an implementation of ICA in spherical harmonic space (Delabrouille, Cardoso & Patanchon, 2003; Cardoso et al., 2008). A parametric model of the sky is fit to the auto- and cross-spectra of n_ν input maps. The model is then

$$C_\ell^{ij} = C_\ell^{\mathrm{CMB},ij} + F_\ell^{ij} + N_\ell^{ij}, \tag{1.41}$$

where C_ℓ^{ij} are the auto- or cross-spectra of the maps i and j, $C_\ell^{\mathrm{CMB},ij}$ are the CMB spectra, F_ℓ^{ij} are the foreground spectra and N_ℓ^{ij} are the noise terms. Often the spectra are binned into n_q ℓ-bins to give C_q^{ij}, where the bin width is Δ_q.

At each ℓ bin the foregrounds are parametrized by a symmetric amplitude matrix (with size $n_{\mathrm{fg}} \times n_{\mathrm{fg}}$) and a mixing matrix (with size $n_\nu \times n_{\mathrm{fg}}$). There are therefore $n_\nu n_{\mathrm{fg}} + n_q n_{\mathrm{fg}}(n_{\mathrm{fg}}+1)/2$ foreground parameters. This model captures the correlation between foreground components. It is often reasonable to treat the input sky maps as independent, in which case the noise matrix N_ℓ^{ij} is diagonal. If unknown, the diagonal elements of this matrix can also be treated as free parameters.

The auto- and cross-spectra are known and form a symmetric $n_\nu \times n_\nu$ matrix for each ℓ bin giving $n_q n_\nu(n_\nu+1)/2$ known parameters.

Assuming that the spherical harmonic coefficients are independent normally-distributed random variables with a variance proportional to the power spectrum, the unknown parameters (θ) can be found by minimising the negative log-likelihood,

$$\phi(\theta) = \sum_{q=1}^{n_q} \Delta_q \left[\mathrm{trace}(C_q(\theta)^{-1}\hat{C}_q) - \log\det(C_q(\theta)^{-1}\hat{C}_q) - n_\nu \right]. \tag{1.42}$$

This method has been used on Archeops data (Tristram et al., 2005a), *WMAP* data (Patanchon et al., 2005) and both simulated and real *Planck* data (Leach et al., 2008; Planck Collaboration et al., 2015a). Planck Collaboration et al. (2014) tested four methods of separating the CMB from its foregrounds on simulated temperature data. They compared NILC, SEVEM, Commander (described in the following section) and SMICA. Of the four methods, SMICA produced maps with the smallest residuals.

1.3.4 Bayesian parametric model fitting

Bayesian methods can be used to simultaneously fit a model of the sky, parametrized by θ, to multiple sky maps, $\mathbf{d}$, by maximising the posterior distribution

$$p(\theta|\mathbf{d}) = \frac{p(\mathbf{d}|\theta)P(\theta)}{p(\mathbf{d})} \propto \mathcal{L}(\theta)p(\theta), \tag{1.43}$$

where $\mathcal{L}(\theta)$ is the likelihood and $p(\theta)$ the prior on θ. The sky model can include the CMB, foregrounds and instrumental effects. The sky models typically contain many free parameters, e.g. amplitudes and spectral properties for each foreground component per pixel. Markov Chain Monte Carlo (MCMC) methods, described in Appendix B, are typically employed to explore the multidimensional and highly correlated parameter spaces. Bayesian methods have the advantage of robustly propagating errors in the maps to the estimates of the CMB and foreground components. As a disadvantage, Bayesian parametric model fitting is computationally very expensive.

Many examples of parametric fitting codes can be found in the literature. Commander (Eriksen et al., 2004, 2008) is the Gibbs sampling code used by *Planck* to perform pixel based parametric fitting and in Planck Collaboration et al. (2016b) it was used to produce all-sky CMB and foreground component maps. They fitted both a sky model and instrumental parameters to the *Planck*, *WMAP* and Haslam maps. In total intensity their sky model included synchrotron, free-free, spinning dust, thermal dust, SZ, line emission and CMB components. In polarization their sky model included only synchrotron, thermal dust and CMB components. There were not enough low-frequency data points to include the synchrotron spectral index as a free parameter and instead the synchrotron spectrum was determined by a frequency-shifted GALPROP template (Planck Collaboration et al., 2015d; Orlando & Strong, 2013). More low-frequency maps are required to gain extra degrees of freedom in the fits and allow the synchrotron spectral parameters to vary in the model.

Extra low-frequency surveys should also help reduce the large degeneracies that currently exist between the synchrotron, free-free and AME components. COMMANDER 2 is currently in development and is designed to make best use of developments in linear solver technology, parallelism and sky/instrument modelling[5]. Other parametric fitting packages include FGFIT (Eriksen et al., 2006) and MIRAMARE (Stompor et al., 2009).

1.4 Observations

1.4.1 CMB temperature fluctuations and E-mode power spectrum

The CMB was discovered in 1964 by Arno Penzias and Bob Wilson (Penzias & Wilson, 1965). They detected an isotropic 3.5 K temperature excess and the interpretation of this signal as the CMB was made in a companion paper (Dicke et al., 1965). Woody & Richards (1981) used a balloon-borne spectrophotometer to measure the CMB frequency spectrum. They found it was qualitatively consistent with a 3 K blackbody but their statistical analysis was limited by systematic effects. The COBE satellite was the first experiment to confirm the blackbody nature of the CMB (Mather et al., 1990) and also made the first statistical detection of its temperature fluctuations (Smoot et al., 1992). Subsequent ground and balloon experiments such as VSA (Watson et al., 2003; Taylor et al., 2003; Scott et al., 2003), BOOMERANG (de Bernardis et al., 1999, 2000), MAXIMA (Lee et al., 1999; Hanany et al., 2000), CBI (Padin et al., 2002, 2001), DASI (Halverson et al., 1998; Kovac et al., 2002) and Archeops (Benot et al., 2002; Benoît et al., 2004; Tristram et al., 2005b) measured the main features of the CMB TT, TE and EE spectra including the acoustic peaks and Silk damping.

The *WMAP* satellite, launched in 2001, measured the CMB power spectrum from $\ell = 2$ to $\ell = 1000$ at a sensitivity greater than previously managed by other experiments and was the first experiment to measure the spectrum across such a large range of angular scales simultaneously (Hinshaw et al., 2013; Bennett et al., 2013). More recently the *Planck* satellite was launched in 2009. *Planck* has mapped the microwave sky and measured the CMB TT, TE and EE power spectra with unprecedented accuracy (Planck Collaboration et al., 2016a). The *Planck* best fit ΛCDM spectrum to the CMB TT power spectrum is shown in Figure 1.8 (*top*). The corresponding

[5] http://commander.bitbucket.org/

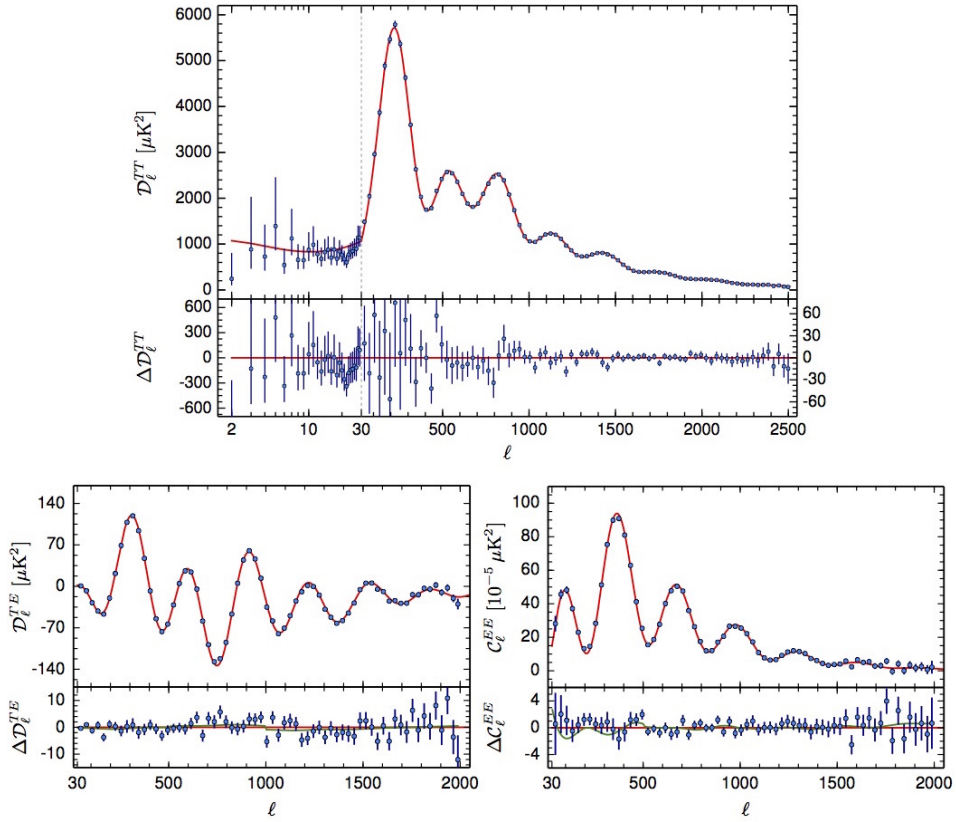

Figure 1.8: The TT (*top*), TE (*bottom left*) and EE (*bottom right*) angular power spectra of the CMB from Planck Collaboration (2015). The points are estimates of the CMB power spectrum and the red line is the best fit ΛCDM spectrum to the TT data. The bottom panels in each plot are the residuals to the fit. Figure reprinted from Planck Collaboration, A&A, 594, A13, 2016, reproduced with permission ©ESO.

TE (*bottom left*) and EE (*bottom right*) spectra calculated from the best fit TT cosmological parameters are plotted over the polarization data along with the residuals. The TE and EE residuals are larger than expected from theory, suggesting there are systematics in the data, possibly temperature-to-polarization leakage. In combination with other datasets the Planck Collaboration have produced the tightest ever constraints on the parameters of the standard cosmological model and find no significant evidence for failure of ΛCDM (Planck Collaboration et al., 2016a).

Ground based telescopes such as ACT (Swetz et al., 2011) and SPT (Shirokoff et al., 2009) have measured the CMB temperature anisotropies on very small angular scales. Both telescopes have measured the SZ effect from galaxy clusters (Plagge et al., 2010; Fowler et al., 2010).

1.4.2 CMB B-mode power spectrum

Table 1.3 lists constraints on r found by various experiments.[6] The current upper limits on r cannot be improved with more temperature data alone due to confusion from density perturbations (Knox, 1995), and more precise polarization data are needed. There are many currently ongoing and proposed experiments trying to measure the CMB B-mode power spectrum from the ground, balloons and space.

The first direct measurement of the (non-primordial) CMB B-mode polarization was made by the POLARBEAR telescope (The POLARBEAR Collaboration et al., 2014). Their measurements covered the multipole range $500 < \ell < 2100$, which is the part of the spectrum dominated by lensed B-modes. POLARBEAR2 and the Simons Array are currently under development and aim to constrain r to $\sigma_r = 6 \times 10^{-3}$ (Suzuki et al., 2016).

BICEP1 was the first ground-based telescope designed specifically to measure the B-mode signal left by cosmic inflation (Keating et al., 2003; Barkats et al., 2014). BICEP1 was followed by an upgraded BICEP2 (Nguyen et al., 2008; Ade et al., 2014a), the Keck Array (Sheehy et al., 2010; Ade et al., 2016) and now BICEP3 (Ahmed et al., 2014). The anomalously high value of r found by Ade et al. (2014b) from BICEP2 data was a result of mis-evaluating the contribution from polarized foregrounds, in particular thermal dust. Once polarized foreground emission is included in the analysis r is no longer detected. Subsequent analysis has shown the BICEP2 data are fully consistent with thermal dust emission and $r < 0.12$ (Ade et al., 2015; Mortonson & Seljak, 2014; Flauger, Hill & Spergel, 2014).

ACTpol (Thornton et al., 2016) and SPTpol (Austermann et al., 2012) are the polarization sensitive upgrades to ACT and SPT respectively and have measured the CMB and polarization anisotropies on very small angular scales. Current balloon borne experiments measuring the CMB B-mode signal include EBEX (Oxley et al., 2004; Reichborn-Kjennerud et al., 2010), SPIDER (Filippini et al., 2010; Fraisse et al., 2013) and LSPE (The LSPE collaboration et al., 2012). Several space based telescopes have also been proposed such as *LiteBIRD* (Matsumura et al., 2014), *COrE* (The COrE Collaboration et al., 2011) and *PIXIE* (Kogut et al., 2011).

1.4.3 Ancillary surveys

To determine the CMB power spectrum down to low ℓ values requires large-area sky maps. There are no large areas of the sky with sufficiently low levels of polarized

[6] It is important to note that r is highly correlated with other parameters.

Table 1.3: Constraints on the tensor-to-scalar ratio from different datasets, see the references for details.

Dataset	Reference	r	Comments
Planck TT, lowP	Planck Collaboration (2015)	< 0.10 (95% CL)	
Planck TT, lowP, lensing, ext.	Planck Collaboration (2015)	< 0.11 (95% CL)	
Planck TT, lowP, BKP	Planck Collaboration (2015)	< 0.08 (95% CL)	
Planck TT, lowP, lensing, ext., BKP	Planck Collaboration (2015)	< 0.09 (95% CL)	
BICEP2, Keck, *Planck*	Ade et al. (2015)	< 0.12 (95% CL)	
BICEP2	Ade et al. (2014b)	$0.2^{+0.07}_{-0.05}$	Incorrect foreground modelling
BICEP1	Barkats et al. (2014)	< 0.70 (95% CL)	
WMAP, BAO, H_0, eCMB	Hinshaw et al. (2013)	< 0.13 (95% CL)	
*WMAP*7, BAO, H_0, SPT	Story et al. (2013)	< 0.11 (95% CL)	

Table 1.4: Overview of existing and on-going sky surveys at frequencies between 400 MHz and 1 THz that cover at least half the sky and have angular resolutions $\lesssim 1°$. Table from Jones et al. (2017, in prep.).

Survey/ Telescope	Frequency [GHz]	FWHM [arcmin]	Declination Coverage	Stokes[a]	Sensitivity[b] noise	offsets	Status[c]	Reference(s)
Haslam (various)	0.408	51	All-sky	I	1 K	3 K	3	Haslam et al. (1982)
Stockert[d]	1.42	35	$> -30°$	I	9 mK	50 mK	3	Reich & Reich (1986)
DRAO (26-m)[d]	1.4	36	$> -29°$	QU	12 mK	30 mK	3	Wolleben et al. (2006)
Villa Elisa[d]	1.4	35.4	$< +10°$	IQU	9 mK	50 mK	3	Testori, Reich & Reich (2008)
CHIPASS (Parkes)	1.394	14.4	$< +25°$	I	0.6 mK	30 mK	3	Calabretta, Staveley-Smith & Barnes (2014)
GMIMS-HB N	$1.28 \rightarrow 1.75$	30	$> -30°$	IQU	12 mK	?	1	Wolleben et al. (2010)
STAPS (Parkes)	$1.3 \rightarrow 1.8$	15	$< 0°$	IQU	?	?	1	Haverkorn (priv. comm.)
HartRAO	2.326	20	$-83°$ to $+13°$	$I-Q$	25 mK	80 mK	3	Jonas, Baart & Nicolson (1998)
S-PASS (Parkes)	2.3	≈ 9	$< 0°$	IQU	0.1 mK	unknown	1	Carretti (2010)
C-BASS	$4.5 \rightarrow 5.5$	45	All-sky	IQU	0.1 mK	1 mK	0	Jones et al. (2017, in prep.)
QUIJOTE	$11 \rightarrow 19$	≈ 60	$\gtrsim 0°$	$[I]QU$	25 μK	unknown	1	Génova-Santos et al. (2015b)
WMAP	$22.8 \rightarrow 94$	$49 \rightarrow 15$	All-sky	IQU	4 μK	1 μK	3	Bennett et al. (2013)
Planck LFI	$28.4 \rightarrow 70$	$32 \rightarrow 13$	All-sky	IQU	3 μK	1 μK	2	Planck Collaboration et al. (2016a)
Planck HFI	$100 \rightarrow 353$	$10 \rightarrow 5$	All-sky	IQU	$0.2 \rightarrow 0.5$ μK	1–5 μK	2	Planck Collaboration et al. (2016a)
Planck HFI	545, 857	5	All-sky	I	0.4, 0.8 μK	1 μK	2	Planck Collaboration et al. (2016a)

[a] $[I]QU$ denotes surveys where Stokes I is measured but with much larger systematic errors than Stokes Q and U. $I-Q$ denotes a single linear polarization.
[b] Approximate average Stokes I sensitivity in K_{RJ} 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, Reich & Testori, 2004; Testori, Reich & Reich, 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.

foreground emission that foregrounds can be neglected. Remazeilles et al. (2016) show that foreground modelling errors can introduce significant bias in estimates of r. This, along with the experience of BICEP2, demonstrates that a more complete picture of polarized foregrounds is urgently needed.

To separate the CMB from its foregrounds requires maps at a range of different frequencies. Current and ongoing large-area sky surveys at frequencies relevant to CMB research (400 MHz to 1 THz) with resolutions $\lesssim 1°$ are summarised in Table 1.4 (adapted from Jones et al. (2017, in prep.)). Most of the presently available ground-based surveys suffer from low sensitivity and large scan synchronous artefacts, which limit their use in component separation.

In polarization, frequencies below a few GHz suffer from significant Faraday rotation and depolarization which complicates sky modelling and, again, limits their use. We now briefly describe Faraday rotation. A linearly polarized electromagnetic wave

Table 1.5: Typical rotation angles for diffuse emission at high and low Galactic latitude, calculated by Jones et al. (2017, in prep.) using the catalogue of rotation measures of extragalactic radio sources by Taylor, Stil & Sunstrum (2009).

	$\theta_{\rm R}$	
Frequency [GHz]	High latitude	Low latitude
1.4	37°	112°
2.3	14°	42°
5.0	3°	9°

can be considered as the superposition of two circularly polarized waves of equal amplitude but opposite handedness. A plasma in a magnetic field is birefringent and the two circular polarizations will have different phase velocities. This causes the polarization angle of the linearly polarized superposition to rotate as it propagates through the magnetised plasma. In radio observations the rotation angle, $\theta_{\rm R}$, is related to the wavelength of the radiation, λ, and the rotation measure, R, by

$$\theta_{\rm R} = R\lambda^2. \tag{1.44}$$

R is found through the line of sight integral,

$$R = 0.81 \int n_e B_{\parallel} dl, \tag{1.45}$$

where distances are measured in parsecs, n_e is the number density of electrons measured in cm^{-3} and $B_{\parallel}$ is the component of the magnetic field parallel to the line of sight measured in microgauss (Burke & Graham-Smith, 1997). Faraday depolarization is where unresolved fluctuations in n_e and $B_{\parallel}$ decrease the observed polarization fraction. Table 1.5 lists estimates of the typical values of $\theta_{\rm R}$ for diffuse emission at different frequencies. The values were estimated by Jones et al. (2017, in prep.) who used the catalogue of rotation measures of extragalactic radio sources by Taylor, Stil & Sunstrum (2009). Faraday depolarization becomes significant when $\theta_{\rm R}$ is larger than a radian.

The 408 MHz survey of Haslam et al. (1982) is the most reliable low-frequency survey currently available. It still suffers from scan-synchronous residuals but there have been successful attempts to remove these, in particular Remazeilles et al. (2015). The Remazeilles et al. (2015) map is now almost universally used instead of the original and will hereafter be referred to as the Haslam map. The S-band Polarization All-Sky Survey (S-PASS) has mapped the southern sky at a central frequency of 2.3 GHz and resolution of 9" with the Parkes telescope (Carretti, 2010). The Q-U-I

JOint TEnerife survey (QUIJOTE) is an ongoing experiment that will eventually map several regions of the sky between 10–40 GHz at sub-degree resolution (Génova-Santos et al., 2015b). The first region forms a wide Galactic survey covering $20,000\,\mathrm{deg}^2$ and the second is a deeper survey covering $3000\,\mathrm{deg}^2$. Further observations have been made by QUIJOTE for Galactic science purposes, focussing on regions with AME emission and the Galactic centre.

1.5 C-BASS and thesis layout

To accurately measure the signal from inflationary B-modes requires more large-area surveys at a range of frequencies to help remove the polarized foreground emission. In total intensity there are still large degeneracies between the low-frequency foreground components that could be removed with more data.

The C-Band All-Sky Survey (C-BASS) is a project to produce an all-sky map in Stokes I, Q and U at a central frequency of 5 GHz with an angular resolution of 0.73°. To produce an all-sky map from the ground requires two telescopes, one in the northern and one in the southern hemisphere. The key science goal of C-BASS is to improve the removal of polarized foregrounds from the CMB and hence enable the detection of the primordial B-mode signal. Other goals include improving the removal of foregrounds to the CMB in total intensity, measuring the Galactic magnetic field and studying the distribution and nature of anomalous microwave emission.

This thesis details the work I have contributed towards the C-BASS project and, where clearly stated, the work of others in the C-BASS collaboration. When I joined the project, C-BASS North was already making survey observations and commissioning had begun on C-BASS South.

In Chapter 2 I give an overview of the C-BASS survey. It contains a description of the receivers, the survey strategy, the target sensitivity and forecasts from simulations that I ran to demonstrate the constraining power of a C-BASS data point when fitting a pixel-based parametric model of the sky to maps at various frequencies.

In Chaper 3 I summarise the data reduction pipeline, which was largely developed by others, and my work characterising the noise properties of calibrated time-ordered data from C-BASS North.

In Chapter 4 I describe the mapmaking process, which converts calibrated time-ordered data into sky maps. I use simulations to quantify the errors introduced into the C-BASS maps by the mapmaker and describe the modifications I made to the mapmaker to reduce the signal error caused by sub-pixel fluctuations in the sky. I test

the C-BASS North maps for systematics using jackknife tests and present preliminary maps of the northern sky in total intensity and polarized intensity.

In Chapter 5 I analyse the diffuse Galactic emission in total intensity. I present a 3-colour image of the sky that demonstrates the complexity of low-frequency diffuse Galactic emission. I extend the standard T-T plot method and use it to estimate the spectral index across the northern sky between 408 MHz and 5 GHz. I use T-T plots of Barnard's Loop to check the calibration of the C-BASS map against other datasets.

In Chapter 6 I analyse the diffuse Galactic emission in polarized intensity. I produce a map of the spectral index between 5 and 30 GHz, modelling measurements of the polarized intensity as Rician random variables.

In Chapter 7 I summarise this thesis and detail future work.

Chapter 2

The C-Band All-Sky Survey (C-BASS)

This first half of this chapter (Section 2.1) contains an overview of the C-BASS project including descriptions of the telescopes' optics (Section 2.1.1), the continuous comparison architecture (Section 2.1.2), the polarimeters (Section 2.1.3), the overall scan strategy (Section 2.1.4) and the target sensitivity (Section 2.1.5). The second half of this chapter (Section 2.2) contains forecasts that demonstrate the constraining power of C-BASS data using pixel-based parametric component fitting.

2.1 Overview of the survey and the telescopes

A summary of the C-BASS survey and telescopes is given in Table 2.1. The survey is at a central frequency of 5 GHz, which is low enough in frequency for the maps to be largely dominated by synchrotron and free-free radiation, but high enough in frequency that Faraday rotation and depolarization will be small across most of the sky. The peak in the B-mode power spectrum is at a few degrees, which sets the required resolution at $\lesssim 1^\circ$. To produce a map that is sensitive to structure on large angular scales requires a total-power scanning telescope. At 5 GHz, the astronomical signal is strong enough to detect with a single-pixel receiver at high signal-to-noise in a timescale of roughly two years (see Section 2.1.5). To produce an all-sky map from the ground requires two telescopes, one in the northern and one in the southern hemisphere. Photographs of both instruments are shown in Figure 2.1.

The northern system was developed first and the lessons learnt and technological advancements made during its development were applied to the southern system. The major difference between the two systems is that the radiometer/polarimeter in C-BASS North is fully analogue where as in C-BASS South it is digital. The digital

Table 2.1: A summary of the main features of the two C-BASS telescopes.

	C-BASS North	C-BASS South
Observatory	OVRO, California	SKA-Klerefontein
Latitude	37.2°	−30.7°
Optical configuration	Gregorian	Cassegrain
Primary mirror diameter	6.1 m	7.6 m
Central frequency (nominal)	5 GHz	
Bandwidth (nominal)	1 GHz	
Angular resolution (FWHM)	0.73°	
Stokes coverage	I, Q, U	

Figure 2.1: Photographs of the C-BASS North (*left*) and C-BASS South (*right*) telescopes. Images from http://www.astro.caltech.edu/cbass/.

radiometer/polarimeter provides C-BASS South with spectral capability. The two systems also have different primary optics, which means that each telescope required individual optical designs in order for them both to have matched beams.

Like most modern CMB experiments, the C-BASS data products will be HEALPix maps[1]. HEALPix (Hierarchical Equal Area isoLatitude Pixelization) is a method of dividing the sky into pixels of equal surface area. The lowest resolution HEALPix map, $N_{\rm side} = 1$, breaks the sphere into 12 pixels. At higher resolution the sky is broken into $12 \times N_{\rm side}^2$ pixels where $N_{\rm side}$ can take values 2^n where n is any non-negative integer. Some of the results in this thesis have been derived using the HEALPix (Gorski et al., 2005) package.

2.1.1 Optics

The northern antenna was originally designed as a prototype for a concept to replace the large antennas of the Deep Space Network with arrays of small antennas (Imbriale, 2003). It was subsequently donated to the project by the Jet Propulsion Laboratory. It is of Gregorian design with a single-piece primary mirror of diameter of 6.1 m. The southern antenna was designed to communicate with low-earth orbit satellites and was donated to SKA South Africa by Telkom SA. It is a Cassegrain telescope with a segmented primary mirror of diameter 7.6 m. The differences meant both systems required different optical configurations with independently shaped secondary mirrors to produce well-matched beams, which are shown in Figure 2.2 (Holler et al., 2013).

The telescopes have far-out sidelobes due to signal spillover past the edges of the mirrors. As the telescopes scan the sky, the sidelobes scan across the ground and introduce a scan-synchronous contamination into our data. To decrease the sidelobes of the Northern telescope and the resultant ground pickup, absorbing baffles were added around the primary and secondary mirrors (Holler et al., 2013). The absorbing baffles themselves radiate thermally but this radiation is relatively stable and unpolarized. The small increase in thermal noise due to the baffles is preferable to scan-synchronous contamination from the ground. Absorbing baffles are not needed on C-BASS South as it has a larger primary mirror that is heavily under-illuminated, and results in intrinsically lower sidelobes.

The optics of both telescopes have been designed to minimise cross polarization (Holler et al., 2013). Most radio telescopes use metal legs to support the secondary

[1] http://healpix.sourceforge.net/

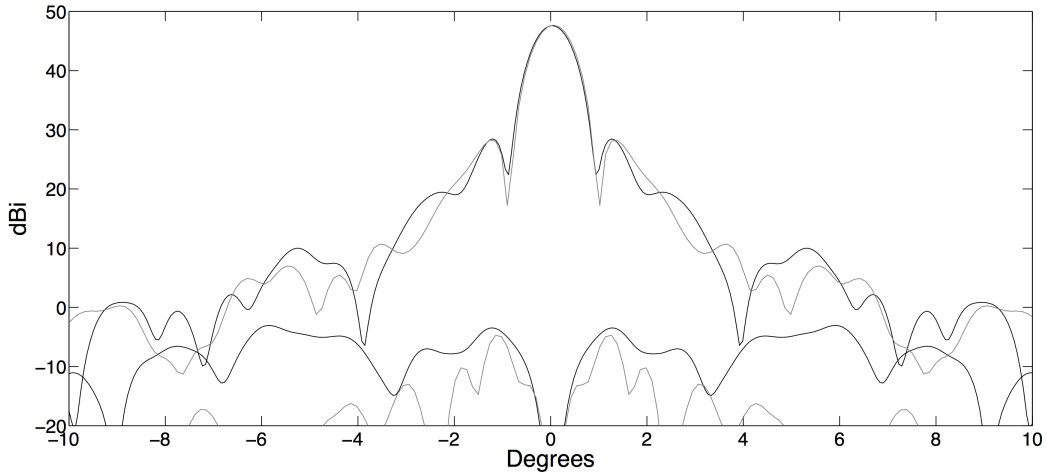

Figure 2.2: The simulated C-BASS North (*black*) and South (*grey*) co-polar and cross-polar beams out to 10°. Figure reprinted from Holler et al. (2013), ©2013 IEEE.

mirrors. These legs break the circular symmetry of the telescope and therefore increase cross polarization by scattering the incoming radiation. To avoid this effect, the C-BASS systems use cones of radio transparent foam to support their secondary mirrors. The C-BASS design maintains the circular symmetry of the system to minimise cross-polarization and avoids scattering the incoming radiation from the ground.

2.1.2 Basic radiometers

This section introduces the concept of system temperature, the effects of gain drifts, and Dicke switched and continuous-comparison radiometers. Both C-BASS receivers employ the continuous-comparison architecture to minimise the effect of gain drifts in the system.

Radio emission from most natural sources takes the form of Gaussian random noise. An antenna couples the electromagnetic radiation from the sky and produces a voltage V, which is proportional to the incident electric field. Like the incoming astronomical signal, the measured output voltage will appear noise-like with zero mean and a variance that is proportional to the incident power P. Radio astronomers commonly measure P in units of temperature. This is convenient because in the Rayleigh-Jeans limit the noise power per unit bandwidth, P_ν, generated by a resistor at temperature T is $P_\nu d\nu = k_B T d\nu$. Appendix A describes the different temperature definitions commonly used in the CMB community.

The system temperature, $T_{\rm sys}$, is a measure of the total noise power generated by a system. To a reasonable approximation the powers from different sources can be

added, giving

$$T_{\rm sys} = T_{\rm CMB} + T_{\rm source} + T_{\rm atm} + T_{\rm ground} + T_{\rm rx} \tag{2.1}$$

where the contributions here are from the CMB, an astronomical source, the atmosphere, the ground and the receiver itself, all temperatures being referenced to the top of the atmosphere. If the receiver consisted of many components connected in series with gains G_1, G_2, G_3... and noise temperatures T_1, T_2, T_3... then the Friis formula gives

$$T_{\rm rx} = T_1 + \frac{T_2}{G_1} + \frac{T_3}{G_1 G_2} + ... \tag{2.2}$$

$T_{\rm sys}$ determines the white noise on a measurement, $\Delta T_{\rm w}$. During an observation of length τ there are $2\Delta\nu\tau$ independent samples where $\Delta\nu$ is the bandwidth. Each sample is drawn from a normal distribution with variance $2T_{\rm sys}^2$. Averaging these samples together reduces the noise on the final measurement and the white noise level is related to $T_{\rm sys}$ via the radiometer equation,

$$\frac{\Delta T_{\rm w}}{T_{\rm sys}} = \frac{1}{\sqrt{\Delta\nu\tau}}. \tag{2.3}$$

To detect the weak diffuse Galactic radio signals requires high levels of amplification and, due to the large gains required, even low levels of gain instability introduce $1/f$ noise that can dominate over the thermal noise of the receiver. Following the derivation in Rohlfs & Wilson (2004), if the system has a gain G then the output power W is

$$W = kT_{\rm sys}G\Delta\nu \tag{2.4}$$

where k is the Boltzmann constant. Any change in G is indistinguishable from a change in $T_{\rm sys}$, such that

$$W + \Delta W = kT_{\rm sys}(G + \Delta G)\Delta\nu \tag{2.5}$$

$$W + \Delta W = k(T_{\rm sys} + \Delta T_{\rm G})G\Delta\nu, \tag{2.6}$$

this gives

$$\frac{\Delta T_{\rm G}}{T_{\rm sys}} = \frac{\Delta G}{G}. \tag{2.7}$$

Ideally $\Delta T_{\rm G} \ll \Delta T_{\rm w}$ and this sets the required stability of the system,

$$\frac{\Delta G}{G} \ll \frac{1}{\sqrt{\Delta\nu\tau}}. \tag{2.8}$$

Gain instabilities can be removed using a Dicke Switch (Dicke, 1946) to make a differential measurement. The system diagram of a simple Dicke switch radiometer

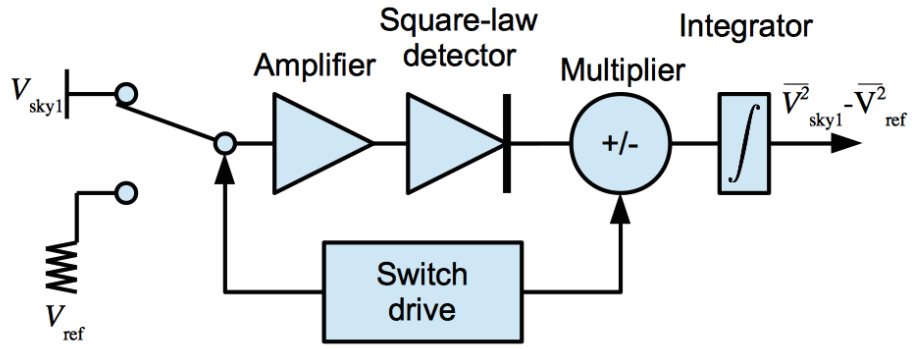

(a) Simple Dicke switch system diagram.

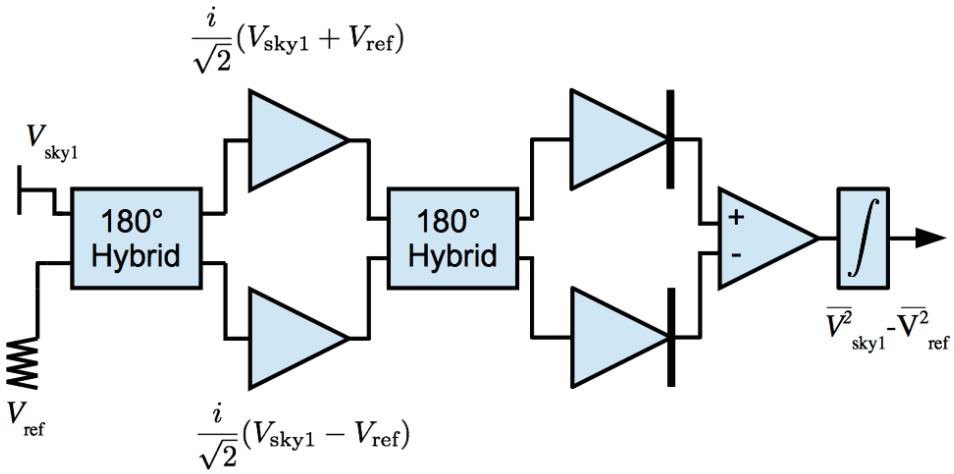

(b) Simple continuous comparison system diagram.

Figure 2.3: Simple system diagrams of gain-stabilised radiometers measuring a single polarization.

measuring a single polarization is shown in Figure 2.3(a). The incoming signal at the front of the system is rapidly switched between the sky and a reference signal. The reference signal is often either a temperature controlled load or other part of the sky. After differencing the sky and reference signal, the gain instabilities in the differential signal are given by

$$\frac{\Delta T_{\mathrm{G}}}{T_{\mathrm{A}}} = \frac{\Delta G}{G}\left(\frac{T_{\mathrm{A}} - T_{\mathrm{ref}}}{T_{\mathrm{sys}}}\right), \tag{2.9}$$

where T_{A} is the antenna temperature of the source and T_{ref} is the temperature of the reference signal. Trivially setting $T_{\mathrm{ref}} = T_{\mathrm{A}}$ gives $\Delta T_{\mathrm{G}} = 0$ and removes the dependence on gain fluctuations. Such a system is called a "balanced system" and the temperature of the reference signal is called the balance temperature.

The removal of the effect of gain drifts comes at a cost. The observing time on

source is reduced by a factor of two, which in turn increases $\frac{\Delta T_{\rm w}}{T_{\rm sys}}$ by a factor of $\sqrt{2}$. There is a further increase in $\frac{\Delta T_{\rm w}}{T_{\rm sys}}$ by a factor of $\sqrt{2}$ because two independent and noisy measurements have been subtracted from one another.

Continuous comparison radiometers are an alternative to the Dicke switch. The system diagram for a simple continuous comparison radiometer is shown in Figure 2.3(b). Continuous comparison radiometers take advantage of the coherence of radio signals before detection. The sky and reference voltage signals are summed and differenced before being sent down separate chains of RF components and then recovered at the end, prior to power detection. A differential measurement can then be made by subtracting the recovered sky and reference signals. The RF summing and differencing is done using 180° hybrids. A 180° hybrid is a device with four ports, two for the input signals and two for the output signals. The two output signals are the linear sum and difference of the two input signals with an extra 180° phase shift introduced between the two outputs.

The continuous comparison architecture recovers a factor of $\sqrt{2}$ in $\frac{\Delta T_{\rm w}}{T_{\rm sys}}$ compared to a Dicke switch as the sky is observed continuously but increases the financial cost of the system as two chains of RF components are required.

To balance a continuous comparison system, $T_{\rm ref}$ must be set to the system temperature in the sky channel at the first hybrid. In a real system, $T_{\rm ref}$ will not be exactly equal to the physical temperature of the reference load because of losses in the connecting cables and the extra thermal noise added by those cables.

Both C-BASS receivers are continuous-comparison systems with temperature-controlled cold loads used as the reference signal. The Northern system diagram is shown in Figure 2.4 and the Southern system diagram in Figure 2.5. They are fundamentally the same system up until the second set of 180° hybrids in C-BASS North. C-BASS South digitises the signal at this point and the subsequent functions are carried out digitally in a field programmable gate array (FPGA). C-BASS North performs the same operations using RF components.

A noise diode is used to inject a calibration signal of constant temperature and polarization into the two sky channels using -30 dB directional couplers. To calibrate drifts in the gain of the system, the power is scaled so that the amplitude of the noise diode is constant. An astronomical calibrator is later used to convert the data to brightness measured in Jy beam^{-1} (Section 3.1.6).

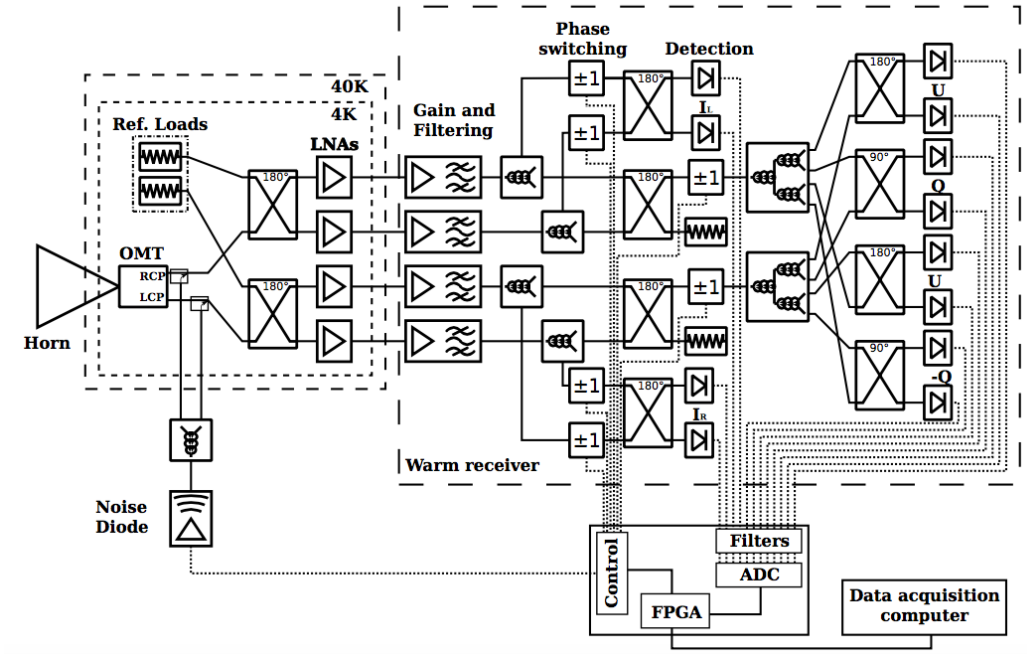

Figure 2.4: C-BASS North system diagram. Figure reprinted from King et al. (2010), with permission of the authors. Note that the L2C is not labelled on the diagram but is immediately after the OMT.

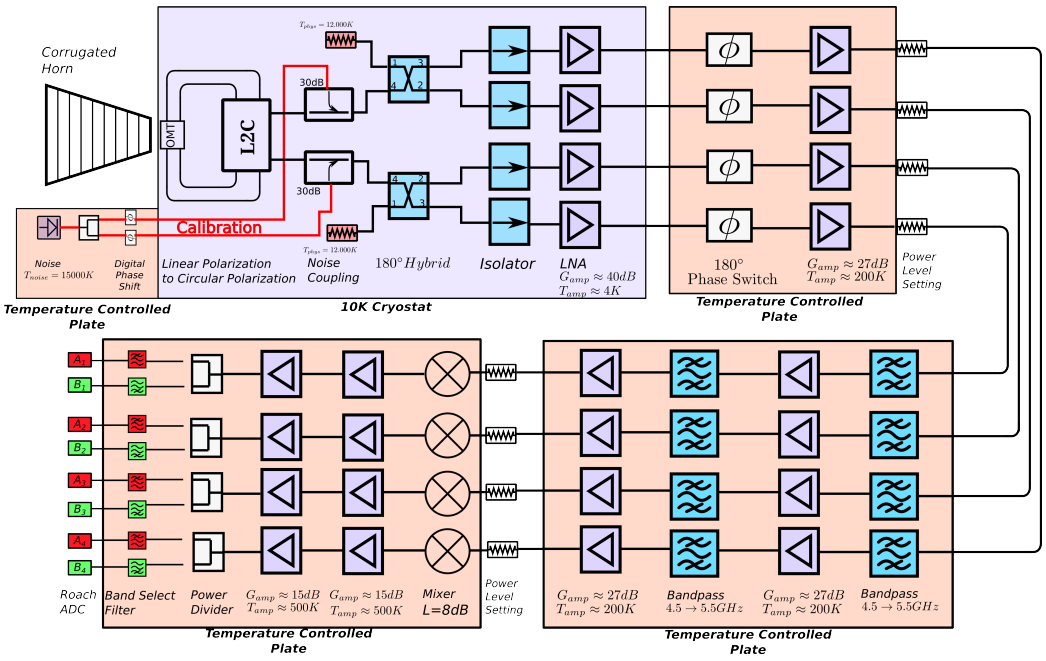

Figure 2.5: C-BASS South system diagram. Figure reprinted from Copley (2013), with permission of the author.

2.1.3 Polarimetry

The final C-BASS data products will be Stokes I, Q and U sky maps. This section shows how the Stokes parameters are related to the components of an electric field and how the C-BASS systems measure them. The northern polarimeter is an analogue system, using RF components to construct the Stokes parameters. The southern system does the polarimetry digitally.

Stokes parameters are related to the complex amplitudes of an electric field in either a Cartesian basis (E_x and E_y) or a circular basis (E_l and E_r) via

$$I = |E_x|^2 + |E_y|^2 = |E_l|^2 + |E_r|^2 \tag{2.10}$$

$$Q = |E_x|^2 - |E_y|^2 = 2\Re(E_l E_r^*) \tag{2.11}$$

$$U = 2\Re(E_x E_y^*) = -2\Im(E_l E_r^*) \tag{2.12}$$

$$V = -2\Im(E_x E_y^*) = |E_l|^2 - |E_r|^2. \tag{2.13}$$

For most astrophysical sources $V = 0$ and $Q, U \ll I$. To construct Q from E_x and E_y requires the subtraction of two large numbers, which is an inherently inaccurate process. Constructing Q from E_l and E_r however uses a correlation which removes uncorrelated noise and results in far a more accurate measurement. For these reasons it is preferable to work with E_l and E_r. The C-BASS systems use linearly polarized feeds but immediately rotate the E_x and E_y channels to E_l and E_r with analogue components.

The northern receiver is a completely analogue system (King et al., 2014b). It sums and differences the sky and load and correlates E_l and E_r to form Q and U with RF components. It only digitises the channels after I, Q and U have been constructed. Digitisation technology advanced sufficiently during the development of C-BASS North that the southern receiver could implement the same architecture with a digital back-end (Copley, 2013).

The first component in the C-BASS North system diagram (Figure 2.4) is the feed horn. The horn guides incident radiation down a tapered wave guide to the orthomode transducer (OMT). The OMT consists of four orthogonal probes that couple to E_x and E_y. The C-BASS OMT, like the optics, was designed to have a very low cross-polarization (Grimes et al., 2007). The signal from the OMT is then sent through a linear to circular converter (L2C). The L2C is made of two 180° hybrids and one 90° hybrid. Together these hybrids make linear combinations of, and introduce phase shifts between, the input signals to convert them from linearly to circularly polarized.

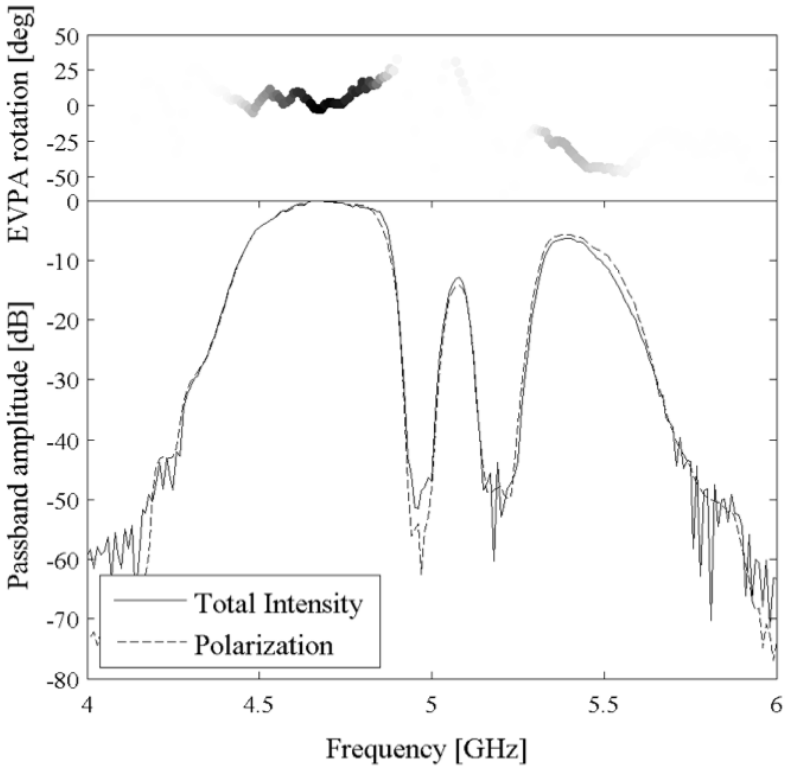

Figure 2.6: The *bottom* subplot shows the C-BASS North passband in total intensity (*solid line*) and linear polarization (*dashed line*). The *top* subplot shows the polarization angle of the system across the band where the intensity of the point is determined by the amplitude of the passband. Figure reprinted from King et al. (2014a), with permission from OUP. For permission to reuse, please contact the rights holder.

On the system diagram the sky channels are labeled LCP (left circular polarization) and RCP (right circular polarization).

Directional couplers then couple in the noise diode signal before the sky and reference load signals enter the first 180° hybrid, which sums and differences the two inputs. The cold reference-loads are fixed to the cold plate inside the cryostat. There are now four channels in the receiver; RCP + load1, RCP − load1, LCP + load2 and LCP − load2. Each of these channels is amplified and passed through filters that define the 4.5–5.5 GHz passband, which is shown in Figure 2.6. The channels are then split to the radiometer and polarimeter sections of the receiver.

The radiometer part of the receiver separates the sky and reference load signals with a further set of 180° hybrids. The powers of the LCP, RCP, load1 and load2 channels are then detected using diodes. We then have two independent measure-

ments of the total intensity,

$$I_1 = |E_l|^2 = I - V \tag{2.14}$$

$$I_2 = |E_r|^2 = I + V. \tag{2.15}$$

The polarimeter part of the receiver also separates the sky and reference load signals with 180° hybrids. Because Q and U are found from correlations, unlike the intensity signals, they do not require stabilisation and so the load signals are dumped in a resistor. The remaining hybrids correlate LCP and RCP to produce Q and U. The system provides two estimates of Q and U but these are not independent.

Figure 2.6 shows how the polarization angle of the northern polarimeter changes with frequency (the polarization angle of an ideal instrument is the same at all frequencies). The measured polarization angle of astronomical emission is therefore a function of its frequency spectrum. Using equations 5 and 6 from King et al. (2014a) we estimate the change in measured polarization angle between emission with a spectral index of $\beta = 0$ and $\beta = 4$ to be 0.5°. This effect is therefore negligible across most of the sky.

Having detected I, Q and U at the detector diodes the signals are then digitised. The C-BASS analogue to digital converters (ADC) sample at 2 MHz but, due to the C-BASS resolution and scan speed, sky signal only appears at frequencies below 40 Hz. The final system sampling rate is therefore reduced to 100 Hz. A series of filters is used to decimate the time series and apply a low pass filter with a cut-off at 50 Hz. The finite width of the filter introduces correlations on very short timescales, Figure 2.7 shows the autocorrelation function of the filter chain. After roughly 90 ms the autocorrelation has dropped to 1% (see Stevenson (2013) for full details).

The Southern receiver, shown in Figure 2.5, digitises the signal much sooner than the Northern receiver. All of the polarimetry and separation of sky and reference load channels is performed digitally. The digital southern polarimeter has spectral resolution, dividing the 1 GHz bandwidth into 128 evenly spaced frequency channels. The spectral back end allows for dynamic RFI flagging (Section 3.1.3), more informed setting of the balance temperature and the potential to measure the Faraday rotation across the band to help measure the strength of the Galactic magnetic field.

2.1.4 Scan strategy

The C-BASS telescopes scan in azimuth at fixed elevations. The lowest elevation that each telescope scans at cuts through either the north or south celestial pole. The

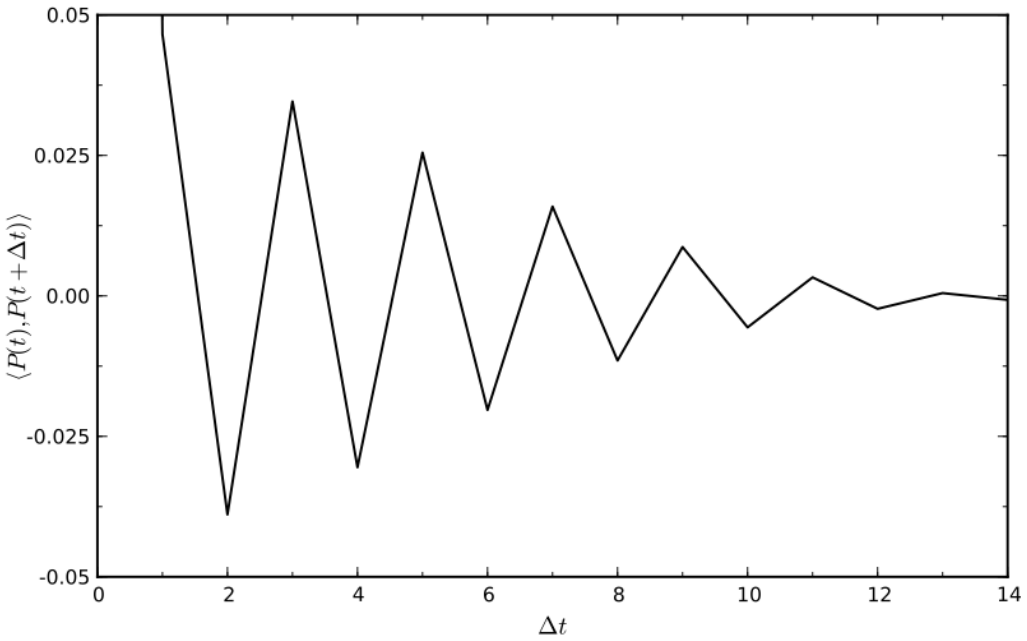

Figure 2.7: Autocorrelation function of the C-BASS North filters. The function is normalised to 1 at $\Delta t = 0$, Δt is in units of 10 ms. Figure reprinted from Stevenson (2013), with permission of the author.

telescopes also scan at elevations 10°, 30° and 40° above the pole. The telescopes scan at five speeds, 3.8, 3.9, 4.0, 4.1 and 4.2 degrees per second.

The telescopes scan in long sweeps of the sky so that they can detect signal across the largest angular scales. By scanning at fixed elevations the variation in the atmospheric emission is minimised and repeatable ground templates can be constructed. There are many advantages to scanning through the celestial poles. Firstly it ensures complete sky coverage. Scans at an elevation equal to the telescope latitude, ϕ, will observe all declinations between the north/south celestial pole and $\pm\,|-90° + 2\,|\phi||$ in the northern/southern hemisphere. The latitudes of the telescopes (Table 2.1) ensure significant overlap in sky coverage from both receivers that will allow for continuity tests between both surveys and will ease cross calibration. Secondly these scans return to the same point in the sky every 360° which provides a check on any drifts in the system or atmosphere.

If each telescope only scanned at the elevation required to cross its respective celestial pole then the integrated observing time per unit solid angle (sky coverage) would be significantly lower at intermediate declinations than at the poles and the lower declination bound. We therefore scan at elevations 10°, 30° and 40° above the pole to fill in these regions. The sky coverage of several C-BASS North scans

at constant elevation are shown in Figure 2.8. To demonstrate the sky coverage of the final C-BASS maps, the integration time per pixel of a map made from the C-BASS North night-time survey scans and simulated C-BASS South scans is shown in Figure 2.9.

Scanning at multiple elevations also increases the angular distribution of scans through a pixel (cross linking). Greater cross linking allows for more reliable reconstruction of the sky during mapmaking (Section 4.1). Scanning at different elevations also maps systematics at fixed frequencies to different angular scales on the sky and this will help to average them away (Section 3.1.5).

The faster the telescopes scan the sky the less impact $1/f$ receiver drifts will have on the resulting data. The fastest both telescopes can scan is roughly four degrees per second. During survey observations the C-BASS telescopes scan at five speeds close to $4^\circ\,\mathrm{s}^{-1}$ so that systematics in the data at fixed frequency map to different angular scales on the sky and hence average down.

The telescopes perform survey scans for approximately 200 minutes, with sky dips at the beginning, middle and end to monitor the atmospheric opacity. The noise diode is fired every two sweeps of the sky to monitor and correct for any gain drifts in the system. The survey scans are bookended by calibration observations, which are used to check the telescope pointing and to convert the data to astronomical units. This complete set of observations is then repeated all day and night (Section 3.1).

2.1.5 Target sensitivity

We aim to detect polarized emission over most of the sky. To determine the required sensitivity, we extrapolated the *WMAP* K-band polarized intensity map to 5 GHz with a spectral index of 3. The polarized intensity at 5 GHz was estimated to be greater than 0.5 mK over 90% of the sky. To achieve a 5σ detection of polarized emission over this fraction of the sky sets the target sensitivity at $0.1\,\mathrm{mK\,beam}^{-1}$. In I this is lower than the confusion limit due to extragalactic radio sources. Using the 1.4 GHz VLA survey (Mitchell & Condon, 1985; Condon et al., 1998) the confusion limit for a Gaussian beam is estimated by Condon (2002) to be

$$\frac{\sigma_c}{\mathrm{K}} \simeq 0.07 \left(\frac{\nu}{\mathrm{GHz}}\right)^{-2.7}. \tag{2.16}$$

This gives $\sigma_c \simeq 0.9$ mK at 5 GHz.[2] Whilst the final C-BASS I maps will be confusion limited, the Q and U maps will not because extragalactic sources are typically only a few per cent polarized.

[2]Note that this is independent of the beam area.

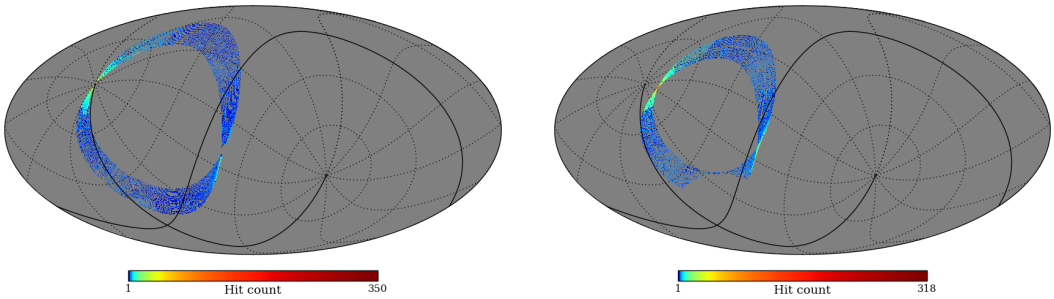

(a) Elevation 37°

(b) Elevation 47°

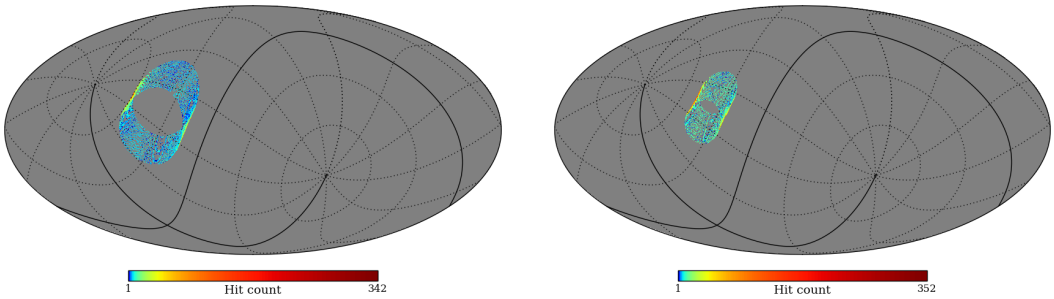

(c) Elevation 67°

(d) Elevation 77°

Figure 2.8: Example coverage of small number of azimuth scans made by C-BASS North at fixed elevations. The hit counts are the number of samples in each $N_{\text{side}} = 512$ pixel. The maps are in Galactic coordinates and the grid lines show lines of constant RA and DEC. The missing pixels in the map of elevation 47° scans are samples within 60° of the sun that have been removed.

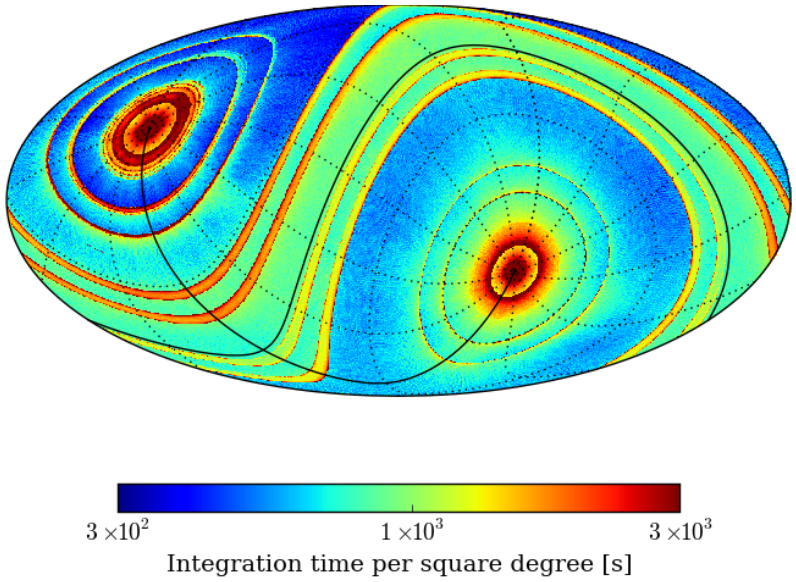

Figure 2.9: The integration time per square degree in an $N_{\rm side} = 512$ map made from C-BASS North night-time survey scans and simulated C-BASS South scans at all elevations. The graticules show lines of constant right ascension and declination.

To find the corresponding target sensitivity in a HEALPix pixel, we multiply the target sensitivity per beam by the square root of the ratio of the beam and pixel solid angles,

$$\frac{\sigma}{\mathrm{mK\,pixel}^{-1}} = \frac{\sigma}{\mathrm{mK\,beam}^{-1}} \times \sqrt{\frac{\Omega_{\rm beam}}{\Omega_{\rm pixel}}}, \tag{2.17}$$

where $\Omega_{\rm beam}$ and $\Omega_{\rm pixel}$ are the solid angles of the beam and HEALPix pixel respectively. The C-BASS beam can be approximated as a Gaussian and such a beam occupies a solid angle

$$\Omega_{\rm beam} = \frac{\pi\theta^2_{\rm FWHM}}{4\ln 2}, \tag{2.18}$$

where $\theta_{\rm FWHM}$ is the full width half maximum measured in radians. For the C-BASS beam, $\theta_{\rm FWHM} = 0.73^\circ$ and so $\Omega_{\rm beam} \simeq 1.839 \times 10^{-4}\,\mathrm{sr}$. Table 2.2 contains the target sensitivity expressed in mK pixel^{-1} for a range of HEALPix resolutions.

Table 2.2: The target sensitivity for HEALPix pixels at different resolutions

$N_{\rm side}$	Pixel Area [sr]	$\sigma_{\rm target}$ [mK pixel^{-1}]
64	2.56×10^{-4}	0.0848
128	6.39×10^{-5}	0.170
256	1.60×10^{-5}	0.339
512	4.00×10^{-6}	0.678
1024	9.99×10^{-7}	1.36

2.2 Component separation forecasts

To demonstrate the impact of C-BASS data on the results of CMB component separation we simulated the parametric pixel-fitting process for individual pixels from a variety of mock data sets, each with different levels of foreground contamination chosen to represent different sky regions, using properties of existing or planned surveys. The surveys included are listed in Table 2.3. Other results may be obtained using different methods. For example Errard et al. (2016) use a Fisher matrix approach to forecast how well future CMB missions will perform. They show significant improvements in B-mode fitting when including C-BASS data with current and future experiments, the noise degradation factor (the amount that foregrounds are increasing the final uncertainty) is decreased by a factor of 2–3 when including C-BASS.

The sky model used to simulate the data was a simplified version of that found in Table 4 of Planck Collaboration et al. (2016b). Our simplified model for total intensity is summarised in Table 2.4 and for polarization in Table 2.5. The intensity model includes CMB, synchrotron, free-free, AME and thermal dust components. The polarization model does not include the free-free or AME components but extends the synchrotron model to include a curvature term, C_s, as an extra free parameter.

For each representative pixel we used the Planck Collaboration et al. (2016b) component separation results to suggest values for the foreground parameters. The CMB amplitude was set from a model power spectrum, with standard cosmological parameters that are listed in Table 1.1 and $r \simeq 0.1$, at the angular scale of the simulated pixel. The intensity measurements were simulated in 1° pixels because all components, including the CMB, are detected at high signal-to-noise on this scale. We simulated B-mode measurements in 3° pixels as this scale gives the maximum pixel-to-pixel fluctuations in the CMB B-mode signal. We did not add noise realisations to the mock data so as to avoid biasing the results of the fitting process by individual realisations. Instead we simply used the assigned noise levels to calculate the likelihood in the fitting process. The sky model was then fitted back to

Table 2.3: Frequencies and either achieved or expected sensitivities of the simulated data points used in the component separation forecasts.

Instrument	Frequency [GHz]	$\sigma^T(1^\circ)$ [μK$_{\rm RJ}$]	$\sigma^P(3^\circ)$ [μK$_{\rm RJ}$]	References
Haslam	0.408	2.5×10^6	...	Haslam et al. (1982)
C-BASS	5.0	73.0	24.0	
WMAP K-band	22.8	5.8	...	Planck Collaboration et al. (2016b)
WMAP Ka-band	33.0	4.2	...	—"—
WMAP Q-band	40.7	3.5	...	—"—
WMAP V-band	60.7	3.8	...	—"—
WMAP W-band	93.5	3.9	...	—"—
Planck 30	28.4	2.5	1.1	—"—
Planck 44	44.1	2.6	1.3	—"—
Planck 70	70.4	3.1	1.5	—"—
Planck 100	100.0	1.0	0.51	—"—
Planck 143	143.0	0.33	0.24	—"—
Planck 217	217.0	0.26	0.20	—"—
Planck 353	353.0	0.20	0.19	—"—
Planck 545	545.0	0.086	...	—"—
Planck 857	857.0	0.032	...	—"—
LiteBIRD 60	60.0	...	0.052	Matsumura et al. (2014)
LiteBIRD 78	78.0	...	0.031	—"—
LiteBIRD 100	100.0	...	0.020	—"—
LiteBIRD 143	143.0	...	0.013	—"—
LiteBIRD 195	195.0	...	0.0070	—"—
LiteBIRD 280	280.0	...	0.0038	—"—

Table 2.4: Total intensity sky model used in the component separation forecasts. Table adapted from Planck Collaboration et al. (2016b).

Component	Free parameters and priors	Brightness temperature, s_ν [K$_{\rm RJ}$]	Additional information
CMB	$A_{\rm CMB} \sim \text{Uniform}(-1, 1\,\text{K})$	$x = \frac{h\nu}{k_{\rm B}T_{\rm CMB}}$ $s_{\rm CMB} = A_{\rm CMB}\frac{x^2e^x}{(e^x-1)^2}$	$T_{\rm CMB} = 2.7255\,\text{K}$
Synchrotron	$A_{\rm s} \sim \frac{1}{X}$ $\beta_{\rm s} \sim \text{Normal}(-3 \pm 0.3)$	$s_{\rm s} = A_{\rm s}\left(\frac{\nu}{\nu_0}\right)^{\beta_{\rm s}}$	$\nu_0 = 0.408\,\text{GHz}$
Free-free	$EM \sim \frac{1}{X}$	$g_{\rm ff} = \log\left(\exp\left[5.960 - \sqrt{3}/\pi \log\left(\nu_9 T_4^{-3/2}\right)\right] + e\right)$ $\tau = 0.05468 T_{\rm e}^{-3/2}\nu_9^{-2} EM g_{\rm ff}$ $s_{\rm ff} = T_{\rm e}\left(1 - \exp^{-\tau}\right)$	$T_4 = T_{\rm e}/10^4$ $\nu_9 = \nu/(10^9\,\text{Hz})$ $T_{\rm e} = 7000\,\text{K}$
AME	$A_{\rm AME} \sim \frac{1}{X}$ $\nu_p \sim \text{Uniform}(5, 70\,\text{GHz})$	$s_{\rm sd} = A_{\rm AME}\left(\frac{\nu_0}{\nu}\right)^2 \frac{F(\nu\nu_{\rm p0}/\nu_{\rm p})}{F(\nu_0\nu_{\rm p0}/\nu_{\rm p})}$	$\nu_0 = 22.8\,\text{GHz}$ $\nu_{\rm p0} = 30.0\,\text{GHz}$ $F(\nu) =$ External template*
Thermal dust	$A_{\rm d} \sim \frac{1}{X}$ $\beta_{\rm d} \sim \text{Normal}(1.6 \pm 0.3)$ $T_{\rm d} \sim \text{Uniform}(3, 70\,\text{K})$	$\gamma = \frac{h}{k_{\rm B}T_{\rm d}}$ $s_{\rm d} = A_{\rm d}\left(\frac{\nu}{\nu_0}\right)^{\beta_{\rm d}+1}\frac{\exp(\gamma\nu_0)-1}{\exp(\gamma\nu)-1}$	$\nu_0 = 545\,\text{GHz}$

* SPDUST2 template from http://pages.jh.edu/~yalihai1/spdust/spdust.html

Table 2.5: Polarized sky model used in the component separation forecasts. Table adapted from Planck Collaboration et al. (2016b).

Component	Free parameters and priors	Brightness temperature, s_ν [K_{RJ}]	Additional information
CMB	$A_{\rm CMB} \sim \text{Uniform}(-1, 1\,\text{K})$	$x = \frac{h\nu}{k_B T_{\rm CMB}}$ $s_{\rm CMB} = A_{\rm CMB} \frac{x^2 e^x}{(e^x - 1)^2}$	$T_{\rm CMB} = 2.7255\,\text{K}$
Synchrotron	$A_{\rm s} \sim \frac{1}{\lvert X \rvert}$ $\beta_{\rm s} \sim \text{Normal}(-3 \pm 0.3)$ $C_{\rm s} \sim \text{Normal}_{-0.5<x<0.5}(0 \pm 0.3)$	$s_{\rm s} = A_{\rm s} \left(\frac{\nu}{\nu_0}\right)^{\beta_{\rm s} + C_{\rm s} \ln(\nu/\nu_0)}$	$\nu_0 = 30\,\text{GHz}$
Thermal dust	$A_{\rm d} \sim \frac{1}{\lvert X \rvert}$ $\beta_{\rm d} \sim \text{Normal}(1.6 \pm 0.3)$ $T_{\rm d} \sim \text{Uniform}(3, 70\,\text{K})$	$\gamma = \frac{h}{k_B T_{\rm d}}$ $s_{\rm d} = A_{\rm d} \left(\frac{\nu}{\nu_0}\right)^{\beta_{\rm d}+1} \frac{\exp(\gamma\nu_0) - 1}{\exp(\gamma\nu) - 1}$	$\nu_0 = 353\,\text{GHz}$

the simulated data, using the PyMC Python package (Patil, Huard & Fonnesbeck, 2010) to explore the posterior distribution with a Metropolis-Hastings algorithm. The Metropolis-Hastings algorithm is described in Appendix B.

Figures 2.10 and 2.11 show the recovered frequency spectra when fitting the total intensity model to simulated Haslam, *WMAP* and *Planck* data (*left*) and then when also including a simulated C-BASS frequency point (*right*) for five example pixels. Both the biases and the variances of the recovered frequency spectra are reduced when a C-BASS data point is included in the fit.

Figure 2.12 shows the marginalised posterior density estimates (PDEs) of the total intensity foreground parameters for the region with significant AME emission. Without the simulated C-BASS data point, estimates of β_s, EM, $A_{\rm AME}$, $\nu_{\rm peak}$, $A_{\rm CMB}$ and $A_{\rm d}$ are clearly biased away from their true values. Without sufficient frequency points to disambiguate the components, the solver finds zero free-free-amplitude, gets the synchrotron spectral index wrong by 10%, underestimates the AME amplitude and overestimates its peak frequency. The resulting estimate of the CMB amplitude is also biased. Including the simulated C-BASS data point removes those biases. This demonstrates how the inclusion of C-BASS data when separating the CMB from its foregrounds may produce a final CMB map that can be used with a smaller mask. It also demonstrates how C-BASS can provide insight into the total intensity emission of AME by more accurately recovering its spectral parameters.

Figure 2.13 shows the recovered frequency spectra when fitting the polarization

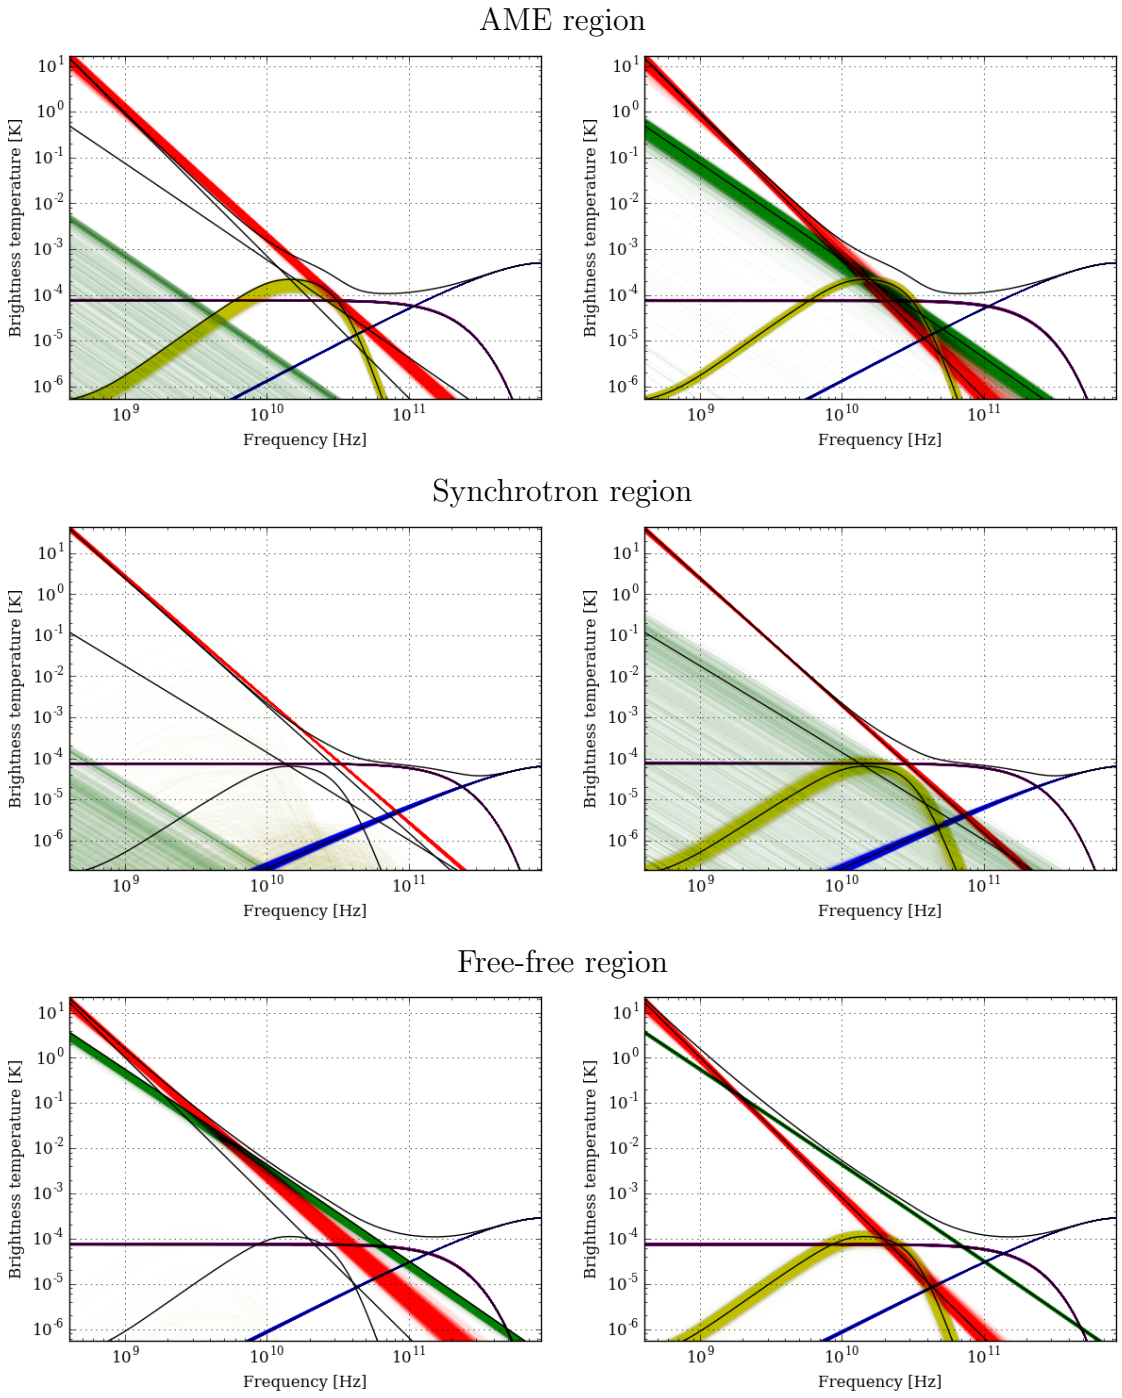

Figure 2.10: Recovered Stokes I frequency spectra of simulated 1° pixels in sky regions with different foreground parameters. The spectra are of example pixels with large amounts of either AME (*top*), synchrotron (*middle*) or free-free (*bottom*) emission. The solid black lines are frequency spectra of the true simulated foreground components. The coloured lines are 5000 randomly drawn samples from the converged MCMC chains. Synchrotron is red; thermal dust is blue; AME is yellow; free-free is green; and CMB is purple. (*left*) Haslam, *WMAP* and *Planck* data. (*right*) Haslam, *WMAP*, *Planck* and C-BASS data.

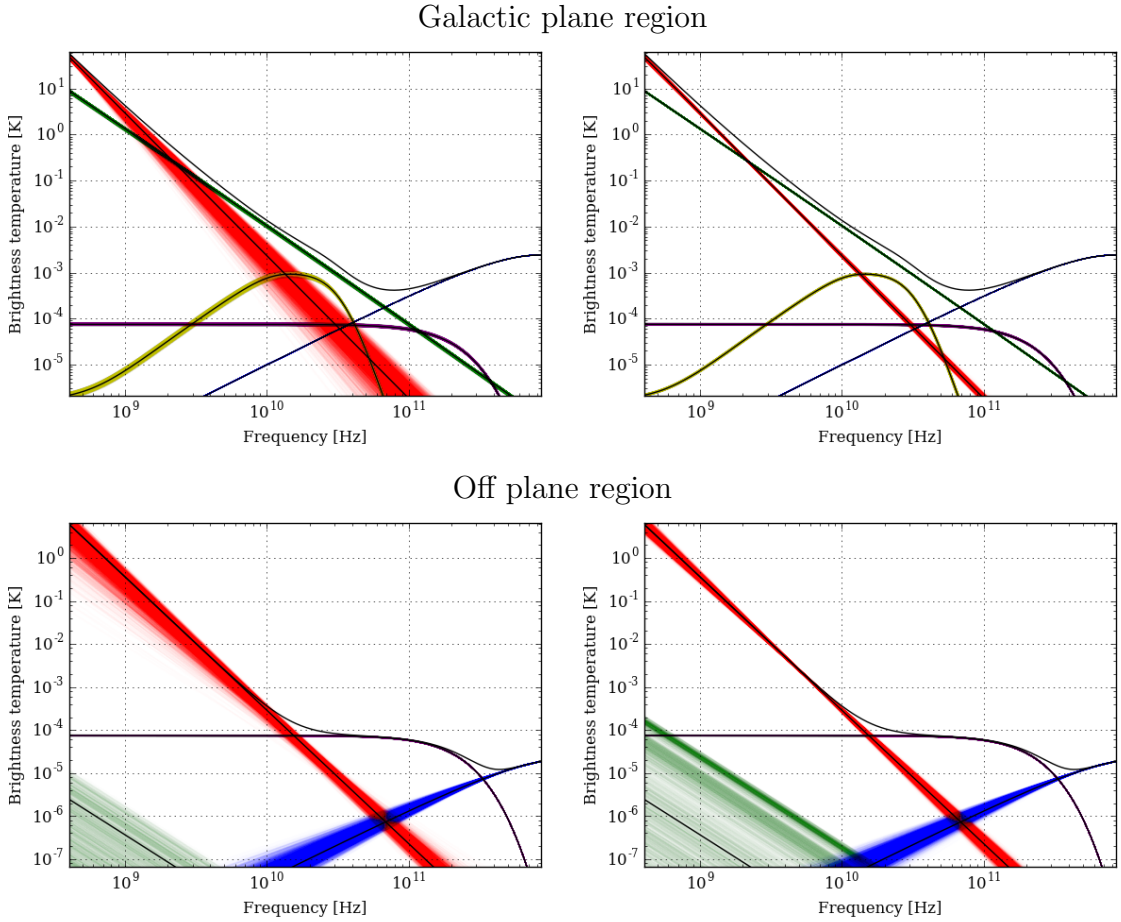

Figure 2.11: Recovered Stokes I frequency spectra of simulated 1° pixels in sky regions with different foreground parameters. The spectra are of an example pixel in the Galactic plane (*top*) and a pixel at high Galactic latitude (*bottom*). The solid black lines are frequency spectra of the true simulated foreground components. The coloured lines are 5000 randomly drawn samples from the converged MCMC chains. Synchrotron is red; thermal dust is blue; AME is yellow; free-free is green; and CMB is purple. (*left*) Haslam, *WMAP* and *Planck* data. (*right*) Haslam, *WMAP*, *Planck* and C-BASS data.

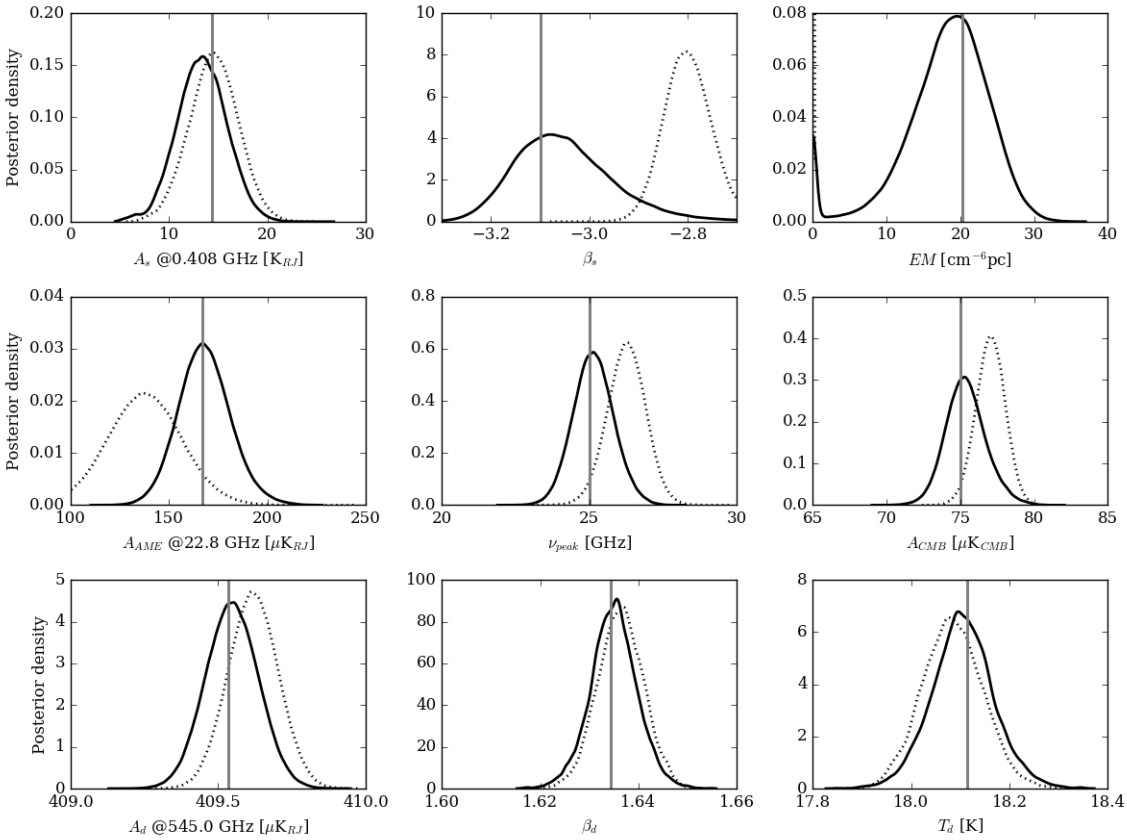

Figure 2.12: PDEs of the Stokes I component parameters for a simulated 1° pixel in a sky region with significant AME. The dashed lines are the PDEs when only including Haslam, *WMAP* and *Planck* data points in the fit. The solid lines are the PDEs when the C-BASS data point is included. The vertical lines are at the true parameter values used to simulate the data.

model to simulated *Planck* and *LiteBIRD* data (*left*) and then when also including a simulated C-BASS frequency point (*right*) for a pixel at high Galactic latitude with low foreground emission. Two simulated datasets were produced, with synchrotron curvature parameter $C_\mathrm{s} = \pm 0.3$. The inclusion of a C-BASS data point will make the greatest difference to the recovered foreground parameters when $C_\mathrm{s} > 0$ and the CMB and synchrotron components have the most similar spectral shapes. The PDEs for the foreground parameters when $C_\mathrm{s} = +0.3$ are shown in Figure 2.14. Using only the *Planck* and *LiteBIRD* frequencies (with no observations below 30 GHz), the data are insufficient to constrain β_s or C_s and these parameters are determined entirely by their priors. This mis-fitting of the synchrotron component biases the estimated CMB B-mode amplitude. Including a C-BASS data point helps constrain the synchrotron spectral parameters which reduces the bias in the estimate of the CMB amplitude by a factor of 2.

These simulations are simplistic but show that C-BASS is necessary to accurately recover the CMB and its foregrounds in polarization. Realistic emission spectra will be more complicated than those simulated here and extra frequency channels, beyond C-BASS, *Planck* and *LiteBIRD*, will likely be needed to verify their spectral shapes.

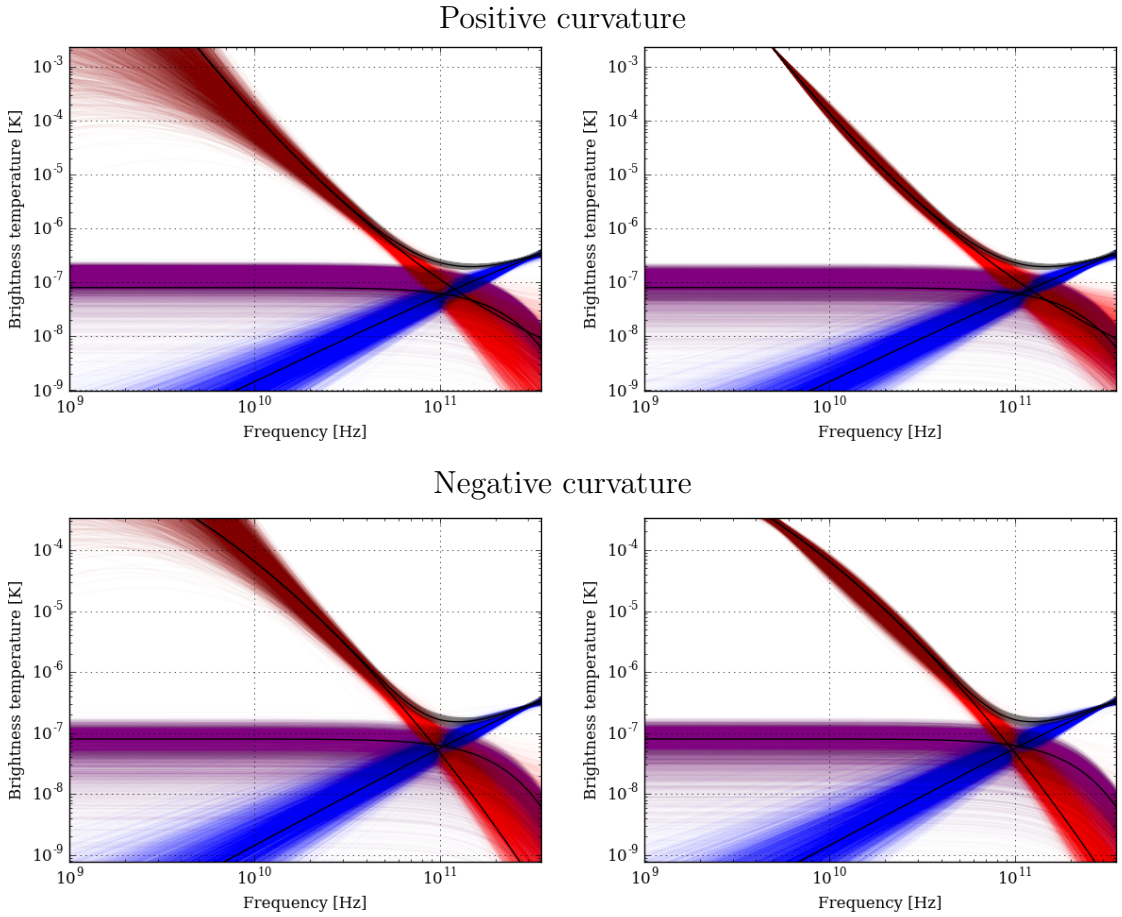

Figure 2.13: Recovered B-mode frequency spectra of a simulated 3° pixel in a sky region with low foreground emission. The spectra are for simulations with positive (*top*) and negative (*bottom*) synchrotron spectral curvature. The solid black lines are frequency spectra of the true simulated foreground components. The coloured lines are 5000 randomly drawn samples from the converged MCMC chains. Synchrotron is red; thermal dust is blue; and CMB is purple. (*left*) *Planck* and *LiteBIRD* data. (*right*) *Planck*, *LiteBIRD* and C-BASS data.

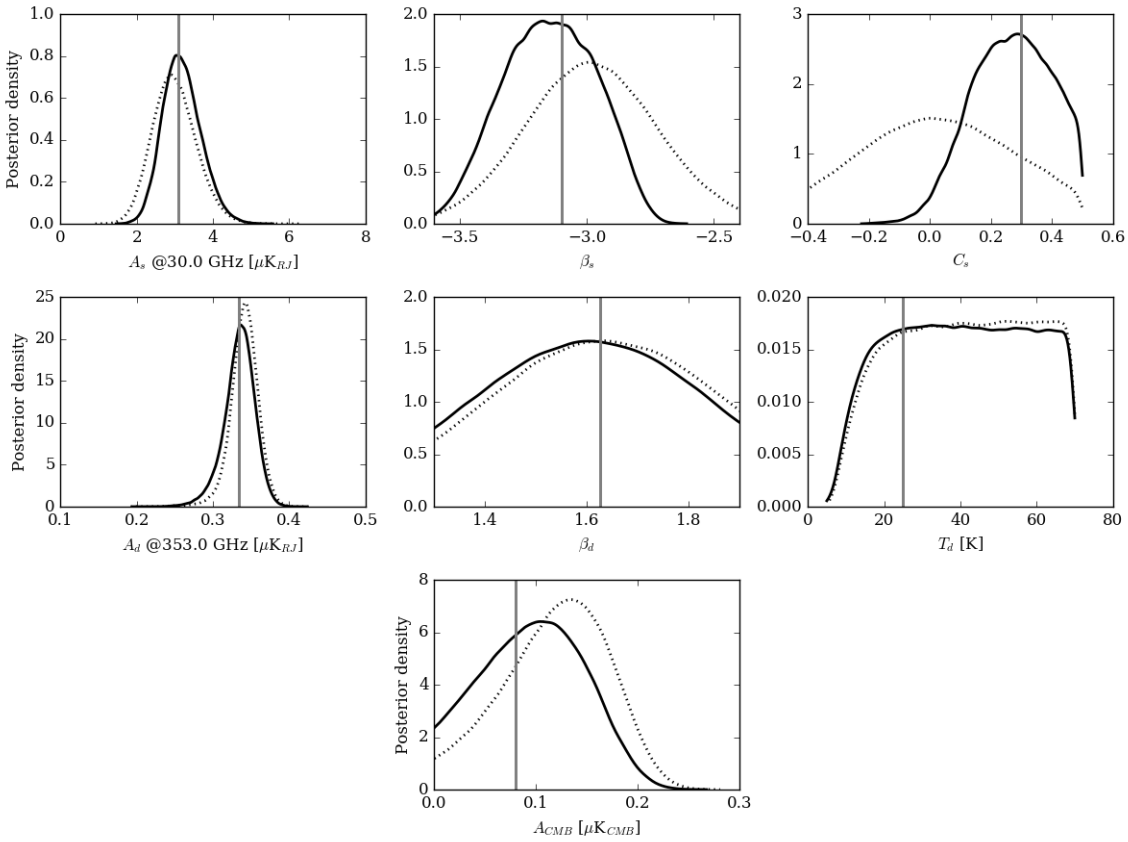

Figure 2.14: PDEs of the B-mode polarization component parameters for a simulated 3° pixel in a sky region with low foreground emission and synchrotron component with positive spectral curvature. The dashed lines are the PDEs when only including *Planck* and *LiteBIRD* data points. The solid lines are the posterior density estimates when the C-BASS data point is included. The vertical lines are at the true parameter values used to simulate the data.

Chapter 3

C-BASS North Data Analysis 1: Time Ordered Data

Section 3.1 contains an overview of the post-processing data pipeline that calibrates raw time-ordered data (TOD) from the telescope and corrects for several systematic effects, specifically ground pickup and a 1.2 Hz contaminating signal introduced by the cold head. It was largely developed before I joined the project and is well documented in Irfan (2014).

In Section 3.2 we define the subsets of C-BASS North data that have been used in the rest of this Thesis.

Section 3.3 shows how we estimated the noise properties of the calibrated TOD. The noise properties are needed in order to produce accurate sky maps with well characterised noise properties and are also used to identify and diagnose problems with the system.

3.1 Data pipeline

Raw C-BASS data are saved as binary files in a non-standard format (DAT files) and contain the uncalibrated TOD from the receiver, the nominal telescope pointing, general housekeeping data such as the temperatures of cryogenic components and also weather information. The data pipeline takes the raw DAT files as inputs and saves the fully calibrated output data in FITS files. The output FITS files contain fully calibrated I, Q and U timestreams along with the actual telescope pointing and a series of flags that identify the status of both the telescope and the data at each time sample. The pipeline is fully automated and, other than the mapmaking steps, implemented using MATLAB R2012b[1]. Figure 3.1 contains a flow chart summarising

[1] https://uk.mathworks.com/products/matlab.html

the key steps in the C-BASS data reduction pipeline. More detail is given on each step in the MATLAB pipeline in the following subsections. The mapmaker is described in Chapter 4.

Partially calibrated Stokes I data from an example standard observation taken from the pipeline after the astronomical calibration step are shown in Figure 3.2. The data have been colour coded to identify key parts of the observation and this figure will be referred to throughout the description of the pipeline. In standard observing mode, survey scans are bookended by a series of calibration observations. In this example observation the telescope surveyed the sky at two elevations, first at 37° and then at 47°. The elevations of the survey scans are labeled on Figure 3.2. The reduced data from the survey scans in this example observation would be output to two FITS files by the pipeline. The first FITS file would contain the scans at elevation 37° and the second FITS file would contain the scans at elevation 47°.

Figure 3.3 is a zoom in to part of the elevation 37° scans shown in Figure 3.2. During survey observations the telescope scanned in azimuth at a fixed elevation of either 37°, 47°, 67° or 77°. It took the telescope $\sim$ 90 s to complete a full 360° sweep of the sky and, typically, every two sweeps the noise diode was fired (for some periods during the observing run it was fired less frequently). The noise diode events are highlighted *red*. The bright spikes in antenna temperature are Galactic plane crossings.

3.1.1 Pointing correction

There are small discrepancies between the actual direction that the telescope was pointing and the direction that the encoders recorded it was pointing. If these pointing errors are large they will distort the final C-BASS map by smearing out features in the sky. The discrepancies are caused by several effects including:

- Deformation of the telescope structure as a function of elevation because of its own weight.
- Zero point offsets of the encoders from their nominal values.
- The azimuth and elevation axis drives may be non-orthogonal and not aligned with their nominal directions.

The global pointing model accounts for these effects and maps the encoder azimuth and elevation to the actual azimuth and elevation. The global pointing model is

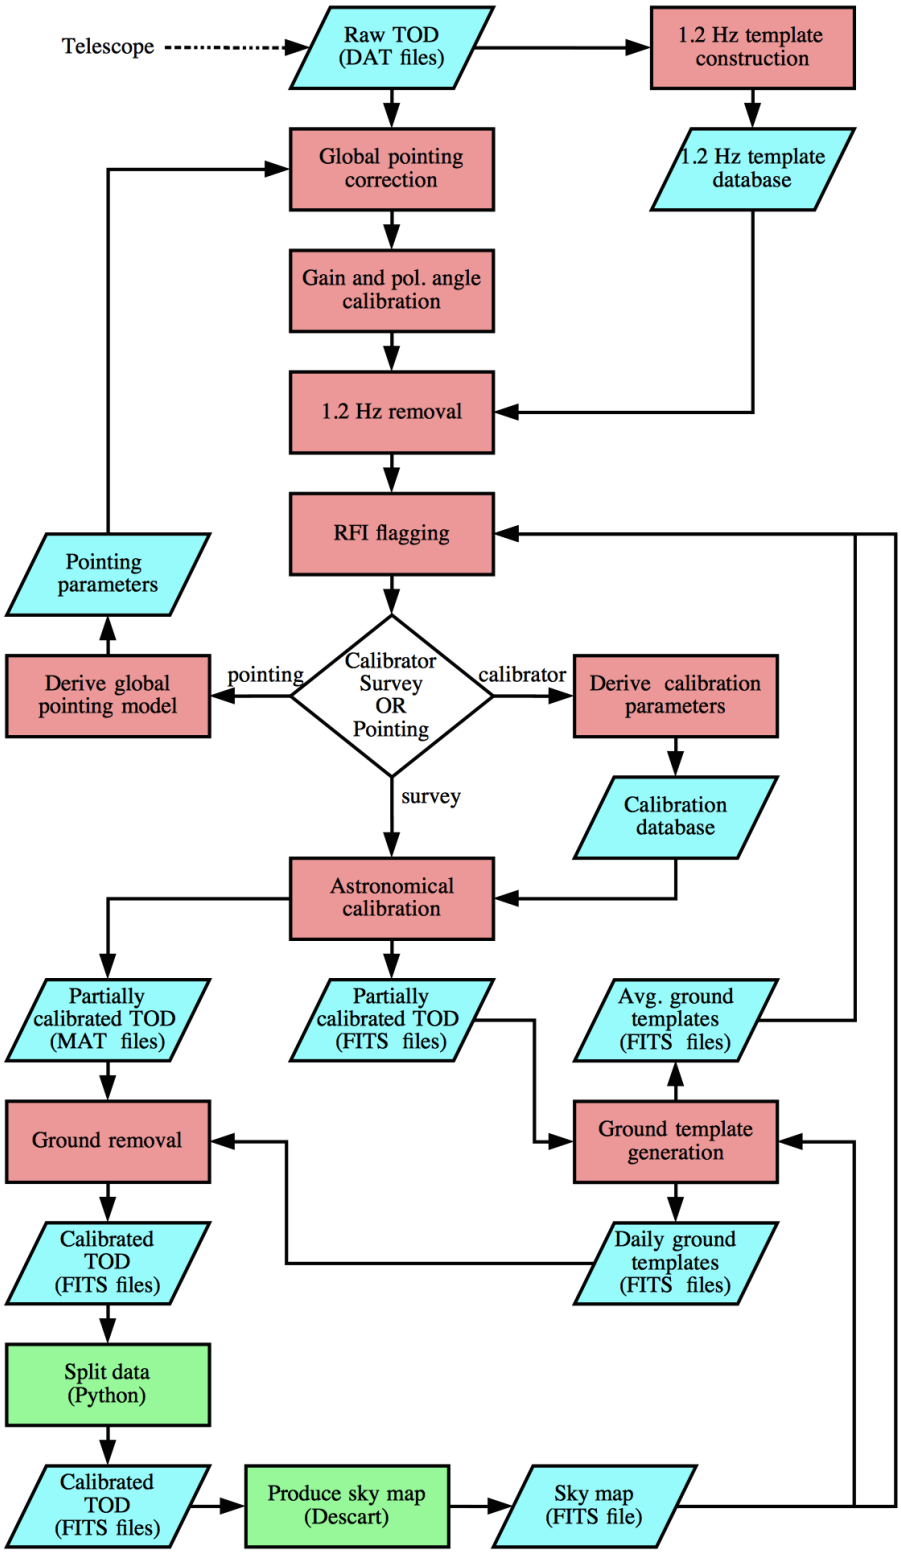

Figure 3.1: Flow diagram indicating the key steps in the C-BASS data pipeline. The *red squares* are MATLAB subprocesses, the *green squares* are mapmaking steps and the *cyan parallelograms* are either inputs or outputs.

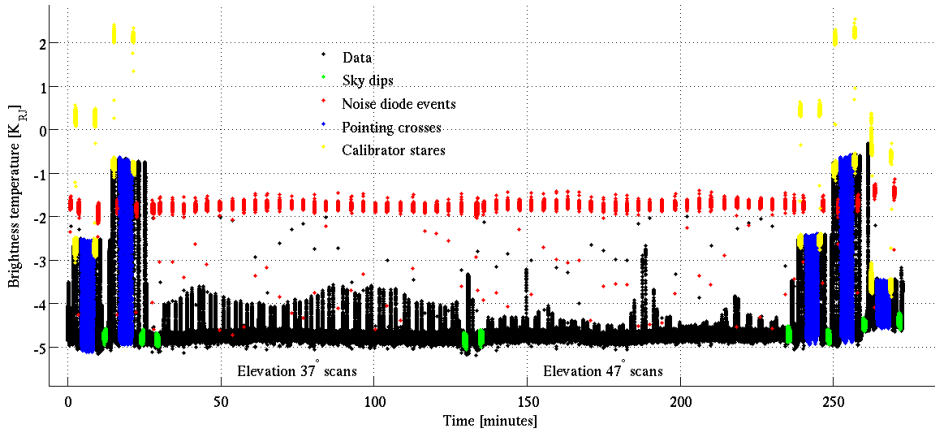

Figure 3.2: Partially calibrated Stokes I TOD from an example set of C-BASS North observations. The samples are colour coded to identify different stages of the observation.

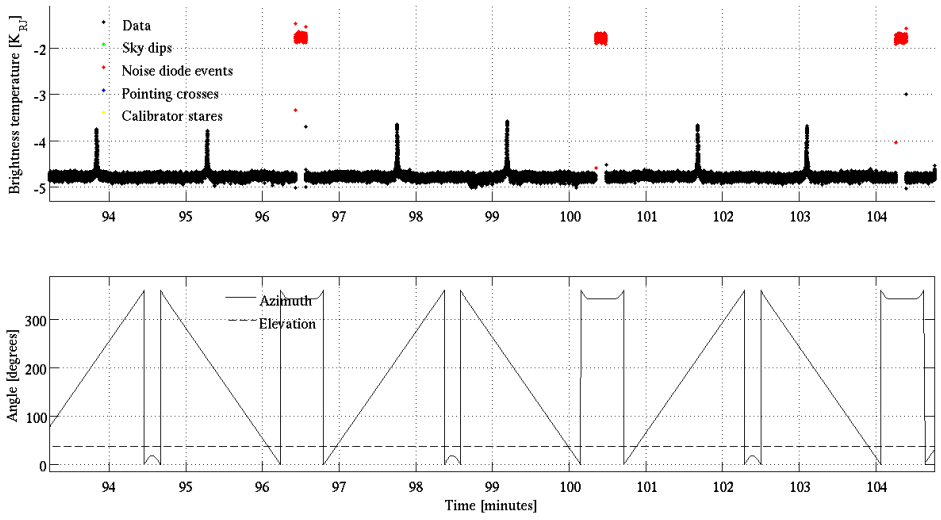

Figure 3.3: Partially calibrated Stokes I TOD (*top*) and telescope pointing (*bottom*) of survey scans from the example C-BASS North observations in Figure 3.2. The samples are colour coded to identify different stages of the observation.

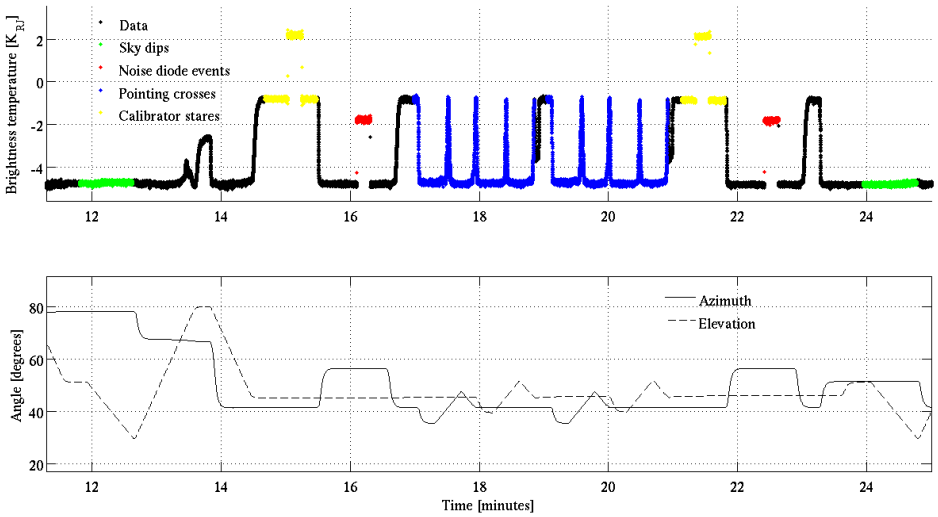

Figure 3.4: Partially calibrated Stokes I TOD (*top*) and telescope pointing (*bottom*) of calibrator observations from the example C-BASS North observations in Figure 3.2. The samples are colour coded to identify different stages of the observation.

stored in a database that is periodically updated. This subroutine applies the pointing corrections to the TOD and also tests its accuracy.

The global pointing model is tested by fitting the C-BASS beam to pointing crosses of bright point sources.[2] The pointing crosses in Figure 3.2 are highlighted *blue* and, from *left* to *right*, are of CygA, CasA, CygA, CasA and then finally of the radio point source J0225+621. A zoom in to part of the calibration observations is shown in Figure 3.4, it includes the first pointing cross observations of CasA. During a pointing cross the telescope scans through a radio point source in azimuth at fixed elevation and then in elevation at fixed azimuth. The C-BASS beam is fitted to these data and the distance between the known location of the point source and the position recovered by the fit is called the pointing offset. If the pointing offset is larger than $5'$ then the global pointing model needs updating. The pointing model is saved for use in the astronomical calibration subroutine (Section 3.1.6).

3.1.2 Gain and polarization angle calibration

Gain and phase drifts in the raw data need to be monitored and corrected for. Gain drifts affect the amplitudes of all the measured Stokes parameters and relative phase

[2]We use Taurus A (TauA), Cygnus A (CygA), Cassiopeia A (CasA), Virgo A(VirA), M42 and J0225+621 as pointing calibrators.

drifts between RCP and LCP will change the apparent polarization angle. In both C-BASS systems this correction is achieved using a temperature-stabilised noise diode that provides a fixed-amplitude reference signal in both intensity and polarization (King, 2009). During survey observations the noise diode was typically fired for a few seconds at the start of alternate 360° sweeps of the sky.

This subroutine scales the amplitudes of the noise diode events to unity and then interpolates the multiplicative gain factors across the TOD. The noise diodes provide gain monitoring on the timescale of minutes and allow us to remove long timescale drifts. After this step in the pipeline the TOD is calibrated in "noise diode units".

In the northern system the phase of the injected noise diode signal is set so that it should appear only in Stokes Q and not Stokes U. A Mueller matrix is applied to rotate the nominal Q and U channels so that the noise diode only appears in the Q channel.[3]

3.1.3 RFI flagging

Transient RFI is caused by human activities. For example passing aeroplanes, mobile phones and wifi from the nearby town of Bishop all emit radiation in the C-BASS band and this radiation contaminates our data. The RFI flagging subroutine attempts to identify these events so that we can exclude them from further analysis.

To operate most effectively, the RFI flagging subroutine requires outputs from previous pipeline runs to use as inputs. Without these it runs in deglitching mode where it simply identifies unphysical spikes in the data with a σ-cut and flags them as contaminated. The deglitching runs on a filtered copy of the data where structure below 4.45 Hz is removed. The cut removes samples that are four times the standard deviation of the filtered data away from its mean value. The flagged samples are excluded from all subsequent analysis. If outputs from previous pipeline runs are available then this subroutine performs a more sophisticated analysis that is described in detail in Heilgendorff et al. (2017, in prep.) and summarised below.

It is not straight forward to identify RFI events with a simple σ-cut in I data because the signal-to-noise of astronomical sources can be very high. In the more sophisticated analysis, the RFI flagger compares continuous chunks of TOD, $\mathbf{d}$, to a

[3]There is currently a software bug in the rotation applied by the Mueller matrix and the noise diode events are rotated into the nominally calibrated U channel. The nominally calibrated Q channel is in reality $-U$ and the nominally calibrated U channel is in reality Q. This error only affects the recovered polarization angles, which we do not make use of in this thesis. The polarized intensities are still correct. The results in this thesis are not materially affected by this error, any reference to Stokes Q should be read as $-U$ and reference to Stokes U should be read as Q.

model time stream, $\mathbf{d}_{\rm model}$. Any samples that are significantly discrepant from the model are flagged as contaminated and excluded from all further analysis. $\mathbf{d}_{\rm model}$ consists of an astronomical sky signal, $\mathbf{d}_{\rm sky}$, and a scan synchronous ground signal, $\mathbf{d}_{\rm ground}$.

$\mathbf{d}_{\rm sky}$ is a function of the telescope pointing in celestial coordinates, RA & DEC, and is generated by extracting values from a sky map, initially produced from C-BASS data that has been run through the pipeline with the RFI flagger in deglitching mode. $\mathbf{d}_{\rm ground}$ is also a function of the telescope pointing but in horizontal coordinates, AZ & EL, and is generated by extracting values from daily average ground templates (Section 3.1.4), again initially produced with the RFI flagger in deglitching mode.

$\mathbf{d}_{\rm model}$ takes the form

$$\mathbf{d}_{\rm model}(t) = G_{\rm sky}\mathbf{d}_{\rm sky}({\rm RA}, {\rm DEC}) + G_{\rm ground}\mathbf{d}_{\rm ground}({\rm AZ}) + A \tag{3.1}$$

where $G_{\rm sky}$ and $G_{\rm ground}$ are multiplicative gain factors and A is a constant offset. Equation 3.1 is fitted to $\mathbf{d}$ in a least-squares sense. Initially 6-σ outliers are flagged and the process is iterated until the residuals have an acceptable skew and kurtosis. The maximum acceptable absolute skewness is 0.1 and maximum acceptable absolute kurtosis is 4.0, these values have been tuned for C-BASS data. No RFI flagging is performed close to the brightest astronomical point sources, as sub-pixel fluctuations in the sky are significant and $\mathbf{d}_{\rm sky}$ will be a poor approximation of the true sky signal. In future work we will fit the C-BASS beam to these bright point sources and use the fitted values to generate $\mathbf{d}_{\rm sky}$ in these regions.

This is an iterative process and once the whole pipeline has been run with the RFI flagger in deglitching mode the output sky map and ground templates can be used as inputs to this subroutine.

3.1.4 Ground removal

Thermal emission from the ground is picked up in the telescope's far-out sidelobes. Due to the scan strategy this emission appears as a scan synchronous signal that varies slowly in time with the changing emissivity of the ground and has a typical amplitude $< 10\,{\rm mK}$. We attempt to reduce this contamination by constructing ground templates and subtracting these from the TOD in an iterative process. Like the RFI flagging subroutine, the ground-subtraction subroutine is most effective if the outputs from previous runs of the pipeline can be used as inputs. The ground-subtraction subroutine is described in detail and tested in Figueras (2016).

Because the ground signal is weak compared to the bright calibrator sources, we only attempt to remove the ground contamination from survey scans and not from the calibration observations. Ground templates are produced from the survey scans, treating each elevation separately, by binning the TOD into 360 azimuth bins and masking out data within a Galactic mask that excludes the brightest 5% of the sky. The daily rotation of the celestial sphere should cause the faint sky signal in each azimuth bin to average away and the stronger sky signals, which will not, average away, are excluded by the Galactic mask. Ground templates are constructed for each day and these daily templates are subtracted from the corresponding TOD. The choice of constructing the ground templates on a daily basis is a compromise between the signal-to-noise ratio of the template and the stability of the ground signal with time.

Once an initial map has been produced from reduced C-BASS data with this level of ground subtraction, it can be used to subtract an estimate of the sky signal from the TOD before binning to produce ground templates with less sky signal contamination.

This method cannot distinguish between the average sky signal and average ground signal at each declination, which are degenerate. This is further justification for masking the brightest regions of the sky when constructing the ground templates, to ensure that the average un-masked sky signal at each declination is close to zero.

The I, Q and U average ground templates for thirty consecutive days of C-BASS North survey scans are shown in Figure 3.5. The ground templates do not vary significantly over the thirty day period. The variation that is present is caused by changes in the ground emissivity, which is likely due to the changing ground temperature and moisture levels.

The overlap region between the Northern and Southern surveys allows for a cross check of the ground subtraction and will be used as a final check on the subtraction once the Southern data are fully available. Each telescope should experience vastly different ground profiles due to both the local topography and climatic conditions. The Northern telescope is in a valley and the Southern telescope is in a relatively flat desert at the highest point for 10–100 km.

3.1.5 1.2 Hz contamination

There is a periodic contaminant in the CBASS North TOD at a frequency of ~ 1.2 Hz and its harmonics. The source of this 1.2 Hz signal is believed to be microphonic vibrations of the receiver components caused by the pumping action of the cold head (the refrigeration cycle is 1.2 Hz). The contamination is not at exactly 1.2 Hz as the

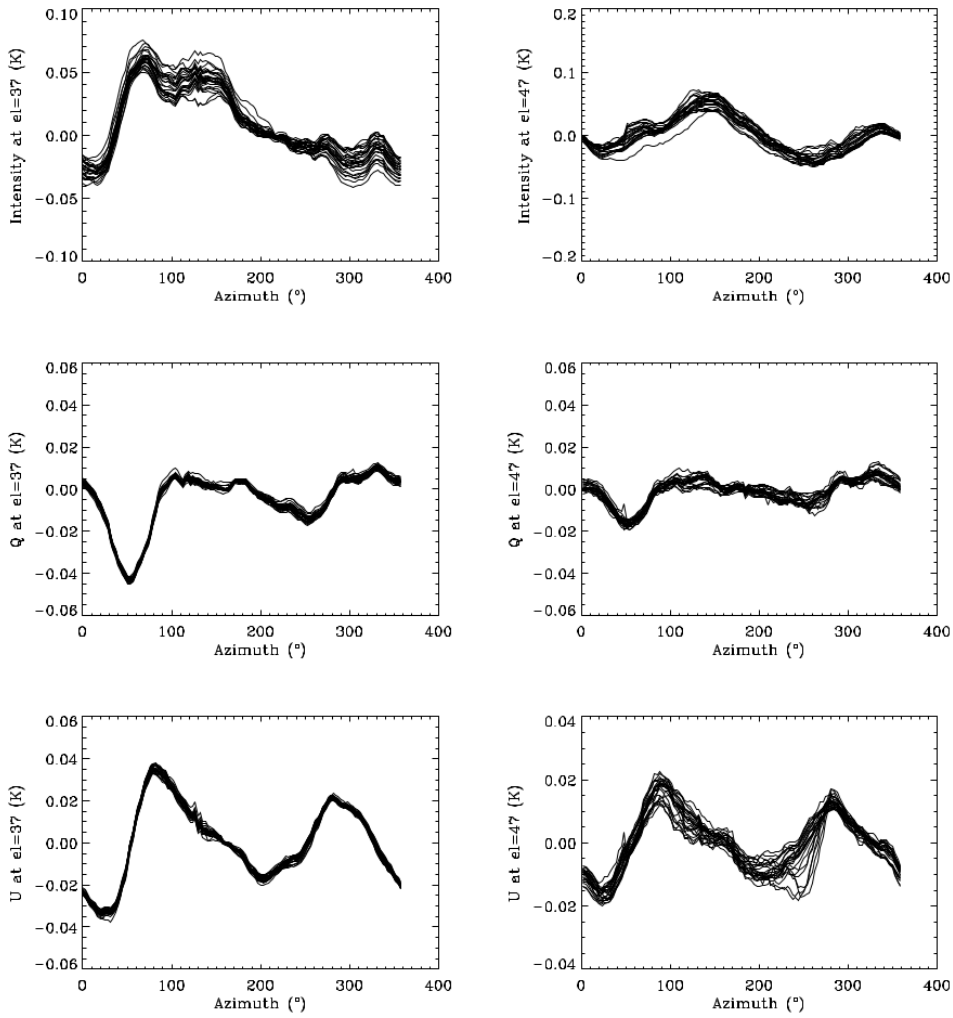

Figure 3.5: The average ground templates for thirty consecutive days, constructed from C-BASS North survey scans. The plots show the I, Q and U templates from *top* to *bottom*. The *left* column of ground templates are constructed from elevation 37° scans, the ground templates in the *right* column are from elevation 47° scans. Figure reprinted from Irfan (2014), with permission of the author.

cold head motor cycles once every 50 power grid cycles and therefore any frequency or phase fluctuations in the mains power will change the frequency of the contaminating signal. The power grid at C-BASS North typically has a frequency of 60 Hz to within a few parts per million (Fink & Beaty, 2000). The frequency of the mains power is adjusted by up to 20 mHz to keep AC synchronous clocks aligned with UTC (NAESB, 2015).

The astronomical signal in C-BASS TOD appears at frequencies below ~ 20 Hz and simply filtering out a band around 1.2 Hz would remove the astronomical signal on

the corresponding angular scales. To remove this contamination while preserving the sky signal, templates of the 1.2 Hz signal are fitted to the TOD and then subtracted from it. This method is extensively documented in Stevenson (2013).

The templates are constructed by folding the TOD onto itself with a timescale of $\frac{1}{1.2\,\mathrm{Hz}} \sim 0.83\,\mathrm{s}$. The data are folded over time periods where the phase of the contaminating signal was constant. These templates are saved to a database. The contamination is removed by loading the templates from the database and fitting them in amplitude, phase and frequency to the data before subtracting them.

To ensure that any residual 1.2 Hz contamination does not always affect the same angular scales, we scan the telescope at multiple speeds to help average the contamination away.

3.1.6 Astronomical calibration

For the C-BASS data to be useful in scientific analysis it must be calibrated from noise diode units to a temperature scale, ultimately referenced to an astronomical source. The calibration is estimated using the calibrator stares and pointing crosses that are highlighted *yellow* and *blue* respectively in both Figure 3.2 and the zoomed in Figure 3.4.

During a calibrator stare the telescope is pointed at the nominal position of the calibrator, as measured by the azimuth and elevation encoders, and the noise diode is fired. The astronomical calibration subroutine calculates the antenna temperature of the amplitude of the noise diode event in Kelvin from Stoke I data. Appendix A contains a definition of antenna temperature and its relationship to the brightness temperature of the sky. The scaling from noise diode units to astronomical units is assumed to be the same in Q and U as in I, this is a reasonable approximation but likely wrong in detail because of the separate hardware chains for each Stokes parameter.

Because of small pointing offsets (estimated in the pointing correction subroutine), the telescope will not have been pointed at the exact centre of the point source. The expected antenna temperature at this distance from the point source is calculated using the C-BASS beam and the peak flux of the calibrator from the literature.

We use Taurus A (TauA) as our primary calibrator because it is visible from both observing sites, it is the brightest 5 GHz source that is unresolved by the C-BASS beam and is highly polarized. When TauA is not visible other astronomical sources are used such as the supernova remnant CasA, as shown in the zoomed in plot of the calibrator observations (Figure 3.4).

To accurately calibrate the data the atmospheric opacity must also be accounted for using

$$T = T' e^{\tau_0 / \sin(\theta_{\rm el})}, \tag{3.2}$$

where T is the atmosphere corrected antenna temperature, T' is the measured antenna temperature, τ_0 is the opacity at zenith and $\theta_{\rm el}$ is the elevation angle. τ_0 is found by fitting Equation 3.2 to data from sky dips, which are highlighted *green* in Figure 3.2. During a sky dip the telescope slews in elevation towards the ground at a fixed azimuth. Typically $\tau_0 \sim 0.7\%$ and so this correction is small.

Sky dips were made after each set of calibration observations to correct for the atmospheric opacity at the time the observations were made. The scalings from noise diode units to atmosphere corrected antenna temperature in Kelvin are saved in a database.

Further sky dips were made at the start and end of each set of survey scans (roughly every 100 minutes). These are used to correct for the atmospheric opacity at the time the survey scans were made. To calibrate the survey scans, the scalings from noise diode to antenna temperature are loaded from the calibration database and interpolated across the TOD - using the opacity from the survey sky dips to correct for the atmosphere.

3.1.7 Pipeline summary

The raw DAT files are automatically run through the pipeline. First only the calibration data are reduced to produce the calibration database. Then the survey data are run through the pipeline in an iterative fashion to improve the RFI flagging and ground subtraction.

The MATLAB pipeline outputs the fully calibrated survey TOD in FITS files, which are used as inputs to the mapmaking steps (Section 4.1). The split data routine, written in Python, splits the scans into continuous chunks and saves the start and end times of each chunk in a table extension in the reduced data FITS files. The tables also leave space in which to record the noise properties of each chunk. A scan will be split if, for example, samples are flagged because they are contaminated by RFI or the telescope was pointing too close to the sun or moon. The longest continuous chunks of data are therefore scans separated by noise diode events and last roughly 180 s but could be arbitrarily short. The mapmaker outputs I, Q and U sky maps along with covariance matrices, again in FITS files.

Table 3.1: The subsets of C-BASS observations that are used in the rest of this thesis.

Name	Daytime/nighttime	Elevation	Sun cut
E37S60	both	37°	60°
E37N	nighttime	37°	–
E37D	daytime	37°	0°
E37DS60	daytime	37°	60°
E37DS90	daytime	37°	90°
E3747S60	both	37°, 47°	60°

3.2 C-BASS data subsets

C-BASS North made survey observations between November 2012 and April 2015, with breaks for maintenance. Between November 2012 and the middle of July 2014 the telescope surveyed the sky with scans at elevations 37° and 47°. Between the middle of July 2014 and April 2015 the telescope surveyed the sky with scans at elevations 67° and 77°. The higher elevation scans significantly increased the depth of observations at intermediate declinations.

Survey observations were made all day and all night. During the day, the Sun is the brightest astronomical source at 5 GHz by a significant margin. Even when the sun is in the telescope's far out sidelobes it is still bright enough to significantly contaminate the C-BASS data. The RFI subroutine (described in Section 3.1.3) identifies data where the Sun was in the telescope's main beam or first few sidelobes. We apply an extra level of Sun flagging during the split data subroutine based on an angular distance cut. We initially thought that flagging data within a 60° radius around the Sun would be sufficient but the Jackknife tests in Section 4.4.2.2 we show that this is not the case. We therefore use maps made only from nighttime scans for scientific analysis.

In this thesis we use various subsets of the C-BASS data, summarised in Table 3.1, for different parts of the analysis. The most restrictive subset, E37N, is used for scientific analysis.

3.3 TOD noise properties

An accurate description of the noise on calibrated TOD is vital when attempting to produce sky maps with well understood errors. The noise properties are expected to change with time, possibly caused by the receiver components ageing, telescope repairs and the weather. This section sets out how we estimated the noise properties

of C-BASS data by fitting a noise model to periodograms of short chunks of the data. These noise properties can then be passed to the mapmaker to improve the final maps. They can also be used to measure the level of residual 1.2 Hz contamination in the C-BASS data.

Two different parametrizations of the noise model are described in Section 3.3.1. The statistical properties of the periodogram are set out in Section 3.3.2. In Sections 3.3.3 and 3.3.4 we describe two different methods of fitting a noise model to C-BASS data and present the noise parameters that they recover.

The first method (Section 3.3.3) recovers estimates of the noise parameters by fitting a noise model in a least-squares sense to binned log-periodograms of the TOD. It requires long continuous chunks of TOD, which we get from dedicated NCP stare observations.

The second method (Section 3.3.4) uses Bayesian inference and MCMC methods to recover the posterior distributions of the noise parameters. It works reliably on far shorter chunks of data than the first method and therefore can be used to estimate the noise properties from the survey scans and doesn't require dedicated observations. We test this method on simulated E3747S60 data and use it to estimate the noise properties of the E37S60 dataset.

3.3.1 Noise model

The noise on data from total power radio telescopes, such as C-BASS, typically has both a white (equal power at all frequencies) and a red ($1/f$, excess noise power at lower frequencies) component. Receiver electronics produce white noise via the thermal agitation of electrons. Extra $1/f$ noise is introduced by both drifts in the receiver gain and in the emissivity of the atmosphere. Noise comprised of these two components is most often characterised by a three parameter model,

$$P_{xx} = \sigma_w^2 \left(1 + \left(\frac{\omega}{\omega_{\text{knee}}}\right)^{-\alpha}\right), \tag{3.3}$$

where P_{xx} is the power spectrum of the noise, σ_w is the white noise level, ω_{knee} ($\omega = 2\pi f$) is the knee frequency at which the $1/f$ and white components have equal power and α is the power law that typically lies between 1 and 2. Note that P_{xx} is normalised such that its mean value is equal to the variance of the data.

This parametrization makes it difficult to distinguish between changes in the white and red components. For example, if σ_w changes then so will ω_{knee} even if the amount

of $1/f$ noise in the signal has remained constant. Instead, Equation 3.3 can be re-parametrized to

$$P_{xx} = \sigma_w^2 + \sigma_r^2 \left(\frac{\omega}{\omega_0} \right)^{-\alpha}, \tag{3.4}$$

where σ_r is the red noise level at a reference frequency ω_0. This parametrization allows for a simpler comparison of the noise properties over time as the white and $1/f$ components have been disentangled.

3.3.2 Statistical properties of the periodogram

The periodogram, I_N, is an estimator of the true underlying power spectrum, P_{xx}, of TOD and it can be quickly calculated using Fourier transforms. To accurately fit models to the periodogram requires knowledge of its statistical properties, which are summarised in this section. The main features of the periodogram estimator, which will be derived below, are that:

- I_N is a biased estimator of P_{xx} but in most cases this bias is small.
- $I_N(\omega)$ is not distributed normally about $P_{xx}(\omega)$.
- Binning I_N provides an estimate of P_{xx} with a distribution is closer to Gaussian.
- Smaller bins are needed when working in logarithmic space.
- There is an additive bias between $\ln(I_N)$ and $\ln(P_{xx})$.

An infinitely long, noise dominated timestream can be considered as a stationary random process $\{x_n\}$ where $-\infty < n < \infty$. The true underlying power spectrum, P_{xx}, is related to the true underlying autocorrelation function, ϕ_{xx}, via

$$P_{xx}(\omega) = \sum_{m=-\infty}^{\infty} \phi_{xx}(m) e^{-j\omega m}. \tag{3.5}$$

Continuous chunks of C-BASS TOD are of finite length, so we can only calculate an estimator of the autocorrelation function from a sequence of N consecutive elements of $\{x_n\}$. The biased autocorrelation estimator, c_{xx}, is defined as

$$c_{xx}(m) = \frac{1}{N} \sum_{n=0}^{N-|m|-1} x(n)x(n+m). \tag{3.6}$$

with an expectation value

$$\mathrm{E}\left[c_{xx}(m)\right] = \frac{N-|m|}{N} \phi_{xx}(m). \tag{3.7}$$

I_N can be found from c_{xx} by [4]

$$I_N(\omega) = \sum_{m=-(N-1)}^{N-1} c_{xx}(m) e^{-j\omega m}. \tag{3.8}$$

The expectation value of I_N is

$$\mathrm{E}\left[I_N(\omega)\right] = \sum_{m=-(N-1)}^{N-1} \mathrm{E}\left[c_{xx}(m)\right] e^{-j\omega m} \tag{3.9}$$

$$= \sum_{m=-(N-1)}^{N-1} \left(\frac{N-|m|}{N}\right) \phi_{xx}(m) e^{-j\omega m}. \tag{3.10}$$

Due to the bias of c_{xx} and the finite limits of the summation, I_N is a biased estimator of P_{xx}.

Equation 3.10 can be interpreted as the Fourier transform of a windowed auto-correlation sequence (Oppenheim & Schafer, 1975). The window, $w(m)$, is known as either the Triangular, Fejer, or Bartlett window. The window $w(m)$ and its Fourier transform, $W(m)$, are defined as

$$w(m) = \begin{cases} \frac{N-|m|}{N} & |m| < N \\ 0 & \text{else} \end{cases} \tag{3.11}$$

$$W(\omega) = \mathrm{FT}\left[w(m)\right] = \frac{1}{N} \frac{\sin^2(\omega N/2)}{\sin^2(\omega/2)}, \tag{3.12}$$

which gives

$$\mathrm{E}\left[I_N(\omega)\right] = \frac{1}{2\pi} \int_{-\pi}^{\pi} P_{xx}(\theta) W(\omega - \theta)\, d\theta. \tag{3.13}$$

Due to this convolution, $I_N(\omega)$ includes contributions from P_{xx} at frequencies other than just ω. The size of this effect depends on the shape of P_{xx}. For example, a true red-coloured spectrum would have power transferred from lower to higher frequencies and vice versa for a blue-coloured spectrum. As $N \to \infty$, $W(\omega)$ gets narrower and therefore I_N becomes less biased. In general, if $c_{xx} \simeq 0$ for $m << N$ then the bias is not significant. In the rest of this work I will assume that this is the case and that I_N is an unbiased estimator of P_{xx}.

For zero-mean Gaussian noise the ratio of $I_N(\omega_k)$ to $P_{xx}(\omega_k)$ obeys a chi-squared distribution with two degrees of freedom

$$\frac{2I_N(\omega_k)}{P_{xx}(\omega_k)} \sim \chi^2_2, \tag{3.14}$$

[4] From the Wiener-Khinchin theorem this is equivalent to $I_N(\omega_k) = \frac{1}{N}|X(\omega_k)|^2$ where $X = \mathrm{FT}(x)$.

where $\omega_k = k2\pi/N$ for $k = 0, 1, ..., N/2 - 1$[5]. This result can be quickly derived for white noise with variance σ_x^2 by considering the Sine and Cosine transforms of $\{x_N\}$, $A(\omega_k)$ and $B(\omega_k)$ respectively.

$$I_N(\omega_k) = A^2(\omega_k) + B^2(\omega_k). \tag{3.15}$$

Since A and B are independent random variables with normal distributions centred on zero with standard deviation $\sigma_x/2$,

$$\frac{I_N(\omega_k)}{\sigma_x^2/2} = \left(\frac{A(\omega_k)}{\sqrt{\sigma_x^2/2}}\right)^2 + \left(\frac{B(\omega_k)}{\sqrt{\sigma_x^2/2}}\right)^2 \tag{3.16}$$

$$\frac{2I_N(\omega_k)}{\sigma_x^2} \sim \chi_2^2. \tag{3.17}$$

Substituting $P_{xx}(\omega) = \sigma_x^2$ gives Equation 3.15. This result still roughly holds for processes that do not have a flat power spectrum as long as the aliasing and biasing effects are not significant (Papadakis & Lawrence, 1993). Equation 3.14 then gives

$$\mathrm{E}\left[I_N(\omega_k)\right] = P_{xx}(\omega_k) \tag{3.18}$$

$$\mathrm{VAR}\left[I_N(\omega_k)\right] = P_{xx}^2(\omega_k) \tag{3.19}$$

Equation 3.19 shows that $I_N(\omega_k)$ is not a consistent estimator of $P_{xx}(\omega_k)$ and Equation 3.17 shows that its probability distribution is highly non-Gaussian.[6] I_N can be smoothed or binned to provide an estimator with a smaller variance and a probability distribution closer to Gaussian via the central limit theorem. Non-overlapping bins are preferable to smoothing so as to maintain the independence of the estimate of $P_{xx}(\omega_k)$ for each ω_k.

Papadakis & Lawrence (1993) use simulations of curved red noise to find that bins averaging at least fifty consecutive elements of $I_N(\omega_k)$ are required for the binned estimates to be normally distributed about the true power spectrum values and caution that the bias induced by this binning is significant at low frequencies. They instead suggest working in logarithmic space as $\ln(I_N)$ has a more symmetric probability distribution than I_N and the variance of $\ln(I_N)$ estimates do not depend on P_{xx} (unlike the variance of I_N). Remembering that $\mathrm{E}\left[g(x)\right] \neq g\left(\mathrm{E}[x]\right)$, the expectation value and

[5] At the Nyquist frequency $I_N/P_{xx} \sim \chi_1^2$.

[6] The variance of a consistent estimator tends to zero as N tends to ∞.

variance of $\ln(I_N)$ can be derived from Equation 3.14,

$$\ln\left(I(\omega_k)\right) = \ln\left(P_{xx}(\omega_k)\right) + \ln\left(\frac{\chi_2^2}{2}\right) \tag{3.20}$$

$$\mathrm{E}\left[\ln\left(I(\omega_k)\right)\right] = \ln\left(P_{xx}(\omega_k)\right) + \mathrm{E}\left[\ln\left(\frac{\chi_2^2}{2}\right)\right] \tag{3.21}$$

$$= \ln\left(P_{xx}(\omega_k)\right) - \gamma \tag{3.22}$$

$$\mathrm{VAR}\left[\ln\left(I(\omega_k)\right)\right] = \mathrm{VAR}\left[\ln\left(\frac{\chi_2^2}{2}\right)\right] \tag{3.23}$$

$$= \frac{\pi^2}{6} \tag{3.24}$$

where $\gamma \simeq 0.5772157...$ is the Euler-Mascheroni constant. Note that $\ln(I_N)$ is a biased estimator of $\ln(P_{xx})$ however the bias can be removed by the addition of a constant.

The logarithmic estimates can now be binned to reduce the variance and provide estimators with more Gaussian probability distributions (once the additive constant bias has been corrected for). The simulations in Papadakis & Lawrence (1993) show that the arithmetic mean of twenty consecutive elements of $\ln(I_N)$ with the geometric mean of the corresponding values of ω provide unbiased estimates of $\ln(P_{xx}(\omega))$ that are normally distributed about $\ln(P_{xx}(\omega))$ except at inflection points where a small bias is introduced.

3.3.3 Least-squares method

The method of least-squares is a fast and commonly used method of fitting models to data that are subject to Gaussian random errors in the y ordinates. To obtain estimates of the noise parameters, the noise model can be fitted in a least-squares sense in log space to binned periodograms of the time ordered data, accounting for the additive bias. The fitting should only be done up to frequencies of order $10 \times f_{\mathrm{knee}}$ to avoid over-weighting the white noise part of the spectrum.

At least two data points in the binned periodogram must be dominated by the $1/f$ component, and therefore at frequencies below f_{knee}, in order to constrain α and σ_r. Before binning this requires at least forty samples of the periodogram to be below f_{knee} and therefore requires continuous stretches of time ordered data lasting longer than $\frac{40}{f_{\mathrm{knee}}}$. For example, if $f_{\mathrm{knee}} = 0.1\,\mathrm{Hz}$ then the $1/f$ component can only be constrained from observations lasting at least 400 s. Normal survey scans are not this long and their noise properties can not be estimated using this method. It can only be used on dedicated observations with continuous chunks of data much longer than a survey scan.

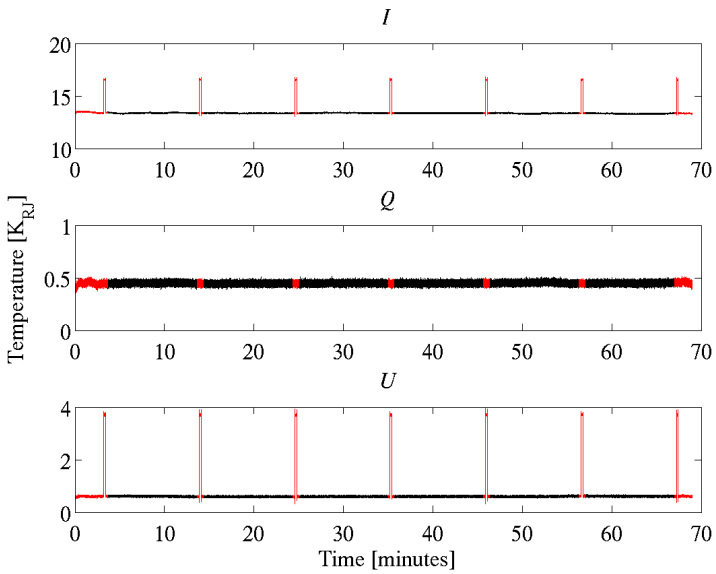

Figure 3.6: Stokes I, Q and U TOD from an example NCP stare. During this observation the noise diode was fired every 10 minutes, the noise diode events are highlighted *red*.

The noise properties can be estimated from NCP stares where the telescope is set to stare at the NCP for arbitrarily long periods of time, firing the noise diode at regular intervals, to determine the noise power down to very low frequencies. Reduced data from an example stare where the noise diode was fired every 10 minutes are shown in Figure 3.6.

The NCP is the point around which the celestial sphere rotates and so the sky signal at this azimuth and elevation does not change with time. Contributions to the $1/f$ signal during these stares will come from both within the receiver and from water vapour in the atmosphere. The receiver architecture is designed to minimise the contributions to the $1/f$ noise from gain fluctuations so atmospheric contributions should dominate. $1/f$ drifts during NCP stares will be smaller than those whilst the telescope is scanning. This is because when the telescope scans across the sky the beam moves through patches of higher/lower density water vapour more quickly than when those patches of water vapour are drifting through the beam due to wind.

NCP stares were made intermittently throughout the Northern observing campaign, the dates on which stares were made are shown in Figure 3.7. Because they are unevenly spread across the campaign, the noise properties during these observations are not expected to be representative of the noise during all survey observations.

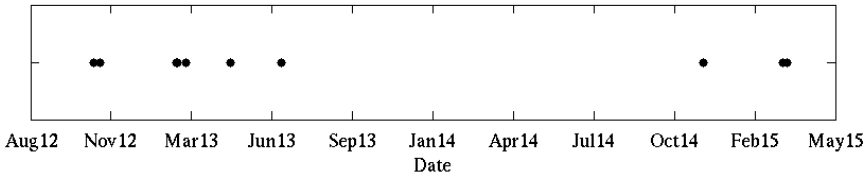

Figure 3.7: The dates of NCP stares made by C-BASS North. They are not evenly spread across the observing campaign.

3.3.3.1 Tests on simulated data

The least-squares method was tested on simulated NCP stare observations by generating 750 realisations of noise and comparing the noise statistics returned by the method to the true values used to generate the TOD. Each realisation was simulated by generating a sequence, y, of 6×10^4 samples of zero-mean white noise with unit variance. The Fast Fourier Transform (FFT) of y was multiplied by the square root of the model P_{xx} and the inverse FFT taken to give 10-minute realisations of the noise sampled at 100 Hz with the correct statistical properties. The noise model in Equation 3.3 was used with $\alpha = 1.5$, $f_{\rm knee} = 0.1$ Hz and $\sigma_w = 0.03$ K. We subtracted the best fit straight line from each continuous stretch of data before calculating its periodogram and fitting the noise model. Figure 3.8 shows histograms of the recovered noise statistics in *grey* and the true simulation values in *red*. The method set out above has recovered unbiased estimates of the noise statistics.

3.3.3.2 Recovered noise statistics from NCP stares

Remembering that the properties of the noise in NCP stares will not be representative of the noise on survey observations, we used the least-squares method outlined above to fit the noise model in Equation 3.3 to these data. Table 3.2 shows the means, medians and standard deviations of the recovered noise statistics for all three Stokes parameters.[7] Using Equation 2.3, assuming a noise equivalent band width of 0.489 GHz and the median values of σ_w, gives a system temperature of 36 K in Stokes I and 31 K in Q and U. These values are in line with the expected system temperature of 32 K (King et al., 2014a). The large spread in parameter values is not too worrying because NCP stares were often made when the telescope was not fully functional and are therefore not representative of the survey data.

[7]The TOD is constructed from the averages of the receiver channels; $I = (I_1 + I_2)/2$, $Q = (Q_1 + Q_2)/2$ and $U = (U_1 + U_2)/2$

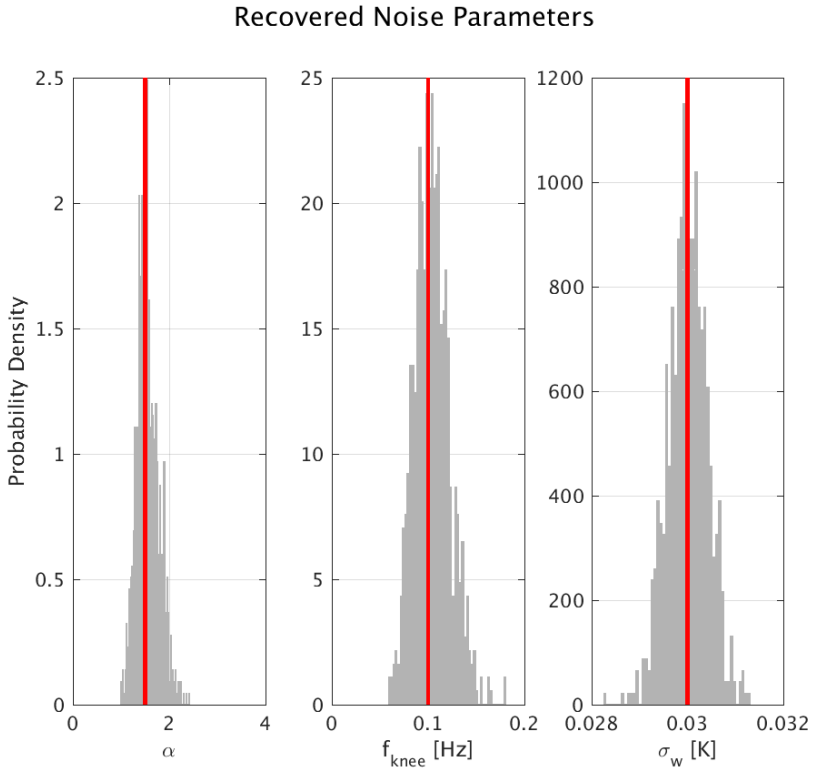

Figure 3.8: Histograms of recovered noise parameters from 750 realisations of $1/f$ plus white noise using the least-squares fitting method. The vertical *red* lines are at the parameter values used to simulate the TOD.

Table 3.2: Means, medians and standard deviations of the estimated noise statistics of NCP stares made with C-BASS North.

Stokes	Parameter	Mean	Median	Standard deviation
	α	2.2	1.6	1.8
I	$f_{\rm knee}$ [Hz]	0.40	0.17	0.46
	σ_w [K]	0.026	0.023	0.010
	α	2.3	1.5	2.3
Q	$f_{\rm knee}$ [Hz]	0.17	0.045	0.28
	σ_w [K]	0.022	0.020	0.012
	α	1.9	1.5	1.9
U	$f_{\rm knee}$ [Hz]	0.29	0.064	0.65
	σ_w [K]	0.022	0.020	0.012

3.3.4 Bayesian method

Ideally we would like to estimate the noise properties directly from the survey scans. There are two extra difficulties in estimating the noise parameters from surveys compared to NCP stares. Firstly, contiguous chunks of data are shorter during surveys than during stares and the least-squares method would not be able to constrain the $1/f$ component from them. A different method that maintains the frequency resolution of the un-binned periodograms is needed. Secondly, the sky signal will change with time as the telescope scans across the sky and an estimate of this sky needs to be subtracted from the I TOD. The sky signal in the Q and U channels is sufficiently weak that they are noise dominated and no sky subtraction is required.

Bayesian methods can be used to fit a noise model to periodograms and recover estimates of the $1/f$ component from these shorter lengths of data. The likelihood function is estimated from the distribution in Equation 3.14 by a change of variables, letting $Y = 2I_N/P_{xx}$, the likelihood function, $p(I_N|P_{xx})$, is found by

$$p(Y|P_{xx})dY = p(I_N|P_{xx})dI_N \tag{3.25}$$

$$p(I_N|P_{xx}) = p(Y|P_{xx})\frac{dY}{dI_N} \tag{3.26}$$

$$= \frac{1}{P_{xx}}e^{-I_N/P_{xx}}. \tag{3.27}$$

Equation 3.27 is sometimes known as the Whittle (quasi-)likelihood (Whittle, 1953). After selecting prior distributions for the parameters in the noise model, the posterior probability can be sampled using MCMC methods (for example the Metropolis-Hastings algorithm). We used the noise model in Equation 3.4 instead of Equation 3.3 to reduce the degeneracy between parameters.

3.3.4.1 Testing choice of prior with simulated data

There are three obvious choices for the prior probability distributions: uniform priors; Jeffreys priors, which for scale parameters are uniform in logarithmic space (Jeffreys, 1946); or informative priors, for example normal or log-normal distributions depending on what information we have about the parameters. Appendix B contains further discussion of Bayesian methods and priors.

We tested this process on simulated data with two choices of prior,

$$p(\alpha)p(\sigma_r)p(\sigma_w) \propto \text{constant} \tag{3.28}$$

and

$$p(\alpha)p(\sigma_r)p(\sigma_w) \propto \frac{1}{\sigma_r \sigma_w}, \tag{3.29}$$

with upper and lower limits to ensure the probability density integrates to unity.[8] The second of which (the Jeffreys prior) was used by Vaughan (2010) when analysing *XMM-Newton* data. In this section we present results from simulations that show that the uniform prior, unlike the Jeffreys prior, provides unbiased estimates of the noise parameters from the shortest chunks of data.

The simulated data was generated by taking roughly half of the E3747S60 dataset (C-BASS North elevation 37° and 47° scans with a 60° sun cut) and overwriting the signal with noise realisations. The white noise level was set to $\sigma_w = 25\,\mathrm{mK}$ and the $1/f$ component was modelled with $\alpha = 1.5$ and $\sigma_r = 25\,\mathrm{mK}\,@\,0.1\,\mathrm{Hz}$. This gave us continuous chunks of TOD with the same lengths as the actual survey data. From this simulated dataset, chunks of data shorter than $\sim 30\,\mathrm{s}$ will not produce reliable estimates of α or σ_r as they won't have two points in the periodogram with frequencies $< f_{\mathrm{knee}} = 0.1\,\mathrm{Hz}$, dominated by $1/f$ noise.

We subtracted the best fit straight line from each chunk before calculating its periodogram. The spectrum above 5 Hz was ignored from the fit to avoid over weighting the white noise dominated region of the spectrum, we also ignored the spectrum between 1.18 and 1.22 Hz where residual contamination from the 1.2 Hz systematic may be present in real data. PyMC was used to sample the posterior probability distribution with the Metropolis-Hastings algorithm.

Histograms of the mean parameter values from the converged MCMC chains, run with both sets of priors, are shown in Figure 3.9. The location of the weighted average parameter values are marked with stars and were calculated using

$$\hat{\mathbf{\Phi}} = \left(\sum_i \boldsymbol{C_i^{-1}} \right)^{-1} \left(\sum_i \boldsymbol{C_i^{-1}} \mathbf{\Phi_i} \right), \tag{3.30}$$

where, for the i^{th} chunk, $\mathbf{\Phi_i} = (\alpha_i, \sigma_{ri}, \sigma_{wi})$ is a vector containing the mean noise statistics for that chunk's converged parameter chains and $\boldsymbol{C_i}$ is the covariance between that chunk's converged parameter chains.

Largely unbiased estimates of all the parameters are recovered from chunks longer than 30 s when using the uniform prior (Equation 3.28). When using the Jeffreys prior (Equation 3.29) the recovered parameters are unreliable on chunks less than 60 s long and the recovered estimates of α are significantly biased even from fairly

[8] $0.25 < \alpha < 4.0$, $0. < \sigma_w < 1.$, $0. < \sigma_r < 1.$.

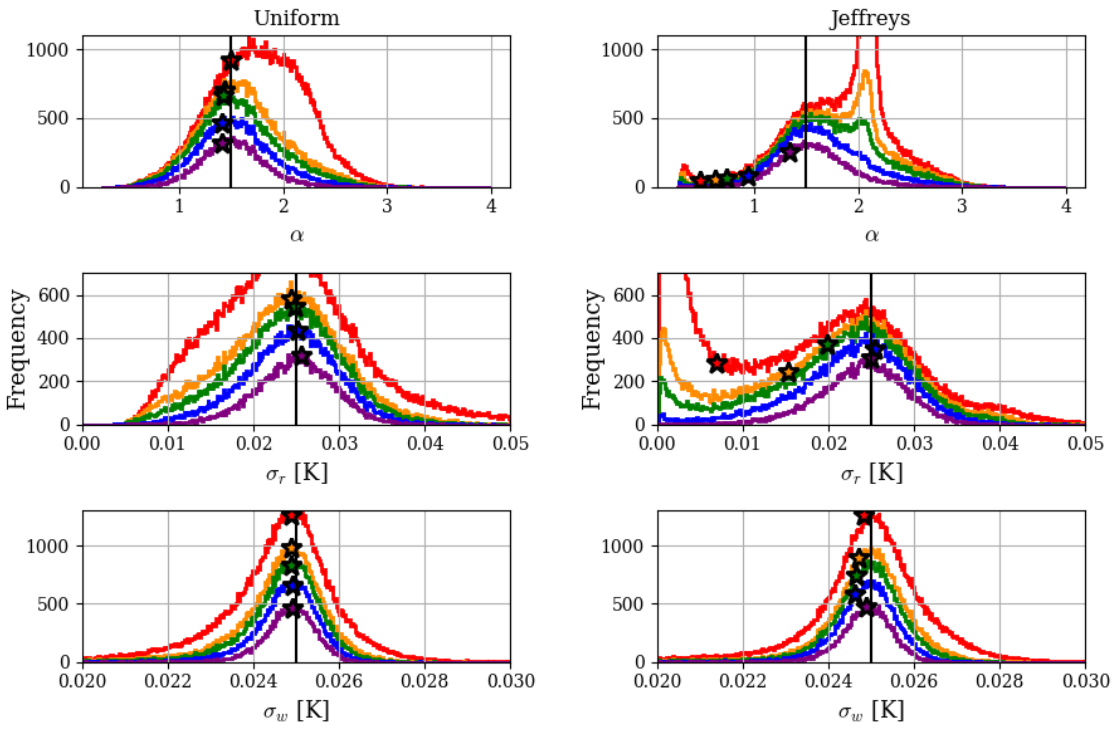

Figure 3.9: Histograms of the mean parameter values of converged MCMC chains from fits to simulated TOD that only included realisations of the noise. *left* using the uniform prior (Equation 3.28), *right* using the Jeffreys prior (Equation 3.29). The *red, orange, green, blue* and *purple* curves are histograms of the parameter estimates from chunks longer than 10, 30, 40, 60 and 100 s respectively. The *stars* are at the weighted averages, accounting for parameter covariance. The vertical *black* lines are at the true parameter values.

long chunks. When the periodogram has a small number of points below $f_{\rm knee}$ the likelihood only weakly constrains the $1/f$ component and the posterior probability is dominated by the prior. The Jeffreys prior drags σ_r to small values and when σ_r is sufficiently small α has the freedom to explore its prior with minimal impact on the posterior. The mean value of α is then halfway between the upper (4) and lower (0.25) limits placed on it.

We therefore use flat priors on all the parameters when fitting the noise model to real TOD. This will enable us to recover unbiased estimates from even short chunks of data. For safety we will ignore the recovered noise statistics from chunks shorter than $\sim 4/f_{\rm knee}$, i.e. we only look at the noise properties from TOD with periodograms having at least three points dominated by the $1/f$ component.

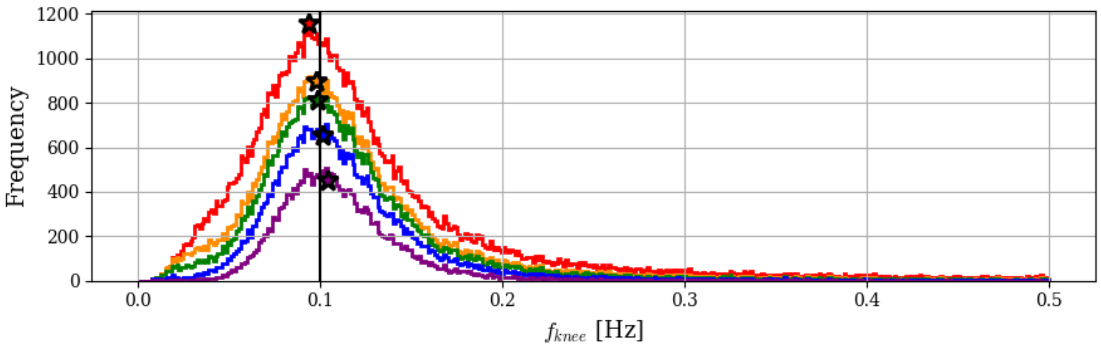

Figure 3.10: Histograms of the mean $f_{\rm knee}$ values calculated from converged MCMC chains from fits to simulated TOD that only included realisations of the noise. The posterior distribution used uniform priors on α, σ_r and σ_w (Equation 3.28). The *red, orange, green, blue* and *purple* curves are histograms of the $f_{\rm knee}$ estimates from chunks longer than 10, 30, 40, 60 and 100 s respectively. The *stars* are calculated from the weighted averages, accounting for parameter covariance. The vertical *black* lines are at the true parameter values.

3.3.4.2 Testing sky subtraction on simulated data

The next difficulty when estimating the noise properties of survey observations is the changing sky signal in the TOD. The polarized sky signal is sufficiently weak compared to the instantaneous sensitivity of the telescope that we do not need to subtract a sky guess from the Q or U timestreams. The total intensity sky signal is much stronger and an estimate of the sky signal must be subtracted from the Stokes I TOD to provide a noise dominated timestream. We remove the Stokes I sky signal using a map made from C-BASS North data. The bilinear interpolation of the four nearest pixels to the sky coordinates of each TOD sample were subtracted from the timestream. This subtraction works poorly when there are significant sub-pixel fluctuations in the sky. If the estimated sky signal changed by more than 25 mK between any two samples then we did not estimate the noise statistics for that chunk.

To test whether the noise parameters could still be accurately recovered from Stokes I survey data when an estimated sky signal has been subtracted, we simulated another dataset and ran it through the process outlined above. As in Section 3.3.4.1 we took roughly half of the E3747S60 dataset (C-BASS North elevation 37° and 47° scans with a 60° sun cut) and overwrote the signal with simulated data. This time the simulation included a sky signal as well as realisations of the noise. The simulated sky signal was generated using the actual pointing of the scans to extract values from a model $N_{\rm side} = 1024$ map smoothed to a Gaussian beam with FWHM

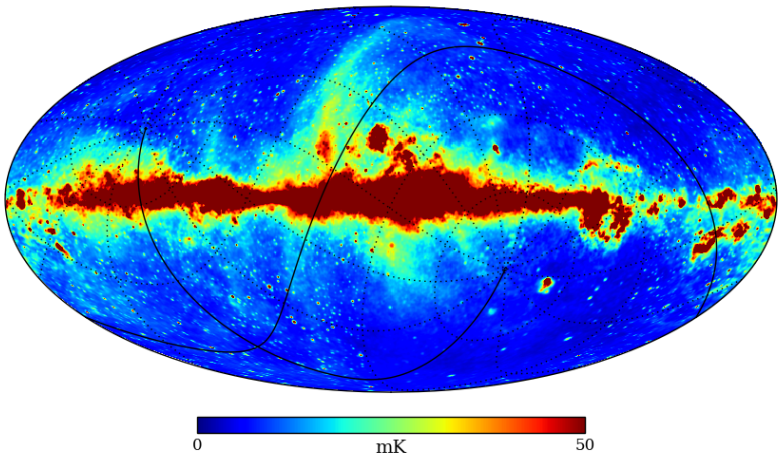

Figure 3.11: Model 5 GHz sky map including diffuse Galactic emission and Northern point sources. The graticules are lines of constant RA and DEC.

0.70° (Figure 3.11). The model sky map was simulated using PySM to generate the 5 GHz diffuse emission (Thorne et al., 2016), point sources were added from the GB6 (Gregory et al., 1996) and S5 (Kuehr et al., 1981) catalogs. The point sources M87, W51, Tau A, Cyg A and Cas A had to be added by hand as they are excluded from the GB6 catalog, their values were taken from Gregory et al. (1996); Baars et al. (1977); Aller & Reynolds (1985); Hafez et al. (2008). We used the same noise parameters as in Section 3.3.4.1 to generate realisations of the noise, the white noise level was set to $\sigma_w = 25\,\mathrm{mK}$ and the $1/f$ component was modelled with $\alpha = 1.5$ and $\sigma_r = 25\,\mathrm{mK}\,@\,0.1\,\mathrm{Hz}$.

We ran this simulated data through the Descart mapmaker to produce an $N_{\mathrm{side}} = 256$ map. This map was used to generate the expected sky signal that we the subtracted from the simulated TOD before estimating the noise statistics. Mapmaking and Descart are discussed in the following chapter.

Recovered parameters from simulated survey scans From each chunk's converged MCMC chains we recorded the means of the marginalised chains, their covariances, the mean f_{knee} and its variance. Histograms of the recovered parameter values are shown in Figure 3.12 and the calculated knee frequencies in Figure 3.13. The residual sky signal in the TOD has biased the estimates of α and σ_w low and σ_r high. For example, the weighted average noise statistics from chunks longer than 40 s deviate from the true values by $\hat{\alpha} - \alpha^{\mathrm{true}} = -0.16$, $\hat{\sigma_r} - \sigma_r^{\mathrm{true}} = 0.58\,\mathrm{mK}$ and $\hat{\sigma_w} - \sigma_w^{\mathrm{true}} = -0.16\,\mathrm{mK}$. These biases are small compared to the spread of the recov-

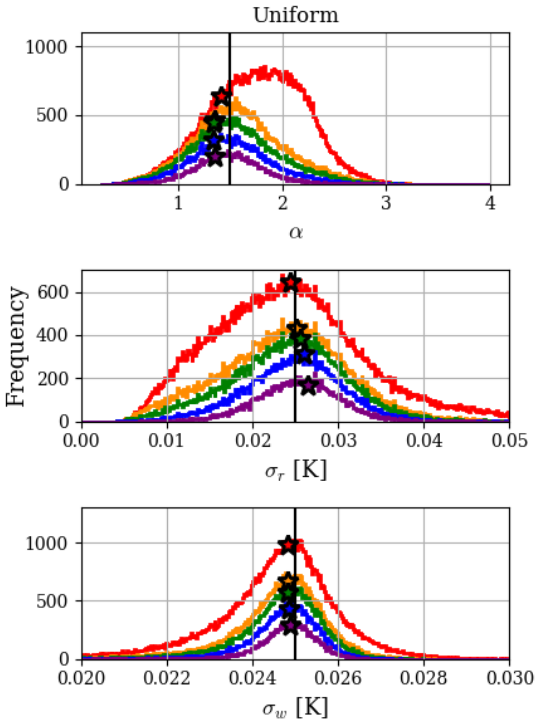

Figure 3.12: Histograms of mean parameter values from converged MCMC chains fitted to simulated TOD including residual sky and noise realisations. The *red, orange, green, blue* and *purple* histograms are of the parameter estimates from TOD chunks longer than 10, 30, 40, 60 and 100 s respectively. The *stars* are at the weighted averages, accounting for parameter covariance, of each histogram. The vertical *black* lines are at the true parameter values.

ered statistics and deemed acceptable. In real data we expect the bias to be similar in magnitude but not the same, as the true sky differs from the model used.

1.2 Hz contamination, null test To estimate the level of residual 1.2 Hz contamination in the TOD we record the average periodogram power between 1.18–1.22 Hz and 1.68–1.72 Hz, $I_N^{1.2}$ and $I_N^{1.7}$ respectively. In the simulated data there is no contaminating signal at 1.2 Hz and so neither $I_N^{1.2}$ nor $I_N^{1.7}$ should be systematically biased above the fitted model spectrum, $P_{xx}^{1.2}$ at 1.2 Hz and $P_{xx}^{1.7}$ at 1.7 Hz. If there were a contamination at 1.2 Hz then $I_N^{1.2}$ would be systematically greater than $P_{xx}^{1.2}$. The distributions of $D^{1.2} = I_N^{1.2} - P_{xx}^{1.2}$ and $D^{1.7} = I_N^{1.7} - P_{xx}^{1.7}$ can be used to estimate the level of residual 1.2 Hz. We calculate $I_N^{1.2}$ and $I_N^{1.7}$ for each continuous chunk of TOD and, to increase the signal-to-noise on the estimated level of contamination, we average together the chunks within each FITS file outputted by the pipeline. This simulated dataset does not include a contaminating signal at 1.2 Hz and therefore both $D^{1.2}$

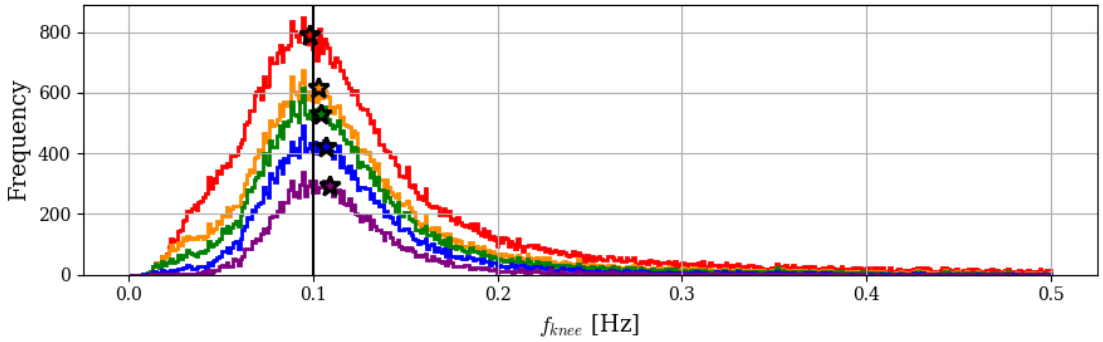

Figure 3.13: Histogram of $f_{\rm knee}$ values calculated from converged MCMC chains fitted to simulated TOD including residual sky signal and noise realisations with uniform priors on α, σ_r and σ_w. The *red, orange, green, blue* and *purple* histograms are of the parameter estimates from TOD chunks longer than 10, 30, 40, 60 and 100 s respectively. The *stars* are calculated from the weighted averages, accounting for parameter covariance, of each histogram. The vertical *black* line is at the true $f_{\rm knee}$ value.

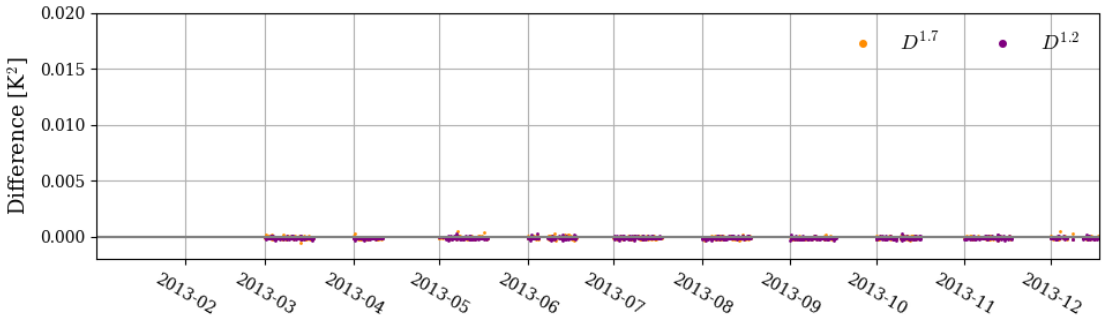

Figure 3.14: $D^{1.2}$ (*purple*) plotted over $D^{1.7}$ (*orange*) against the date of the chunk for the simulated dataset including sky signal and noise realisations.

and $D^{1.7}$ should be centred around zero. $D^{1.2}$ and $D^{1.7}$ are plotted in Figure 3.14 as a function of the date of the observation, both of these statistics are clustered around zero and show no obvious systematic bias.

It is left as future work to calculate the expected distributions of these statistics and to quantify the residual contamination in the time series. In the meantime the distributions of $D^{1.2}$ and $D^{1.7}$ from real data can be compared to these simulated results for a qualitative analysis.

3.3.4.3 Recovered noise statistics from survey scans

For an early data release we are focusing on the C-BASS North elevation 37° survey scans. All of these scans cross the NCP, which provides an easy check on drifts in

the system. To assess the noise on this data we fitted the noise model to the E37S60 dataset (elevation 37° scans with a 60° sun cut) using the process outlined above with an $N_{\rm side} = 1024$ map made from the data to estimate the Stokes I sky signal. We only use the E37N dataset (elevation 37° nighttime scans) for scientific analysis.

Recovered noise statistics The TOD and best-fit noise model for a typical chunk is shown in Figure 3.15. Histograms of the recovered parameter values for each chunk are shown in Figure 3.16. The distributions of recovered noise parameters in Stokes I peak around: 1.5 for α; 40 mK for σ_r; and 22 mK for σ_w. Interestingly Stokes Q typically has a lower white noise level than Stokes U data and its $1/f$ noise exhibits a slightly broader distribution of values. It is important to note that a single weighted average, over all time, for each parameter is not a particularly meaningful number as the noise properties are expected to change with time. The noise properties could change due to the receiver components ageing, telescope repairs and the weather. The noise statistics are expected to be stationary over the 100 minute duration of each FITS file outputted by the C-BASS pipeline and so to increase the signal-to-noise on the recovered parameters we calculated the weighted average for each schedule.

The weighted averages for each schedule are shown in Figure 3.17. Periods of time with outlying noise properties can be clearly seen. For example: Stokes I data from the end of February 2013; I, Q and U data from June and October 2013; around December 2013 the noise properties started to drift before repairs to the telescope were made and they stabilised around January 2014.

The histograms of these weighted averages are shown in Figure 3.18 and the corresponding knee frequencies in 3.19.

1.2 Hz contamination $D^{1.2}$ and $D^{1.7}$ for the real elevation 37° data are plotted in Figure 3.20. Similarly to the simulations, $D^{1.7}$ is clustered around zero. After January 2013, the values of $D^{1.2}$ show a significant bias above zero, the level of which changes over time. This implies there is residual 1.2 Hz contamination in the data and further work should be done to quantify its effect on the final maps.

From these results we can make a crude estimate the variance of the contamination in the TOD. The median values of $D^{1.2}$ in I, Q and U are 1.2×10^{-3}, 2.2×10^{-4} and $1.2 \times 10^{-4}\,\mathrm{K}^2$. This is the average excess power in a band $\Delta\nu = 0.04$ Hz centred on 1.2 Hz. The corresponding variance in the TOD is found by multiplying these numbers by $\Delta\nu / f_{\rm nyquist}$. The standard deviation of the 1.2 Hz signal in the TOD is

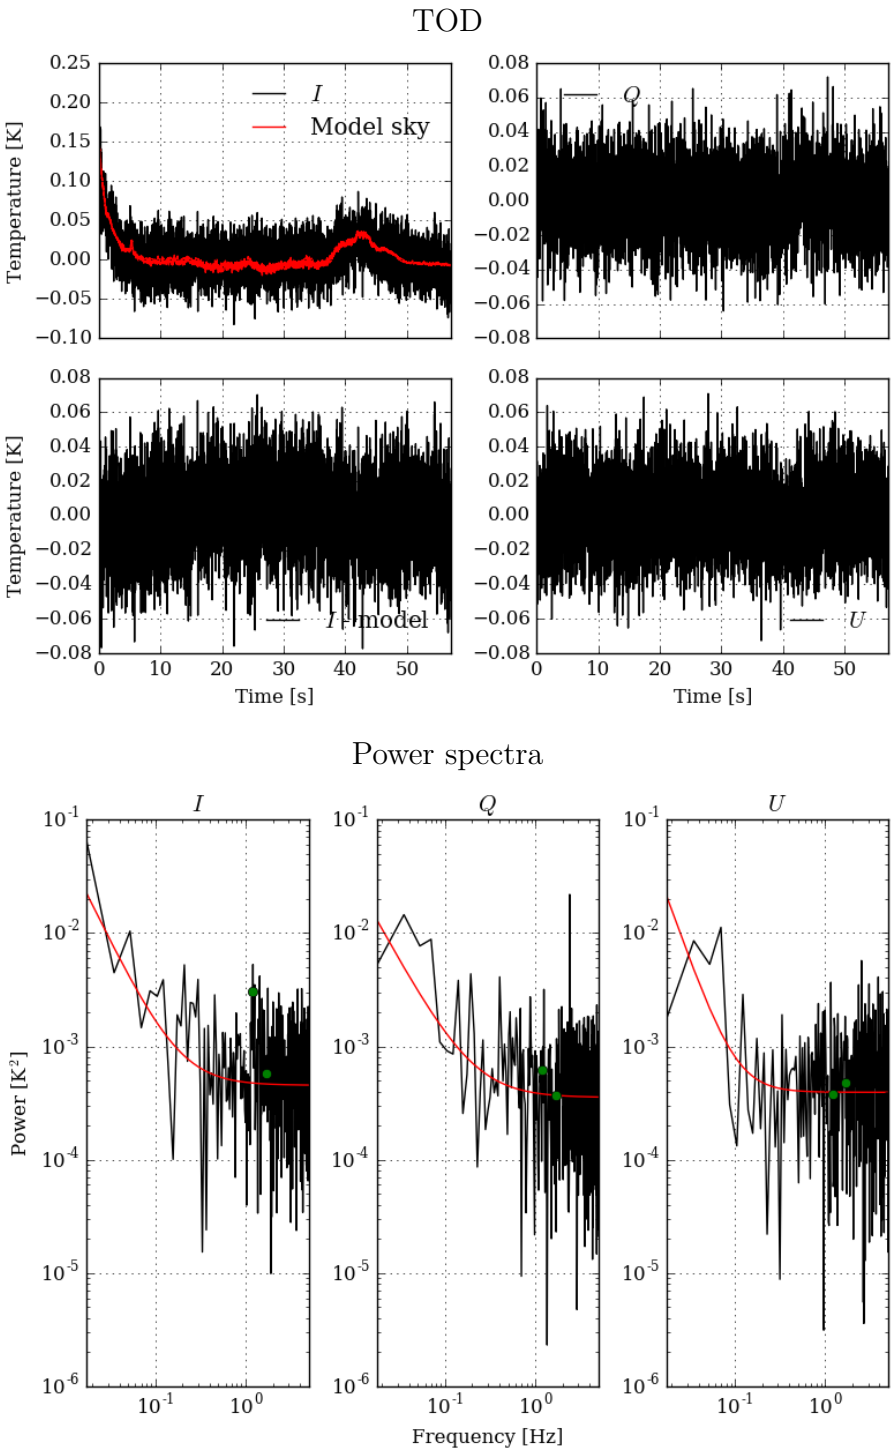

Figure 3.15: The *top* and *middle* plots show a typical chunk of C-BASS North TOD for Stokes I with the model sky overplotted in *red* (*top left*) and the model sky subtracted (*middle left*), Q (*top right*) and U (*middle right*) . The *bottom* three plots show periodograms of this chunk with the best fit (maximum posterior) noise model overplotted in *red* for I (*bottom left*), Q (*bottom middle*) and U (*bottom right*). $I_N^{1.2}$ and $I_N^{1.7}$ are shown by the *green circles*.

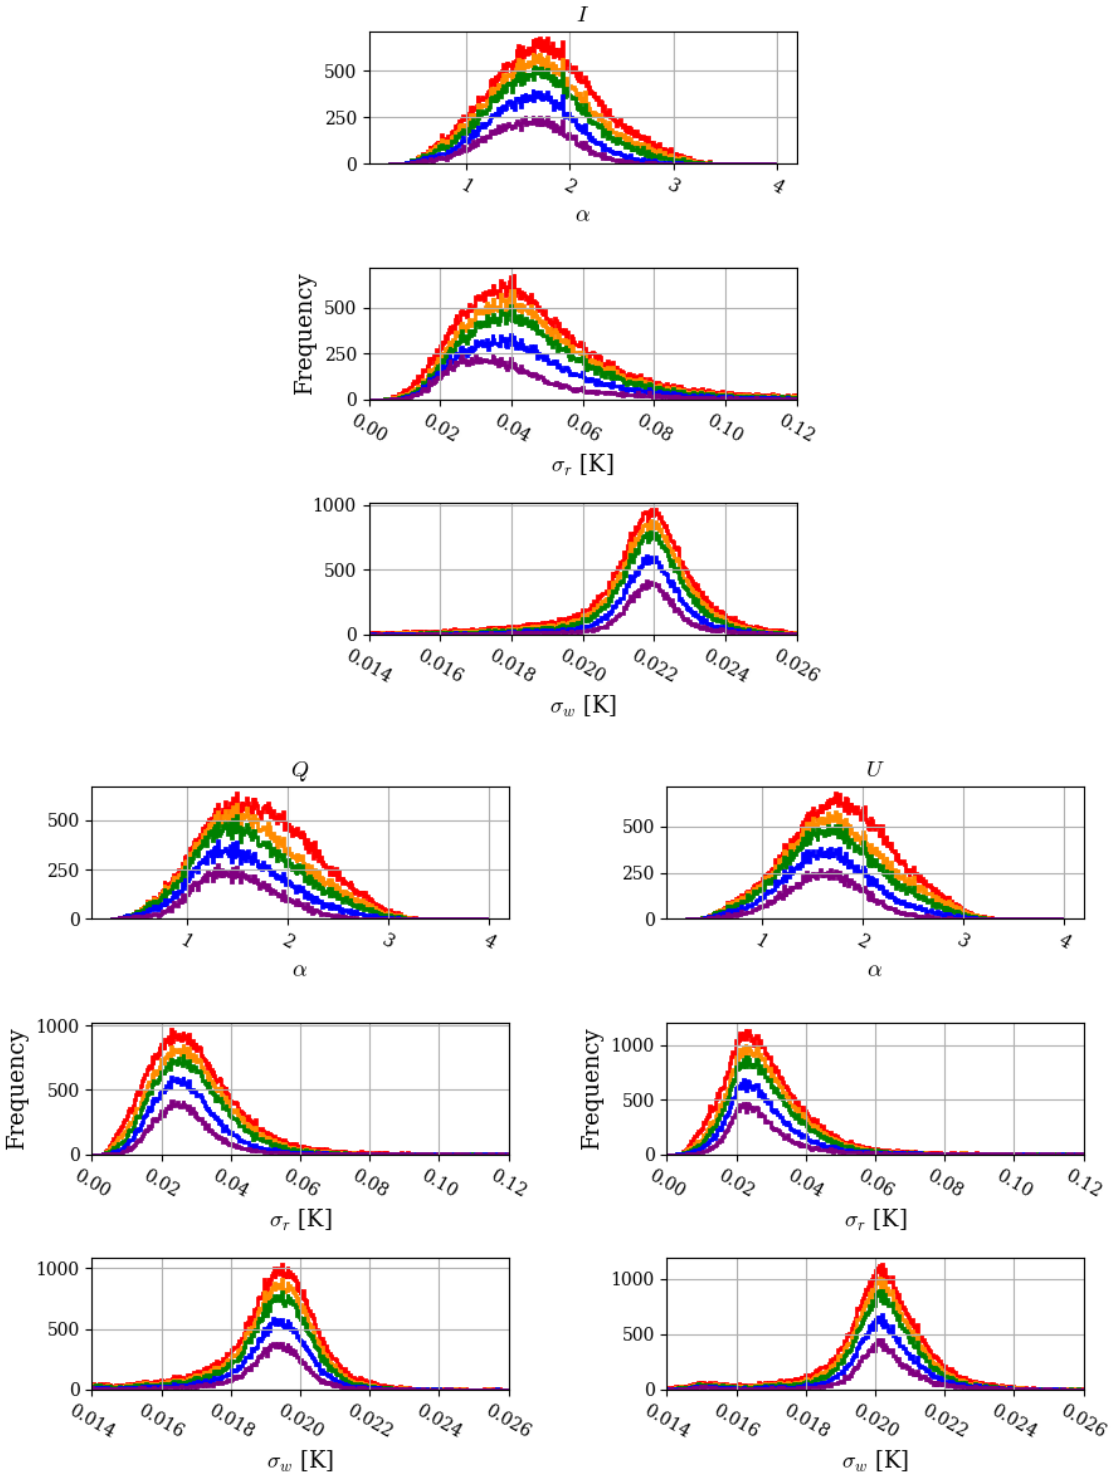

Figure 3.16: Histograms of mean parameter values from converged MCMC chains fitted to C-BASS North TOD, an estimate of the sky signal was subtracted from the Stokes I data. The *red, orange, green, blue* and *purple* histograms are of the parameter estimates from TOD chunks longer than 10, 30, 40, 60 and 100 s respectively.

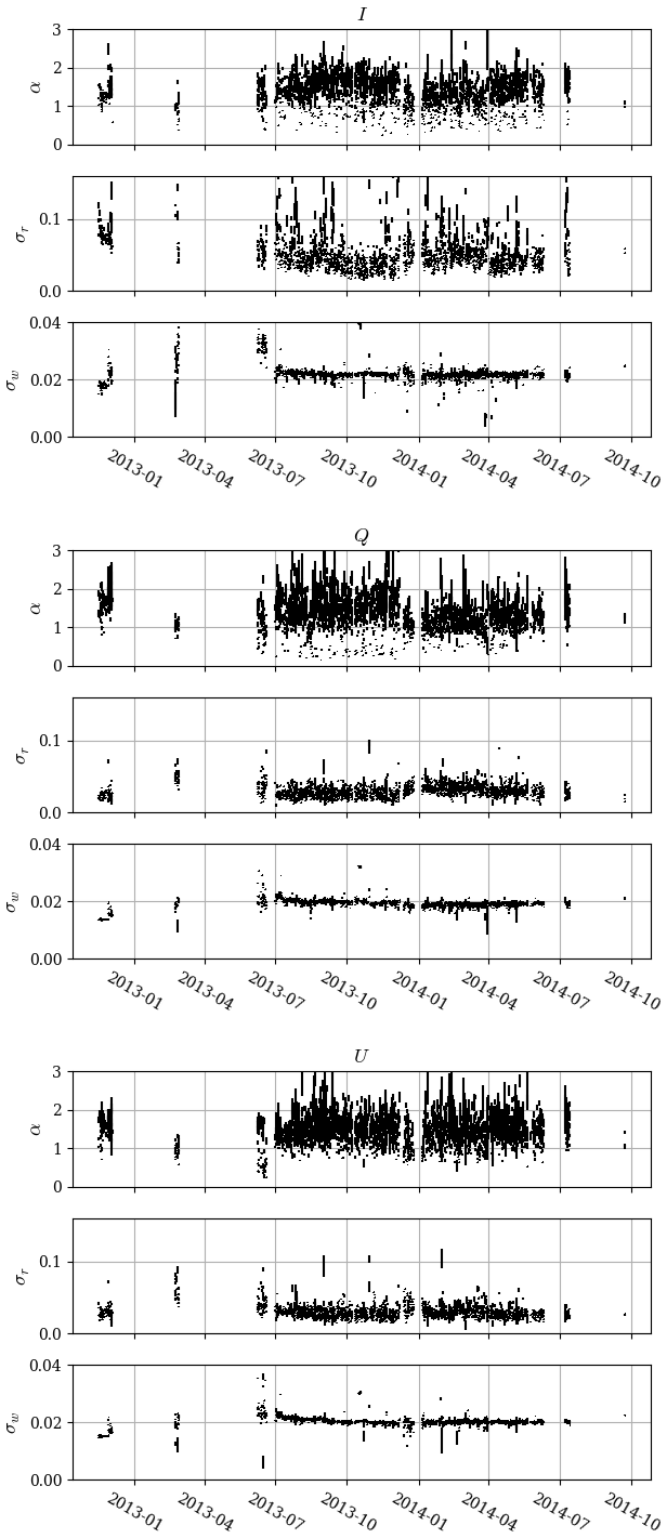

Figure 3.17: Weighted average noise parameters for each schedule with 1σ errorbars. The weighted averages were calculated from chunks longer than 30 s and accounted for parameter covariances.

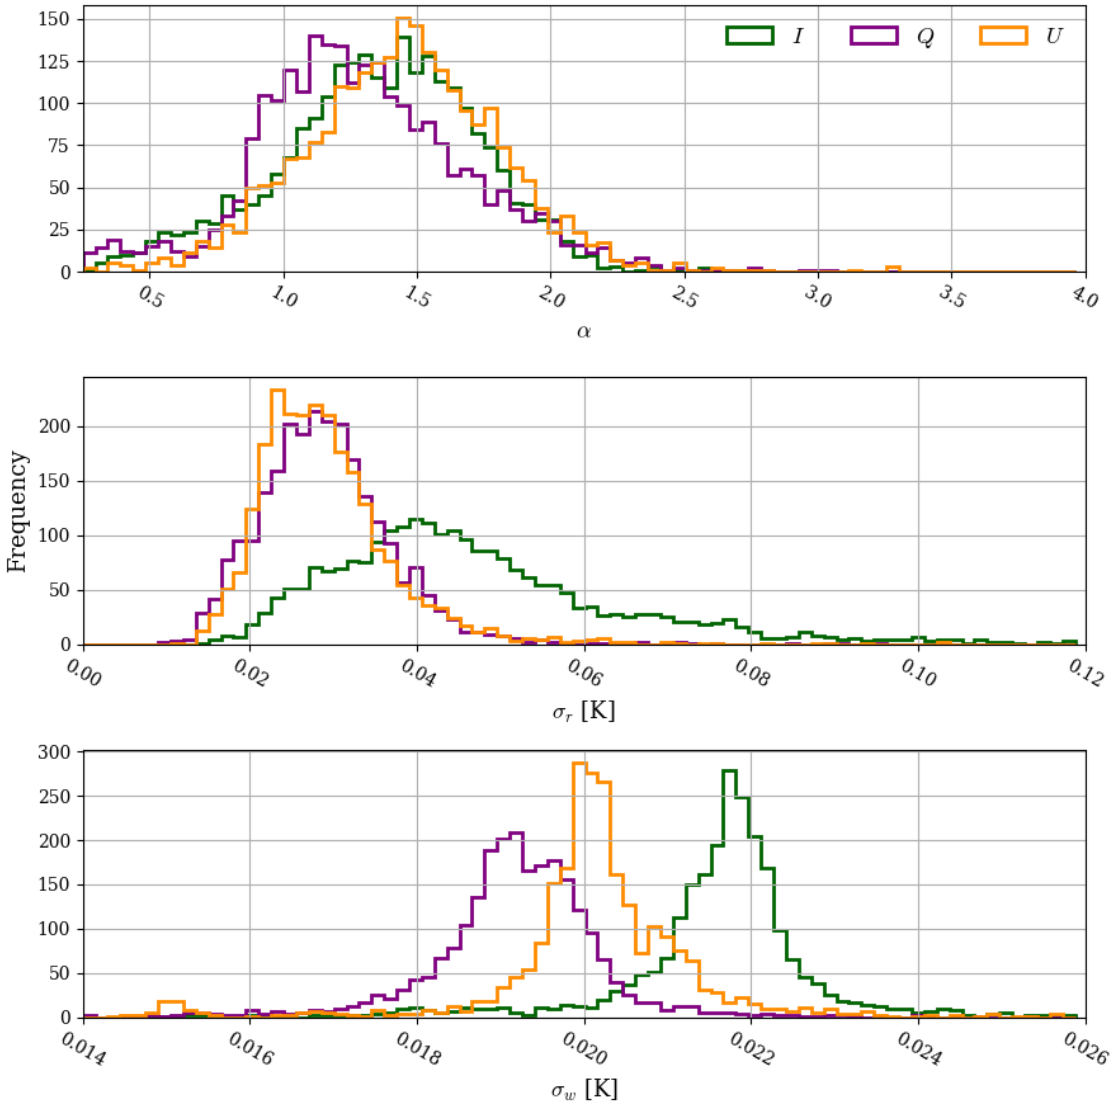

Figure 3.18: Histograms of the weighted average noise parameters for each schedule. The weighted averages were calculated from chunks longer than 30 s and accounted for parameter covariances.

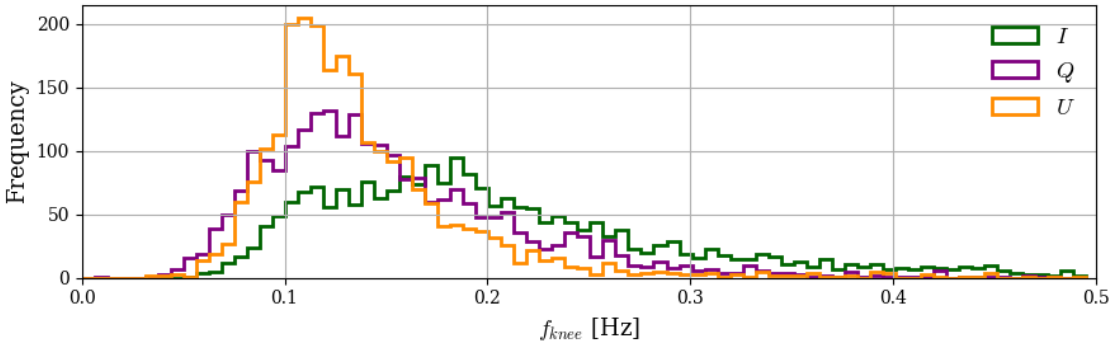

Figure 3.19: Histograms of the weighted average knee frequency for each schedule. The weighted averages were calculated from chunks longer than 30 s.

then 1.0, 0.42 and 0.31 mK in I, Q and U respectively. These values are significantly below the typical white noise levels of $\sigma_w \sim 20$ mK.

3.3.5 Noise statistics summary

We have developed a method of recovering the noise properties of C-BASS TOD and used it on real data. The noise properties can be used to improve the mapmaking process. We also measured the residual contamination in the TOD at 1.2 Hz, further work needs to be carried out to fully quantify its effect although it does not appear significant when compared to the typical white noise variance in the TOD.

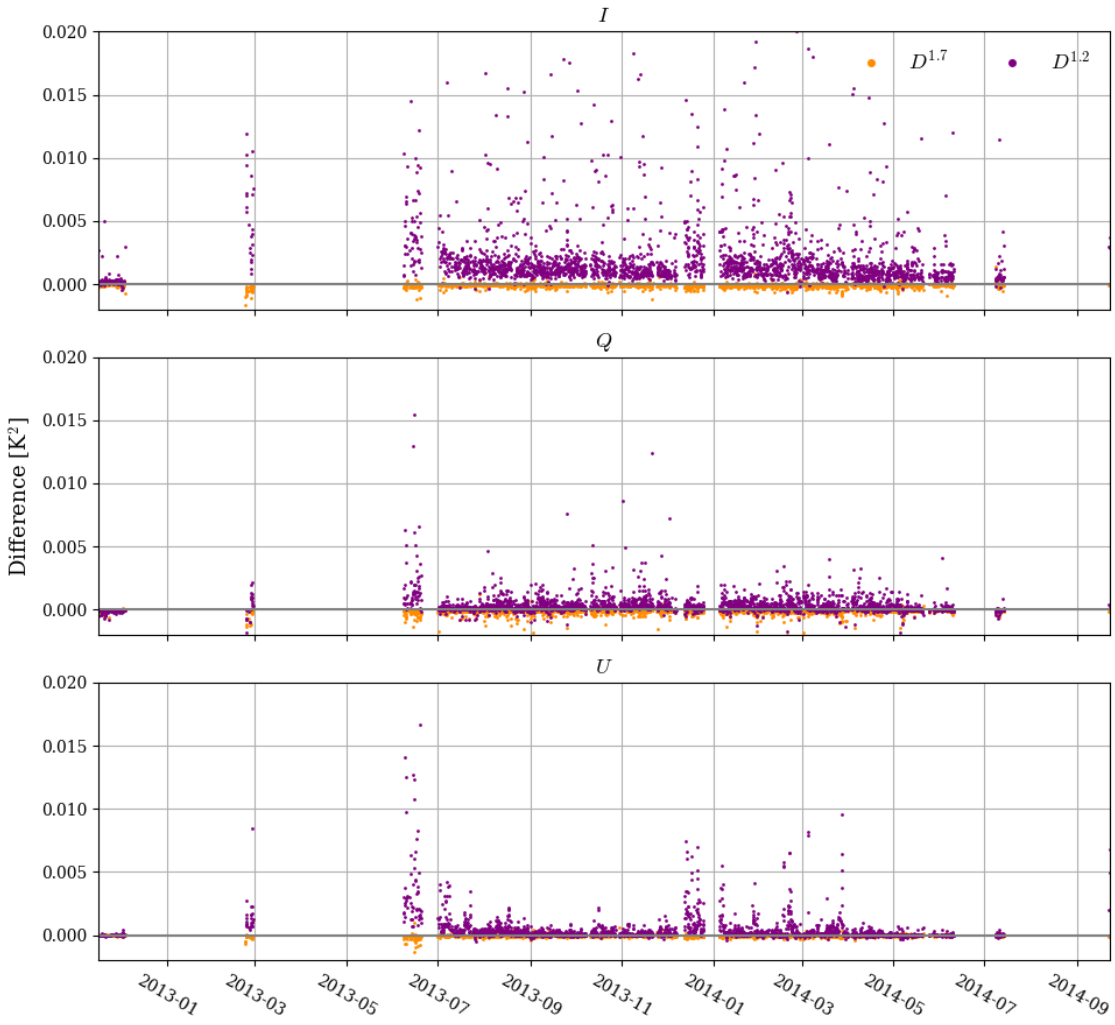

Figure 3.20: 1.2 Hz for all schedules.

Chapter 4

C-BASS North Data Analysis 2: Maps

This chapter presents maps of the C-BASS dataset that was described in Chapter 3 as well as the results of jackknife tests, which have been used to ascertain the data quality. The maps are presented as HEALPix sky maps and are made using the DEStriping CARTographer (DESCART) mapmaker (Sutton et al., 2010, 2009).

Section 4.1 contains an overview of the destriping mapmaking technique, which is used by DESCART, where the $1/f$ noise is modelled as a series of constant offsets. In Section 4.2 we discuss the errors that destriping can introduce into sky maps, specifically signal error caused by sub-pixel fluctuations in the sky and residual correlated noise caused by mis-modelling of the noise. In Section 4.3 we describe the modifications that we made to DESCART to decrease the signal error in C-BASS maps and test those modifications on simulated data. Section 4.4 contains the results of jackknife tests that we used to test C-BASS maps for systematics and to verify the accuracy of the pixel covariance matrices estimated by DESCART. Section 4.5 presents preliminary C-BASS maps that we are using for early science.

4.1 Mapmaking and the destriping technique

Mapmakers attempt to solve the map making equation,

$$\mathbf{d} = \boldsymbol{P}\mathbf{m} + \mathbf{n}, \tag{4.1}$$

where $\mathbf{d}$ is the TOD vector, $\mathbf{m}$ is the true beam-convolved sky map vector, $\boldsymbol{P}$ is a pointing matrix that relates each time sample to a pixel in the map and $\mathbf{n}$ is the noise on the TOD. The noise has a covariance matrix $\boldsymbol{C_N} = \langle \mathbf{n}\mathbf{n}^T \rangle$ that need not be diagonal.

From Equation 4.1, the maximum likelihood estimator of $\mathbf{m}$ and its pixel covariance, $\boldsymbol{C}$, are given by

$$\mathbf{m} = (\boldsymbol{P}^{\boldsymbol{T}}\boldsymbol{C}_{\boldsymbol{N}}^{-\mathbf{1}}\boldsymbol{P})^{-1}\boldsymbol{P}^{\boldsymbol{T}}\boldsymbol{C}_{\boldsymbol{N}}^{-\mathbf{1}}\mathbf{d} \tag{4.2}$$

$$\boldsymbol{C} = (\boldsymbol{P}^{\boldsymbol{T}}\boldsymbol{C}_{\boldsymbol{N}}^{-\mathbf{1}}\boldsymbol{P})^{-1}. \tag{4.3}$$

In the limit of diagonal $\boldsymbol{C_N}$ the problem reduces to naively binning the TOD into pixels.

The maximum likelihood estimator of $\mathbf{m}$ is optimal but, for non-diagonal $\boldsymbol{C_N}$, is slow and memory intensive to calculate. The computational cost is particularly problematic for large datasets (Tegmark, 1997; de Gasperis et al., 2005). As well as using DESCART, the C-BASS collaboration are now working on interfacing NINKASI (Dünner et al., 2012), a maximum likelihood mapmaker, with C-BASS data.

One method of simplifying the problem is to apply a high-pass filter to $\mathbf{d}$, suppressing long time-scale $1/f$ drifts, and then to naively bin the filtered TOD into a map. This filtering approach is fast but distorts the recovered sky signal and biases the estimates of the polarization maps via E to B leakage (Sutton et al., 2010).

An alternative, commonly used, method of simplifying the problem is destriping (Maino et al., 2002a; Keihänen et al., 2004; Keihänen, Kurki-Suonio & Poutanen, 2005; Kurki-Suonio et al., 2009). The destriping technique is an approximation to the maximum likelihood estimator of $\mathbf{m}$ and simplifies the problem by separating the noise into a white component, $\mathbf{n_w}$, and a correlated component, $\mathbf{n}_{\rm corr}$, such that

$$\mathbf{n} = \mathbf{n_w} + \mathbf{n}_{\rm corr}, \tag{4.4}$$

The correlated noise is then modelled as a series of baselines with constant length, λ_b. The baselines are parametrized by a vector of offsets, $\mathbf{a}$, and a projection operator, $\boldsymbol{F}$, that maps the offsets onto the TOD such that $\mathbf{n}_{\rm corr} = \boldsymbol{F}\mathbf{a}$. The mapmaking equation is therefore modified to

$$\mathbf{d} = \boldsymbol{P}\mathbf{m} + \boldsymbol{F}\mathbf{a} + \mathbf{n_w}. \tag{4.5}$$

Destriping mapmakers simultaneously fit for both the sky map and the offsets. The maximum likelihood estimator of $\mathbf{a}$ is

$$\left(\boldsymbol{F}^{\boldsymbol{T}}\boldsymbol{C}_{\boldsymbol{w}}^{-\mathbf{1}}\boldsymbol{Z}\boldsymbol{F} + \boldsymbol{C}_{\boldsymbol{a}}^{-\mathbf{1}}\right)\mathbf{a} = \boldsymbol{F}^{\boldsymbol{T}}\boldsymbol{C}_{\boldsymbol{w}}^{-\mathbf{1}}\boldsymbol{Z}\mathbf{d} \tag{4.6}$$

where $\boldsymbol{C}_a$ is the covariance of the baseline amplitudes, the covariance of the white noise is $\boldsymbol{C_w} = \langle \mathbf{n_w}\mathbf{n_w}^T \rangle$ and $\boldsymbol{Z}$ is the signal cleaning operator that subtracts the naive map from $\mathbf{d}$. The signal cleaning operator is

$$\boldsymbol{Z} = \boldsymbol{I} - \boldsymbol{P}(\boldsymbol{P}^{\boldsymbol{T}}\boldsymbol{C}_{\boldsymbol{W}}^{-\mathbf{1}}\boldsymbol{P})^{-1}\boldsymbol{P}^{\boldsymbol{T}}\boldsymbol{C}_{\boldsymbol{W}}^{-\mathbf{1}}, \tag{4.7}$$

where $\boldsymbol{I}$ is the identity matrix. In "traditional destriping" (all papers prior to Keihänen, Kurki-Suonio & Poutanen (2005)) $\boldsymbol{C_a}$ is ignored and the baselines are assumed to be uncorrelated with one another. In "non-traditional" or covariant destriping $\boldsymbol{C_a}$ is not ignored and knowledge of the correlation between baselines is used (Keihänen, Kurki-Suonio & Poutanen, 2005; Sutton et al., 2009). Covariant destriping is not currently implemented in DESCART and all C-BASS maps are made using traditional destriping.

Other than $\mathbf{a}$, all the terms in Equation 4.6 are known and $\mathbf{a}$ can be quickly solved for using the preconditioned conjugate gradient method (PCG). For details on this method see, for example, Press (2007).

The PCG method is used to solve equations of the form $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$ by first preconditioning the problem using a matrix $\boldsymbol{K}$, such that $\boldsymbol{K}^{-1}\boldsymbol{A}\mathbf{x} = \boldsymbol{K}^{-1}\mathbf{b}$. Ideally $\boldsymbol{K}^{-1} \simeq \boldsymbol{A}^{-1}$ but it should also be fast to compute. The choice of $\boldsymbol{K}$ is a compromise between these two requirements. To solve the mapmaking equation, a reasonable choice is $\boldsymbol{K} = \boldsymbol{F^T C_w^{-1} F} + \boldsymbol{C_a^{-1}}$ (Sutton et al., 2009).

If the data are perfectly described by the model in Equation 4.5, and there are no limits to the numerical precision of the computer solving the problem, then an exact solution would be found after a finite number of steps. It would take as many steps as there are baseline amplitudes to be found. Due to these limitations, however, the exact solution can never be found and instead the PCG algorithm is iterated until either a desired level of convergence is reached or maximum number of steps have been made. The target error is $|\epsilon| = |\boldsymbol{A}\mathbf{x} - \mathbf{b}| < \tau|\mathbf{b}|$, where τ is the PCG tolerance.

Once $\mathbf{a}$ has been estimated, the offsets can be subtracted from $\mathbf{d}$ and the cleaned TOD, $\mathbf{d} - \boldsymbol{F}\mathbf{a}$, can be naively binned to return a destriped map. The destriped map and its pixel covariance are given by

$$\mathbf{m} = (\boldsymbol{P^T C_w^{-1} P})^{-1} \boldsymbol{P^T C_w^{-1}} (\mathbf{d} - \boldsymbol{F}\mathbf{a}) \tag{4.8}$$

$$\boldsymbol{C} = (\boldsymbol{P^T}(\boldsymbol{C_w} + \boldsymbol{F C_a F^T})^{-1} \boldsymbol{P})^{-1}. \tag{4.9}$$

This process is significantly faster than finding the full maximum likelihood solution of Equation 4.3.

The offset length λ_b is the main tuneable parameter in a destriping mapmaker. When λ_b is the length of a single TOD sample then the destriping solution is equivalent to the maximum-likelihood estimate of the sky. Shorter offsets can better model the $1/f$ drifts, but as λ_b decreases the offsets become increasingly correlated with one another and the $\boldsymbol{C_a}$ term becomes important.

Destriping cannot recover the monopole of Stokes I maps. This can be demonstrated by parametrizing the monopole of a map with a constant vector $\mathbf{m}_{\rm mono}$. The monopole adds a global offset to the I TOD such that $\mathbf{d} \rightarrow \mathbf{d} + \boldsymbol{P}\mathbf{m}_{\rm mono}$. This monopole makes no impact on the right hand side of Equation 4.6 because $\boldsymbol{Z}\boldsymbol{P}\mathbf{m}_{\rm mono} = 0$. This means that the monopole in the recovered I map is arbitrary. This is not the case for Stokes Q and U maps. Because of their rotation properties, Q and U monopoles do not transform to global offsets in $\mathbf{d}$ and the destriper can recover their zero levels. Furthermore, C-BASS I data have an arbitrary monopole because the sky brightness is compared to the reference load signal and any monopole is degenerate with the load temperature.

4.2 Destriping errors

We are modelling TOD from C-BASS as the linear sum of three independent components; a sky signal, white noise and $1/f$ drifts. Because destriping is a linear operation, the contributions to the final map from each of these components are also independent and additive. Destriping each of these components will introduce different errors into the final map. Following the nomenclature of Keihänen et al. (2010), the residual error can be separated into signal error and residual noise caused by destriping the sky and noise signals respectively. The residual noise can again be separated into contributions from the white component and the $1/f$ component. The residual noise from the $1/f$ component is called correlated residual noise. In this section we describe the source of these errors.

If TOD contained only sky signal (and optionally white noise), then the naively binned map would be the optimal estimator of $\mathbf{m}$. Signal error is caused by sub-pixel fluctuations in the sky. The destriping solution implicitly assumes that the true sky map is pixelized at the same resolution as the map being produced. This is a poor assumption within pixels where the sky signal is changing rapidly, as is demonstrated in Figure 4.1. In the diagram, multiple TOD samples fall within a single pixel but are not exactly coincident with each other. Destriping mapmakers assume that all TOD samples falling within a pixel measure the same sky signal but can be subject to different realisations of the white noise or baseline offsets. If the sky signal has changed significantly within a pixel then a destriping mapmaker can only interpret the sub-pixel fluctuations as noise and will attempt to fit them with the offsets. The signal error spreads out sharp features in the sky along the scan directions. Destriping at higher resolution will reduce the magnitude of this error (by reducing the level of

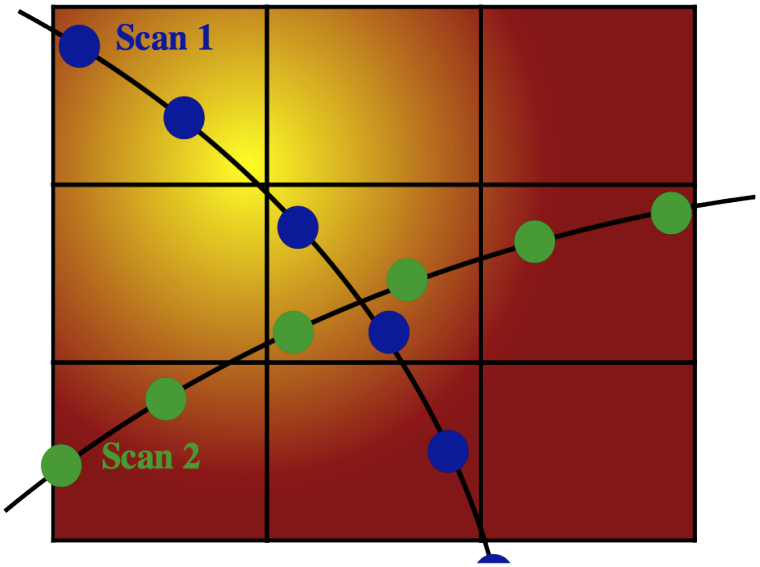

Figure 4.1: Schematic of two scans that cross in the central pixel. The true sky signal is changing rapidly within the pixel (yellow hot, red cold). The circles represent samples along the two scans and are not coincident with one another or the pixel centers.

sub-pixel fluctuations) but the greatest improvements can be found by excluding pixels with significant sub-pixel fluctuations when estimating **a**, but then including them in the final naive binning of the offset subtracted TOD (Keihänen et al., 2010; Tristram et al., 2011). Typically pixel brightness is used as a proxy for the level of sub-pixel fluctuations.

In addition to signal error, residual correlated noise is left in the map because real $1/f$ drifts do not take the form of a series of baselines with constant length. Thus un-modelled drifts will pass into the map.

To most accurately recover the sky and baseline amplitudes requires the pixels to be crossed at a range of angles. In general, the more crossing angles per pixel the better the sky and baseline offsets can be recovered and this is shown in Section 4.3.2.

4.3 DESCART

DESCART is the destriping mapmaker used by the C-BASS collaboration (Sutton et al., 2010, 2009). The pixel covariance matrices produced by DESCART only include the I, Q and U variances and the QU covariance for each pixel (σ_{II}, σ_{QQ}, σ_{UU} and σ_{QU} respectively). The pixel-to-pixel correlations are not estimated because, for any

reasonable $N_{\rm side}$, they are very large matrices ($N_{\rm pix} \times N_{\rm pix}$). DESCART assumes zero IQ and IU covariance when calculating the offset amplitudes and the pixel covariance matrices. From fits of the C-BASS beam to rasters of radio point sources the C-BASS I to QU leakage is measured to be $\sim 1\%$ and so to a very good approximation the IQ and IU covariances are zero.

DESCART produces maps in celestial co-ordinates and, due to the C-BASS scan strategy, it is often easiest to identify systematic errors in maps in this coordinate frame.

We have added a masking option to reduce the size of signal errors in maps produced from high signal C-BASS I TOD. In Section 4.3.1 we show that our modifications have been successful by testing them on simulated data. Currently DESCART only implements traditional destriping, ignoring $\boldsymbol{C_a}$. With maps made from simulated data, in Section 4.3.2 we show that covariant destriping will be needed to create maps free from residual correlated noise at large angular scales. We also show that when including scans at multiple elevations the signal errors are smaller than when making a map from scans at a single elevation due to the larger number of crossing angles per pixel.

4.3.1 Masking to reduce signal error

Signal errors may be significant in maps made from C-BASS data because of the large dynamic range and steep gradients of signal strength when scanning acrosss, for example, the Galactic plane and bright radio point sources. We edited DESCART so that it could exclude TOD samples from masked regions of the sky when estimating the baseline offsets $\mathbf{a}$, but then include those samples during the naive binning of the offset corrected TOD, $\mathbf{d} - \boldsymbol{F}\mathbf{a}$, into a map. In this section we explain our modification and demonstrate its success on simulated data.

We edited DESCART to read in a HEALPix mask saved in a FITS file. In the mask file, pixels affected by significant sub-pixel fluctuations should be set to zero and all other pixels set to one, sky brightness could be used as a proxy for the level of sub-pixel fluctuations. During the estimation of $\mathbf{a}$, TOD samples that come from masked sky regions are ignored by setting their weight to $C_w = 0$. Once $\mathbf{a}$ has been estimated, all of the data are again included and weighted normally when naively binning $\mathbf{d} - \boldsymbol{F}\mathbf{a}$.

To test our modification, we took roughly half of the E3747S60 dataset (elevation 37° and 47° scans with a 60° sun flag) and overwrote the signal with simulated data. We used the noise model in Equation 3.4 and simulated four timestreams, one that

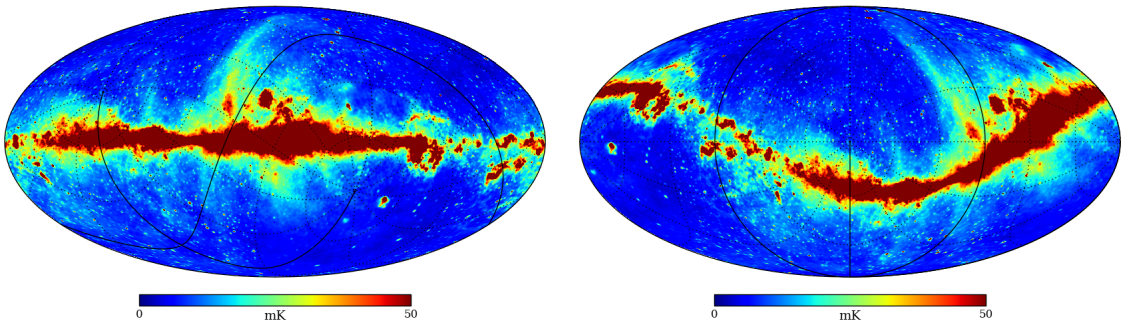

Figure 4.2: Model 5 GHz sky map including diffuse Galactic emission and Northern point sources. The graticules are lines of constant RA and DEC. The sky has been projected in two coordinate systems, Galactic coordinates (*left*) and Celestial coordinates rotated so that the NCP is in the centre of the image (*right*).

only contained realisations of the white noise, one that only contained realisations of the $1/f$ noise, one that only contained sky signal and one with all three components. The white noise level was set to $\sigma_w = 25\,\mathrm{mK}$ and the $1/f$ component was modelled with $\alpha = 1.5$ and $\sigma_r = 25\,\mathrm{mK}\,@\,0.1\,\mathrm{Hz}$. The simulated sky signal was generated using the pointing of the C-BASS North scans to extract values from an $N_{\mathrm{side}} = 1024$ sky map smoothed to a Gaussian beam with FWHM 0.7°. The model sky is shown in Figure 4.2. The diffuse emission in the model sky was simulated using PySM (Thorne et al., 2016) and point sources were added from the GB6 (Gregory et al., 1996) and S5 (Kuehr et al., 1981) catalogs. The point sources M87, W51, TauA, CygA and CasA had to be added by hand as they were excluded from the GB6 catalog, their values being taken from Gregory et al. (1996); Baars et al. (1977); Aller & Reynolds (1985); and Hafez et al. (2008).

The signal error can be estimated by taking the difference between the destriped sky-only map and the naively binned sky-only map, and an example is shown in Figure 4.3 with $\lambda_b = 10\,\mathrm{s}\,(= 1/f_{\mathrm{knee}})$. Because the scan speed is $\sim 4\,\mathrm{deg\,s^{-1}}$ and the telescope scans at elevations of 37° and 47°, baselines of this length have angular scales 32° and 27° on the sky respectively. Signal errors have been introduced into the destriped map along the scan directions spreading from regions of the sky with significant sub-pixel fluctuations, such as near radio point sources and the Galactic plane. The errors have angular sizes that correspond to those of the baselines. Masking the pixels with significant sub-pixel fluctuations will significantly reduce this error.

The subplots in Figure 4.4 show the Descart signal-error maps when destriping in traditional mode with various offset lengths ($\lambda_b = 2.5, 10$ and $20\,\mathrm{s}$) both with (*right*) and without (*left*) masking the brightest 5% of the sky during the estimation of the

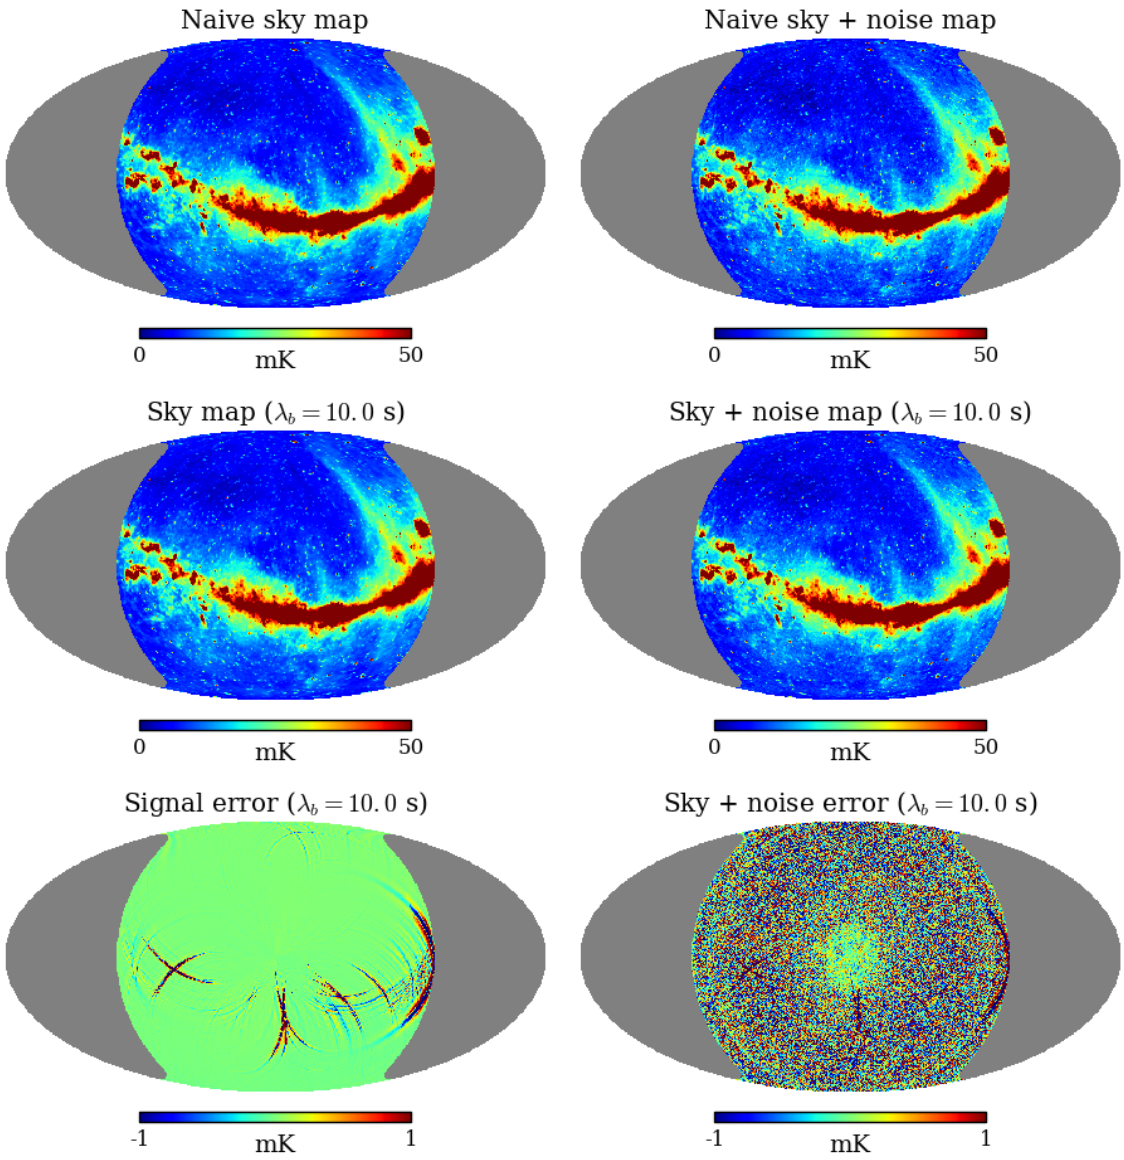

Figure 4.3: DESCART maps made from simulated TOD that only contains a sky-signal (*left*) and TOD that contains sky-signal and noise (*right*). From *top* to *bottom* are the naively binned maps, the destriped maps and then the difference between the destriped maps and the naively binned signal only map, i.e. the signal error on the *bottom left* and the sum of the signal error and residual noise on the *bottom right*. The mapmaker was run with offset length $\lambda_b = 10\,\mathrm{s}\,(= 1/f_{\mathrm{knee}})$ and no masking was used during the estimation of the baseline offsets. Signal errors have been introduced into the destriped map along the scan directions spreading from regions of the sky with significant sub-pixel fluctuations, such as near radio point sources and the Galactic plane. The errors have angular sizes that correspond to those of the baselines.

baseline amplitudes. The 5% mask is shown in Figure 4.5. The signal error is clearly worst at small λ_b and masking the brightest 5% of the sky reduces its magnitude drastically, this is quantified below.

To help understand the total contributions to the final error map, it is useful to plot the residual noise maps. Figure 4.6 shows a naively binned map of simulated TOD that only contain realisations of the $1/f$ noise and a destriped map made from the same data with $\lambda_b = 10\,\mathrm{s}$. The residual correlated noise is much lower in the destriped map than in the naively binned map. Because the $1/f$ drifts do not take the form of a series of constant offsets they cannot be perfectly modelled by the destriper and residual correlated noise is left in the destriped map.

Figure 4.7 shows a naively binned map made from simulated TOD that only contain realisations of the white noise and a destriped map made from the same data with $\lambda_b = 10\,\mathrm{s}$. In the case of white noise only, the naively binned map is the optimal estimator of the sky. The difference between the two maps shows large-scale structure that has been introduced by destriping.

Figures 4.4, 4.6 and 4.7 show the effects of destriping on each of the three components (the sky-signal, $1/f$ noise and white noise). The signal errors are very small when masking the brightest pixels during offset estimation but they are still comparable to the noise errors. We use the angular power spectra of maps of each component to quantify at which angular scales each effect dominates and this is described below.

We used the publicly available PolSpice package to estimate the angular power spectra (Szapudi et al., 2001).[1] We weighted the pixels uniformly, apodized the map to reduce ringing caused by the fractional sky coverage and the resulting spectra were binned with a width $\Delta\ell = 10$.

The angular power spectra of maps made from simulated data with either no masking (*left*) or applying a 5% mask (*right*) are shown in Figure 4.8. The offset length increases from *top* to *bottom*.

The total destriping error is shown in *solid black*. This is the difference between the destriped map made from simulated TOD that contains all three components (sky, white and $1/f$ noise) and the naive map made from sky-only data. The destriped map made from TOD that only contains $1/f$ noise is *solid purple*. The destriped map made from TOD that only contains white noise is *solid orange*. The signal error is *solid cyan*.

For reference, each plot also shows the power spectrum of naively binned maps made from the simulated TOD. The naive map made from TOD that only contains

[1] http://www2.iap.fr/users/hivon/software/PolSpice/

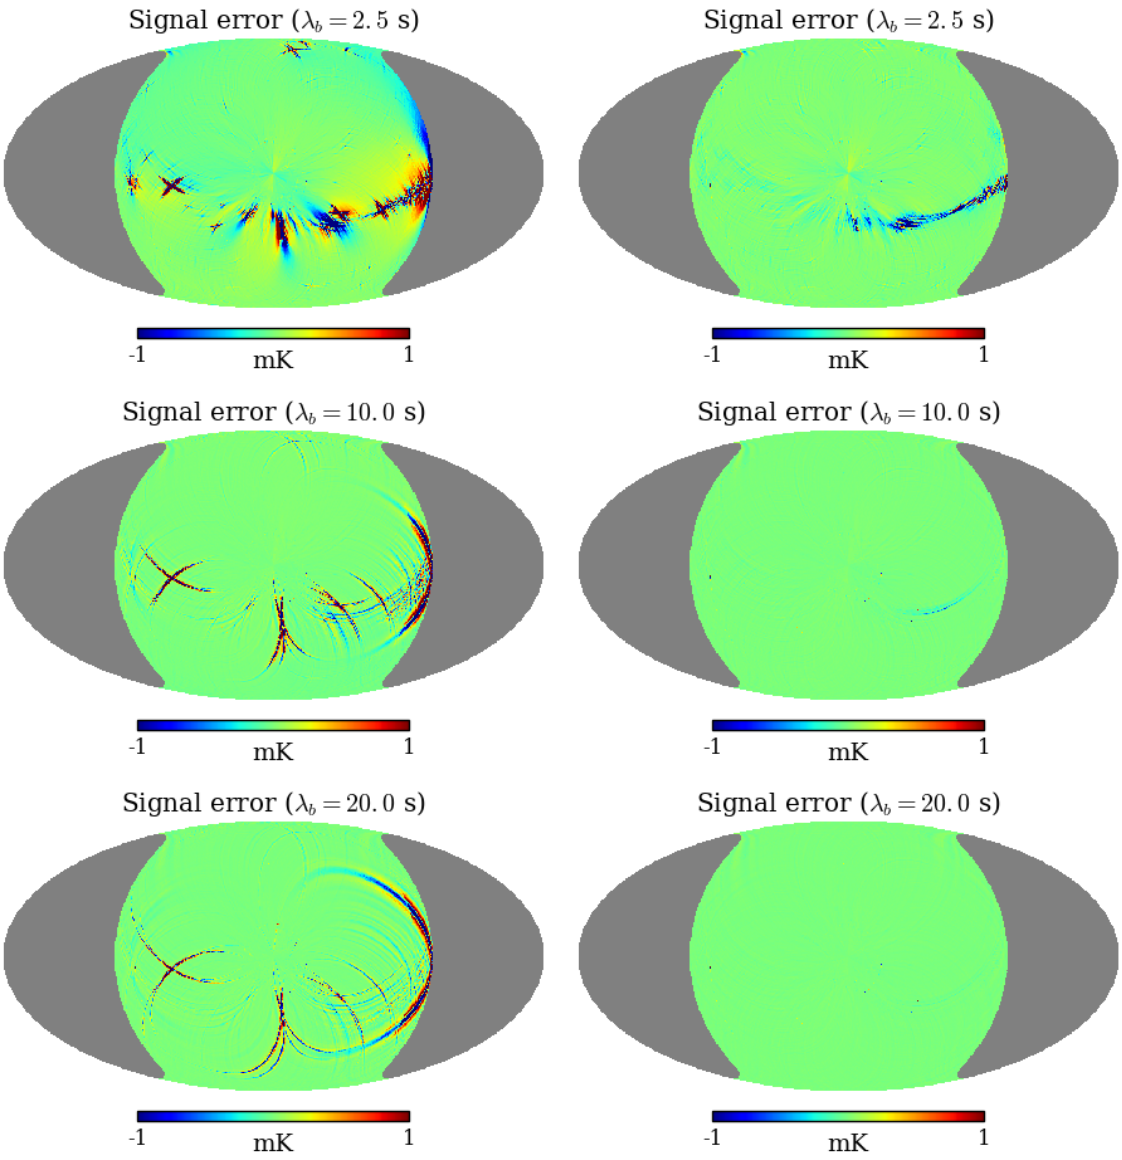

Figure 4.4: The difference between destriped maps and naive maps made from simulated TOD that only contain a sky-signal component. These difference maps show the signal error introduced by destriping. DESCART was run using three different offset lengths, from *top* to *bottom* $\lambda_b = 2.5, 10$ and 20 s. The maps on the *left* were made without masking any pixels during the offset estimation step and those on the *right* excluded data within a 5% brightness mask during offset estimation. Applying the mask greatly reduces the signal error due to destriping.

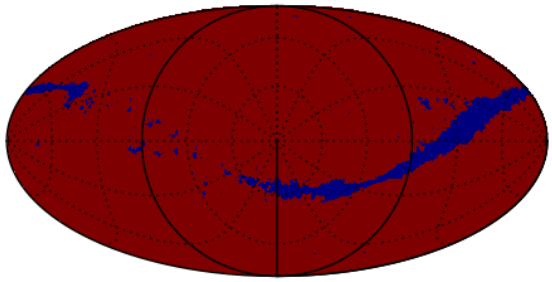

Figure 4.5: Mask excluding the brightest 5% of pixels in the simulated sky map. *Red* is included, *blue* is excluded.

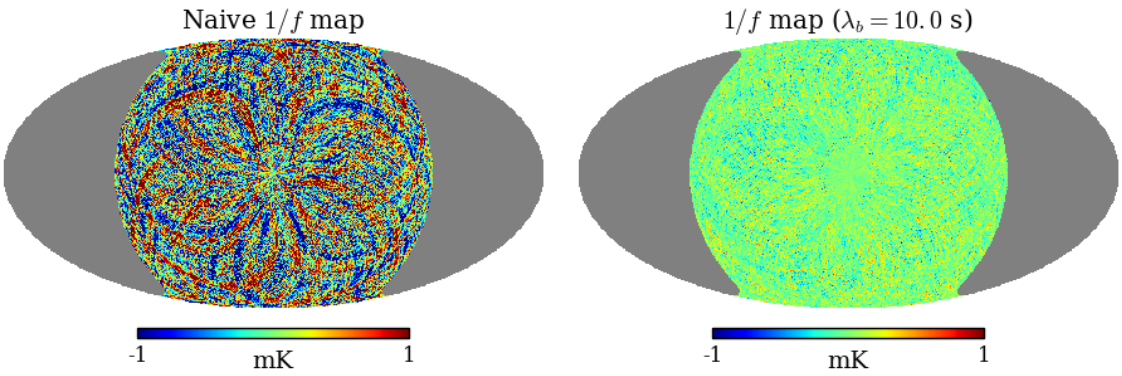

Figure 4.6: DESCART maps made from simulated TOD that only contain $1/f$ noise. The map on the *left* was made by naively binning the TOD, the map on the *right* is the destriped map with $\lambda_b = 10\,\mathrm{s}$ and no masking was used during the offset estimation. There is residual noise in the destriped map but at a much lower level than the naive map.

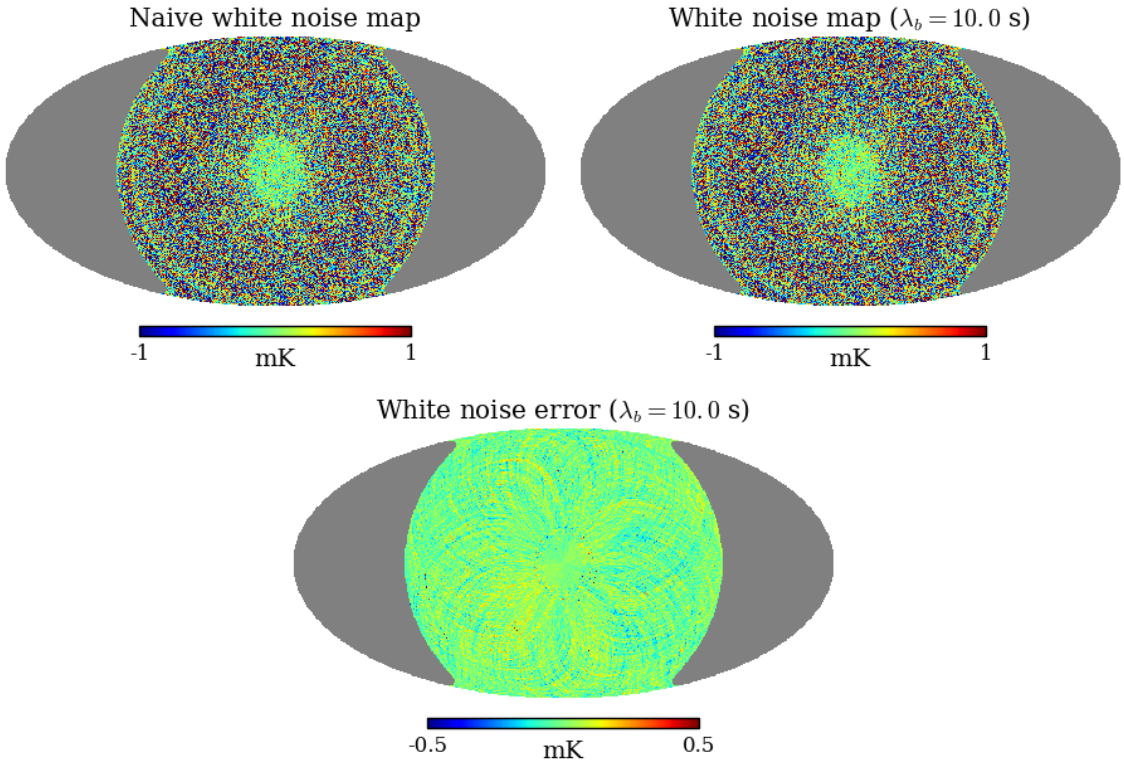

Figure 4.7: DESCART maps made from simulated TOD that only contain white noise. The map on the *top left* was produced by naively binning the TOD, the map on the *top right* is the destriped map with $\lambda_b = 10\,\mathrm{s}$ and no masking was used during the offset estimation. The *bottom* map is the difference between the naive and destriped maps (note the different colour scale). Destriping has introduced structure, along the scan directions, into the white noise map.

sky signal is *solid grey*. The naive map made from TOD that only contains white noise is *dashed orange*. The naive map made from TOD that only contains $1/f$ noise is *dashed purple*. These three lines are identical in all of the power spectrum plots.

The power of the signal error (*solid cyan*) decreases slowly with increasing offset length. Masking during offset estimation reduces the power of the signal error by 1–3 orders of magnitude, depending on ℓ. Without masking, for short offset lengths the signal error dominates the total destriping error across a large ℓ range and at longer offset lengths it only dominates at intermediate ℓ values. When masking, the signal error is subdominant to the residual noise for all but the shortest offset lengths.

The residual $1/f$ noise after destriping (*solid purple*) has orders of magnitude less power than the naively binned $1/f$ noise (*dashed purple*). Destriping has reduced the power in the residual correlated noise component by factors of 10–100 across all multipoles and (for $\lambda_b > 5\,\mathrm{s}$) it is only greater than the white noise component

at $\ell \lesssim 20$. Residual $1/f$ noise dominates at low ℓ and has the least power when $\lambda_b \sim 1/f_{\rm knee}$.

At high ℓ white noise dominates the destriping error and on these small angular scales destriping does not have much effect on the white noise power, the naively binned white noise (*dashed orange*) and destriped white noise (*solid orange*) lie on top of each other. At low ℓ, where white noise is subdominant to the other components, destriping with shorter offset lengths increases its power. When $\lambda_b \lesssim \frac{1}{2f_{\rm knee}}$ the noise is misfit because the correlation between baseline offsets has become significant and the $\boldsymbol{C_a}$ term in Equation 4.6 is no longer negligable. This misfitting introduces extra power in the white noise component at small ℓ. The non-traditional mode, or covariant destriping, in DESCART is currently non-functional and in future work needs to be correctly implemented.

When masking the brightest pixels during the baseline amplitude estimation, the signal error is less than the residual noise at all angular scales and is orders of magnitude lower than when not masking. From these spectra, the optimal value of λ_b corresponds to $\sim 1/f_{\rm knee}$.

4.3.2 Residual correlated noise

To test the level of residual noise in maps made from data with more realistic noise properties we simulated another set of timestreams. For this set of simulations we took the E37S60 dataset (C-BASS North elevation 37° scans) and overwrote the signal with either white noise, $1/f$ noise, simulated sky or all three components. We used the fitted noise statistics from Section 3.3 to generate the realisations of the noise. We ran DESCART in traditional mode and used a 5% mask during the offset amplitude estimation.

In maps made from scans at a single elevation, most pixels are only crossed in two directions. When scans at higher elevations are included, most pixels are crossed in two further directions. Having fewer crossing angles per pixel lessens the mapmaker's ability to recover the sky and manifests particularly clearly near the north celestial pole. Figure 4.9 shows the signal error as a fraction of the true sky for the simulated dataset in Section 4.3.1, which includes scans at elevations 37° and 47°, and this new dataset, with scans at elevation 37° only. The pixel at the north pole has a very high hit count relative to all other pixels in the map and is therefore very heavily weighted. Any signal error in scans cutting through both the pole and the nearby Galactic plane must be counterbalanced by the scans cutting through the pole but not through the bright Galactic plane. When higher elevation scans are included there are extra rings

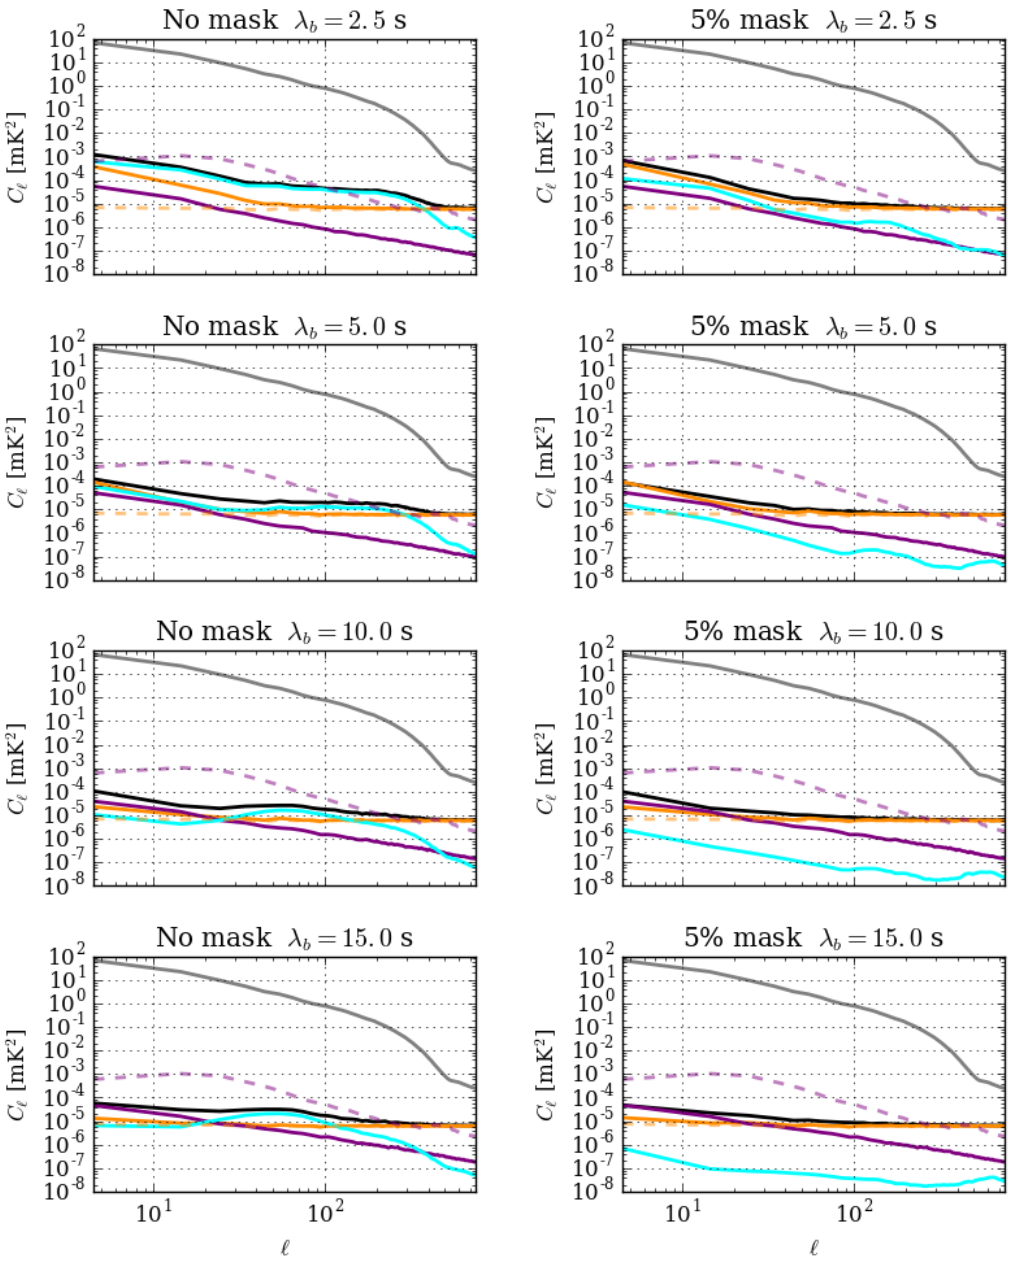

Figure 4.8: Angular power spectra of DESCART maps made from simulated data with either no masking during the destriping step (*left*) or applying a 5% mask (*right*). λ_b increases from *top* to *bottom*.
The angular power spectra of destriped maps made from TOD that only contain $1/f$ noise are *solid purple*, those of white noise are *solid orange* and the signal error is *solid cyan*. The total destriping error is shown in *solid black*.
For comparison, the angular power spectra of naive maps made from TOD that only contain sky signal, $1/f$ noise or white noise are shown in *solid grey*, *dashed purple* and *dashed orange* respectively.
This figure continues overleaf...

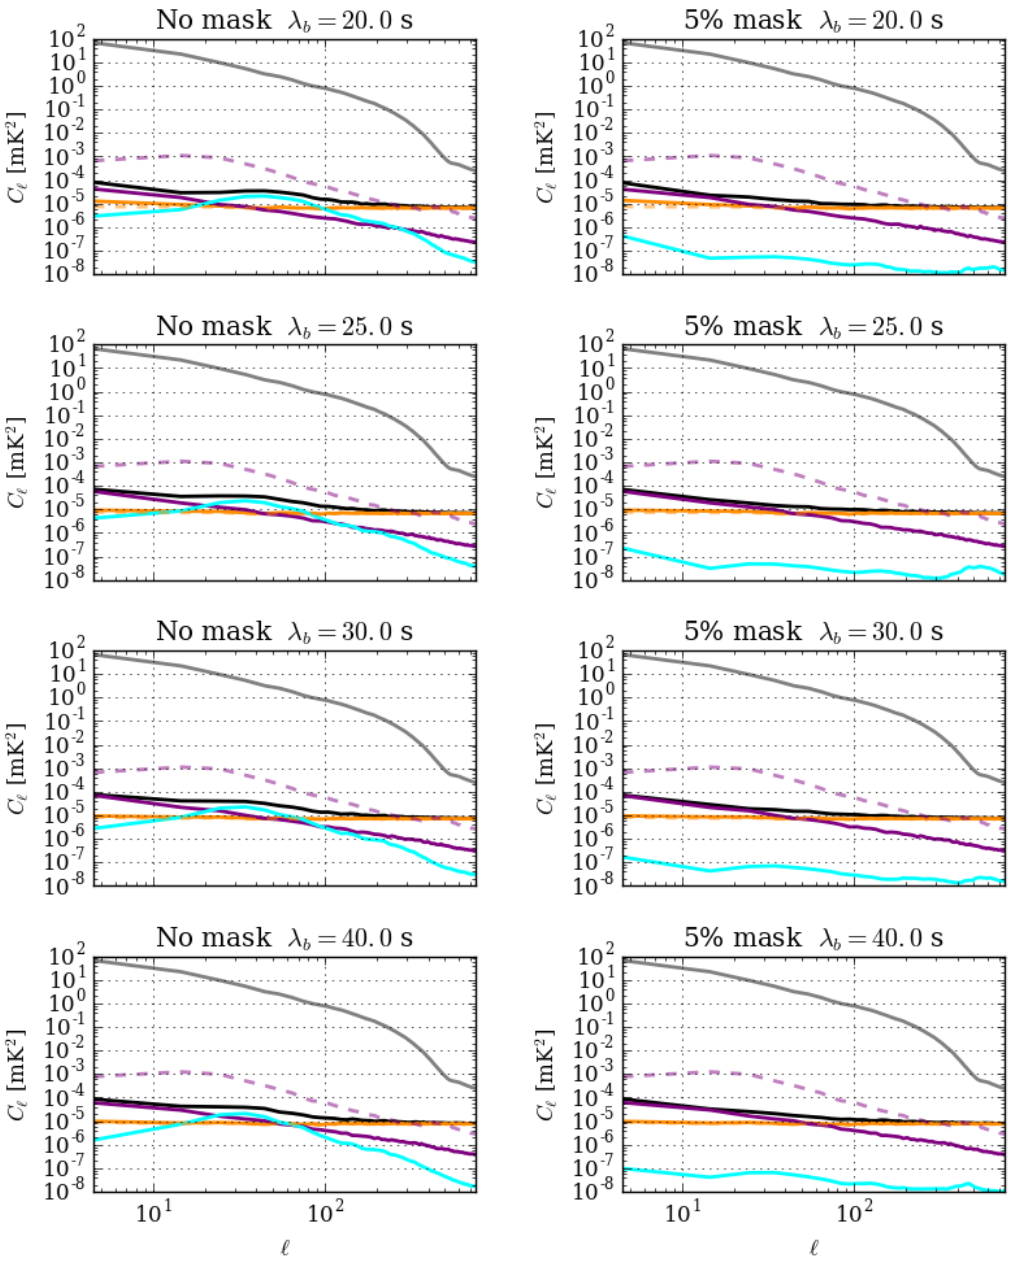

Figure 4.8: Angular power spectra of DESCART maps made from simulated data with either no masking during the destriping step (*left*) or applying a 5% mask (*right*). λ_b increases from *top* to *bottom*.
The angular power spectra of destriped maps made from TOD that only contain $1/f$ noise are *solid purple*, those of white noise are *solid orange* and the signal error is *solid cyan*. The total destriping error is shown in *solid black*.
For comparison, the angular power spectra of naive maps made from TOD that only contain sky signal, $1/f$ noise or white noise are shown in *solid grey*, *dashed purple* and *dashed orange* respectively.

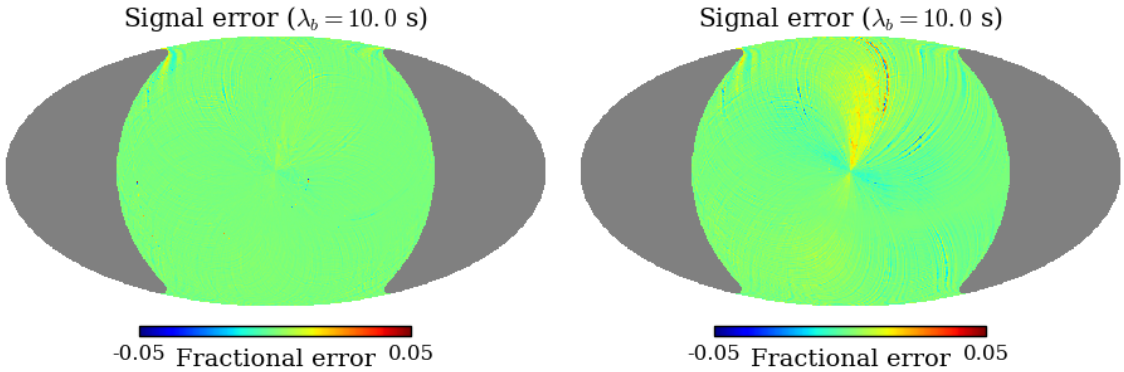

Figure 4.9: Signal error as a fraction of the true sky. *Left* simulated elevation 37° and 47° scans described in Section 4.3.1, *right* simulated elevation 37° scans described in Section 4.3.2. This shows the effect of the additional constraints from the elevation 47° data that increase the number of crossing angles per pixel and reduce destriping errors.

of high hit-count pixels around the north pole with extra crossing angles that can reduce the size of the error.

The angular power spectra of the DESCART maps made from this new set of simulated timestreams, which only include scans at Elevation 37° and used the fitted noise statistics from Section 3.3 to generate the noise, are plotted in Figure 4.10. The spectra clearly show that to accurately model the correlated noise and recover the large angular scales in our maps ($\ell \lesssim 100$) we will need to implement covariant destriping in DESCART. The spectra also show that a baseline offset length between 5–10 s is best able to remove the $1/f$ noise.

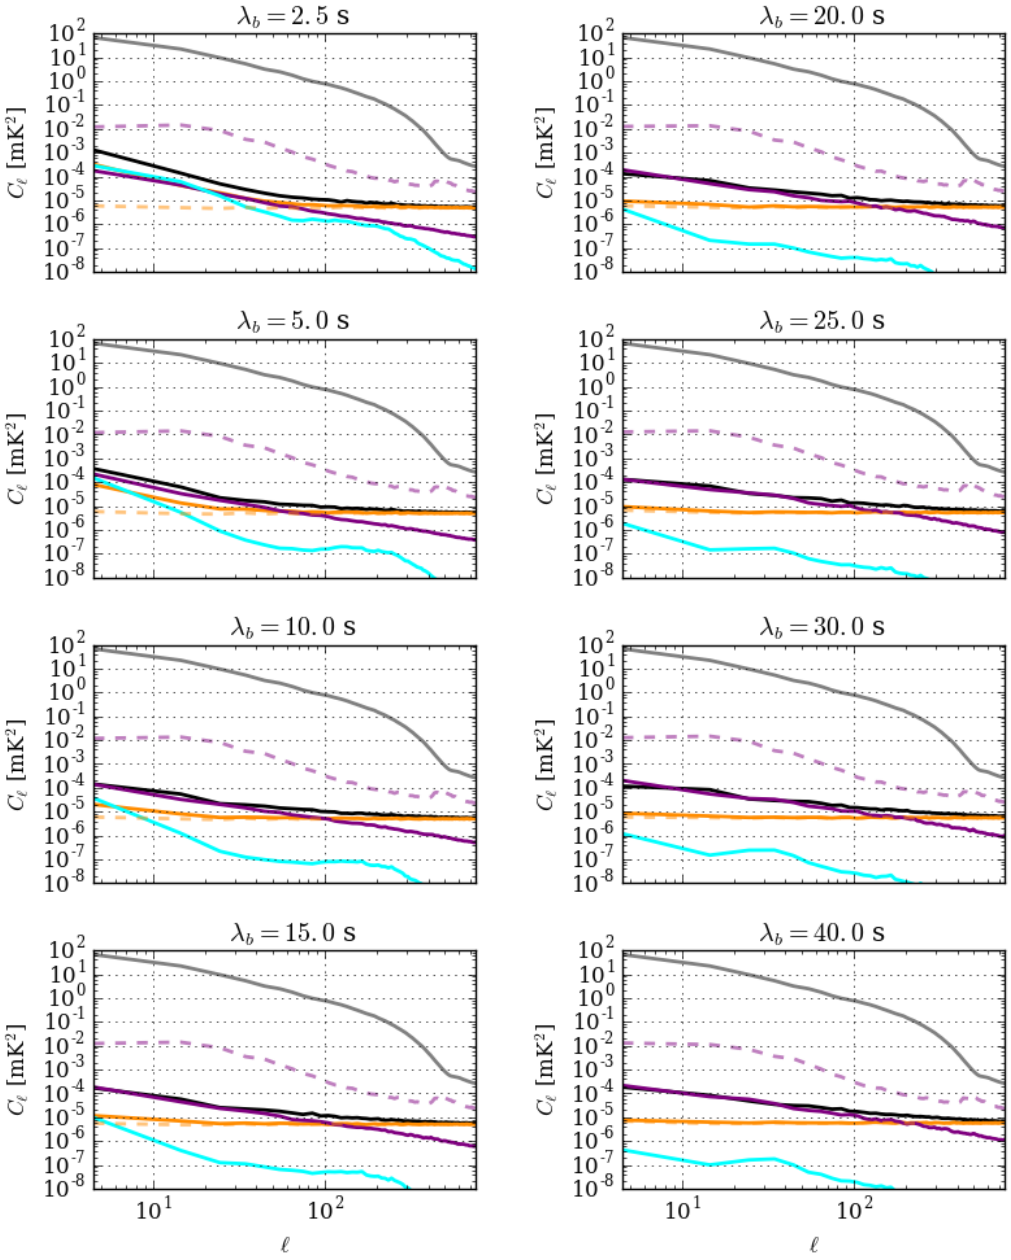

Figure 4.10: Angular power spectra of DESCART maps made from simulated data with noise realisations generated from the actual noise statistics of the C-BASS E37S60 dataset. All of the maps were made masking the brightest 5% of pixels. The various lines are the angular power spectra of the naive sky only map (*solid grey*), the total DESCART error (*solid black*), the residual correlated noise (*solid purple*), the residual white noise (*solid orange*), the signal error (*solid cyan*), the naively binned $1/f$ noise (*dashed purple*) and the naively binned white noise (*dashed orange*). This shows that the optimum baseline length is between 5–10 s.

4.4 Jackknife tests

We performed a suite of jackknife tests on C-BASS maps made from subsets of the E37S60 dataset in order to identify systematic effects in the data and to test the accuracy of the covariance matrices estimated by DESCART. We also tested the level of sun cutting by comparing maps made from the E37N, E37D, E37DS60 and E37DS90 datasets. This section contains a description of the jackknife test procedure (Section 4.4.1) and the results from using the tests on C-BASS data (Section 4.4.2).

In summary the tests show that:

- DESCART under-estimates the white noise power by 10–15%,
- the 60° Sun cut is not sufficient to remove contamination from the Sun's far out sidelobes,
- the gain is consistently calibrated at the 2% level,
- residual ground contamination in the Q and U maps is at the 1-σ level in pixel space,
- the maps are consistent with white noise and sky signal on scales $\ell > 60$ and at $\ell < 60$ there is residual $1/f$ noise, which is below the signal level.

4.4.1 Procedure

In a jackknife test the TOD is split into two independent and roughly-equal halves and maps are made from each half ($\mathbf{m_1}$ and $\mathbf{m_2}$). The difference between the maps is compared to the expected noise, which is determined from the map covariance matrices ($\boldsymbol{C_1}$ and $\boldsymbol{C_2}$). $\mathbf{m_1}$ and $\mathbf{m_2}$ should contain the same sky signal and independent realisations of the noise but might contain different contaminations that will appear in the difference maps. There are too many ways to split the data to test all possible combinations, so instead we chose a subset of the splits to target likely systematic errors.[2]

In order to compare the data in a like-for-like fashion all of the maps were produced in the same way. We ran DESCART in traditional mode with $\lambda_b = 5.0\,\mathrm{s}$ and masked the 5% brightest pixels during baseline amplitude estimation to produce $N_{\mathrm{side}} = 256$ maps. As a final filter to reject bad data from the maps, we discarded chunks that

[2]Whilst not technically jackknife tests, this name has been used in the literature by, for example, *Planck* (Zacchei et al., 2011) and BICEP (Chiang et al., 2010) and so we adopt the same name. They might more accurately be described as null tests.

were more than 10-σ away from the solution after 250 steps of the PCG algorithm. We took the map solution after either 1500 steps of the PCG algorithm or when the residuals were less than one millionth of the initial residual ($\tau = 10^{-6}$), which ever occurred first.

To identify systematic differences between the two data halves the difference maps need to be compared to the expected noise. It is important to note that the covariance matrices output by DESCART only include the I, Q and U variances and the QU covariance for each pixel and assume zero pixel-to-pixel correlations. Because DESCART was run in traditional mode there are still high levels of residual correlated noise in the C-BASS maps on large angular scales that are not accounted for in the covariance matrices. We do not expect to formally pass any of the jackknife tests because of the correlated noise that is unaccounted for in the map noise model.

We present the results from the jackknives in several different forms. Each method of presenting the jackknife test result has advantages and limitations that are described further below.

Difference maps The difference map, $\mathbf{D}$, is found by subtracting one map from another,

$$\mathbf{D} = \mathbf{m_1} - \mathbf{m_2}. \tag{4.10}$$

This is the simplest way of presenting the test result and, whilst simple, difference maps are not intuitive. The white noise level of $\mathbf{D}$ changes with the varying depth of observations across the sky and, without considering the noise levels, the varying level of fluctuations can easily be misinterpreted. Furthermore it is impossible to assess the significance of any structure without comparing it to the expected noise level.

Weighted difference maps The weighted difference map, $\mathbf{D_w}$, is found by scaling the difference map by the expected white noise in each pixel. The p^{th} pixel in $\mathbf{D_w}$ is therefore,

$$\mathbf{D_w}(p) = \frac{\mathbf{D}(p)}{\sqrt{{\sigma_1}^2(p) + {\sigma_2}^2(p)}}, \tag{4.11}$$

where ${\sigma_i}^2(p)$ is the expected white noise in the p^{th} pixel of map $\mathbf{m_i}$. Here we have assumed that the noise in both maps is independent and that their variances add. The pixel values in $\mathbf{D_w}$ are in units of the standard deviation of the expected white noise, $\sigma = \sqrt{{\sigma_1}^2(p) + {\sigma_2}^2(p)}$. Weighted difference maps contain all of the information in a jackknife test. Any systematic differences are easily compared to the noise level and their spatial structure can be used to try and identify the cause.

Histograms of weighted difference maps If $\mathbf{D_w}$ is consistent with white noise then the pixels should be distributed like the standard normal distribution. This can be difficult to judge directly from the weighted difference map but can easily be checked with a histogram of the pixel values. On their own, pixel histograms contain no information about the spatial distribution of any structure.

Reduced χ^2 values The reduced chi-squared value for a jackknife, $\chi^2_{\rm red}$, can be used to condense all the information into a single measure of whether the test has been passed or failed. We calculate $\chi^2_{\rm red}$ for the total intensity jackknives and the polarization jackknives separately via,

$$\chi^2_{\rm red} = (\mathbf{m_1} - \mathbf{m_2})^T (\tilde{\boldsymbol{C}}_\mathbf{1} + \tilde{\boldsymbol{C}}_\mathbf{2})^{-1} (\mathbf{m_1} - \mathbf{m_2})/\nu, \tag{4.12}$$

where ν is the number of degrees of freedom. For the I jackknives; $\nu = N_{\rm pix} - 1$ and for the $p^{\rm th}$ pixel $\tilde{\boldsymbol{C}}_\mathbf{j}(p) = \sigma_{\mathbf{II,j}}(p)$. For the polarization jackknives; $\nu = 2N_{\rm pix} - 1$ and for the $p^{\rm th}$ pixel,

$$\tilde{\boldsymbol{C}}_\mathbf{j}(p) = \begin{pmatrix} \sigma_{\mathbf{QQ,j}}(p) & \sigma_{\mathbf{QU,j}}(p) \\ \sigma_{\mathbf{UQ,j}}(p) & \sigma_{\mathbf{UU,j}}(p) \end{pmatrix}. \tag{4.13}$$

For each jackknife test, if the difference maps are consistent with the noise then the probability to exceed the estimated $\chi^2_{\rm red}$ (PTE) will be drawn from a uniform distribution between 0 and 1. Due to the very large values of ν for a typical C-BASS map, even slight mis-modelling of the noise will result in the data failing the test. We therefore do not currently expect the C-BASS maps to strictly pass the $\chi^2_{\rm red}$ tests because of residual correlated noise on large angular scales that is unaccounted for in the covariance matrices. The values of $\chi^2_{\rm red}$ still allow for a relative comparison between jackknives and we aim to pass these tests once covariant destriping is implemented in DESCART. By condensing all of the information in a jackknife test to a single number for I and a single number for Q and U, we have discarded information that can be used to identify the reason for any failure to pass the test so, while useful, they need to be used in combination with the other means of presenting the results. We only calculate $\chi^2_{\rm red}$ values using pixels further than 10° from the Galactic plane so that small gain differences between the two halves do not dominate the results.

Angular power spectra Systematics on large angular scales are likely to dominate the jackknife results. To determine the angular scales at which our maps are consistent with white noise, we compared the angular power spectra of the difference maps to the expected white noise and the sky signal. The white noise spectra were estimated

from the mean power spectrum of 20 realisations of the expected white noise and the sky signal spectrum was estimated from a sky map made with all the elevation 37° data.

We used the publicly available PolSpice package to calculate the angular power spectra, inverse variance weighting the difference maps and masking certain regions of the sky. The masks, shown in Figure 4.11, excluded the brightest 1% of pixels and also the 5% of pixels with the greatest hit counts. We masked the brightest pixels because small gain errors will leave point like features in the difference maps, which introduce broad features into the power spectra that dominate over a large ℓ range. We masked the pixels with the greatest hit counts because the uneven sky coverage around the NCP and lower declination bound introduces ringing into the power spectrum.

In future cosmological analysis of the polarization data we will most likely be interested in the EE and BB angular power spectra. However, when investigating structure in jackknife maps it is more intuitive to consider Q and U maps as separate scalar fields.

To help guide the eye we scaled the white noise and sky spectra in power to produce the best least-squares fit to the jackknife spectrum. The power spectra show which angular scales have structure that is inconsistent with white noise but not the orientation of that structure.

Gain differences The data halves may be inconsistently calibrated and the sky signal in one map could differ from the other by a multiplicative gain factor. In the difference maps this would appear as a contaminating signal with the morphology of the sky but scaled in amplitude. Any gain difference will therefore be most obvious in the pixels with the greatest sky signal and so we estimate the gain difference from the brightest regions of the sky. We calculated $\chi^2_{\rm red}$ for the difference between two maps from the brightest 1% of pixels after applying a range of multiplicative factors to the second map, treating the I, Q and U maps independently and neglecting any covariance terms. The gain factor that minimised the variance is then our estimate of the relative gain difference between the two maps. In Section 4.4.2 we summarise the results in a table of gain factors and $\chi^2_{\rm red}$ values and for each test also show plots of $\chi^2_{\rm red}$ against the gain factor with the locations of the minima identified.

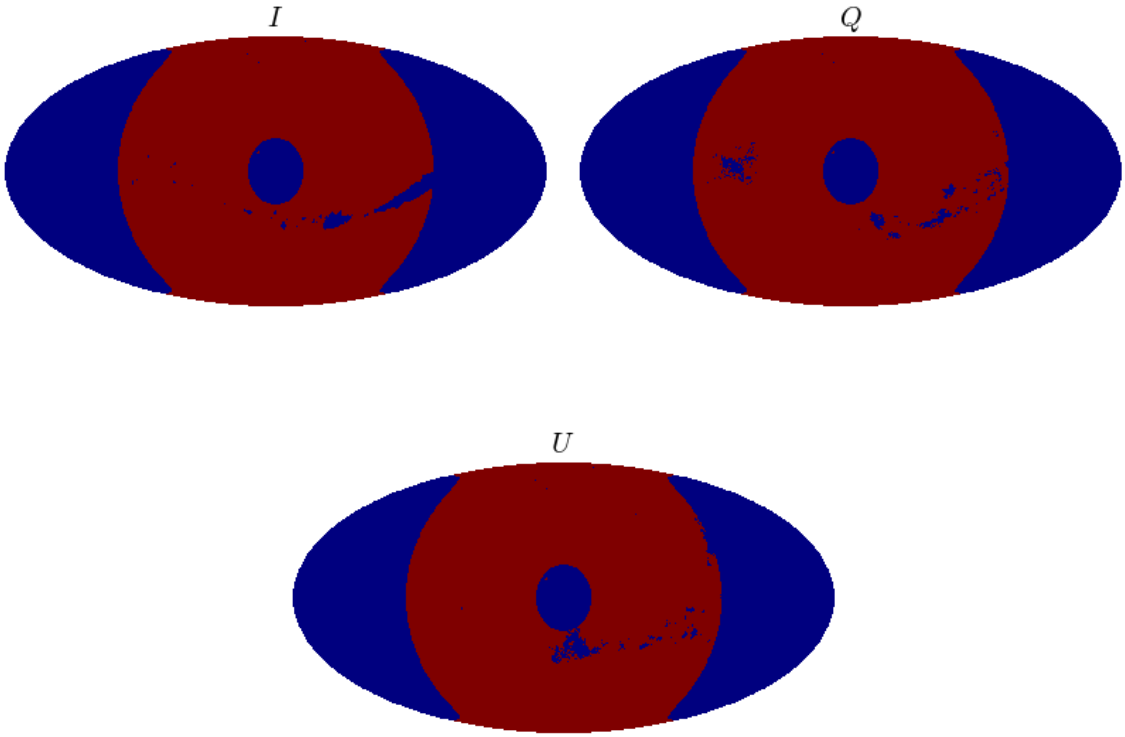

Figure 4.11: Masks applied to jackknife maps when estimating angular power spectra using PolSpice. The maps are in celestial coordinates rotated so that the NCP is at the centre of the projection. The *red* regions are included and the *blue* regions are excluded. The masks exclude the brightest 1% of the sky and the 5% of pixels with the greatest hit counts.

4.4.2 Jackknife results

The five jackknife tests described in this document are an even/odd split in Section 4.4.2.1, day/night splits in Section 4.4.2.2, 1$^{\text{st}}$ half/2$^{\text{nd}}$ half in Section 4.4.2.3, seasons in Section 4.4.2.4 and receiver channels in Section 4.4.2.5. We draw conclusions from the results from these tests in Section 4.4.2.6 and explain the extra contributions to the pixel uncertainties that we will include in scientific analysis. Except for the day/night test, all maps were made from the E37S60 dataset.

Table 4.1 contains a summary of the jackknife tests that we have performed on the C-BASS North elevation 37° datasets. For each test the table contains the targeted systematic(s), χ^2_{red} and ν for I and polarization.

Table 4.2 contains a summary of the gain differences between the data halves, estimated from the brightest 1% of pixels. We did not estimate the gain differences

from the season jackknives as large fractions of the sky are missing.

Table 4.1: A list of the jackknife tests that we performed on the C-BASS North elevation 37° data.

	I		Q, U		
Jackknife	$\chi^2_{\rm red}$	ν	$\chi^2_{\rm red}$	ν	Systematic
Even/odd schedules	1.14	417814	1.12	835629	null test
Day/night	...	...	...	...	ground, pointing & sun
no sun flag	16.59	417809	7.11	835619	
60° sun flag	1.72	417787	1.24	835575	
90° sun flag	1.88	416792	1.22	833585	
1st half/2nd half	2.16	417757	1.30	835515	long timescale drifts
Seasons	...	...	...	...	ground
summer/spring	1.85	348906	1.36	697813	
summer/autumn	1.69	358945	1.28	717891	
summer/winter	2.70	377905	1.27	755811	
spring/autumn	1.65	356338	1.28	712677	
spring/winter	1.47	375319	1.24	750639	
autumn/winter	1.95	385342	1.26	770685	
Receiver channels⋆	1.11	417822	0.11	835645	receiver chains

⋆ Data halves not independent in polarization.

Table 4.2: The gain factors that minimise the $\chi^2_{\rm red}$ of the difference between the brightest 1% of pixels in Stokes I, Q and U. The table lists the gain factor, the minimum value of $\chi^2_{\rm red}$ and the number of degrees of freedom, ν.

Jackknife	I			Q			U		
	Gain	$\chi^2_{\rm red}$	ν	Gain	$\chi^2_{\rm red}$	ν	Gain	$\chi^2_{\rm red}$	ν
Even/odd schedules	**1.01**	8.01	7864	**1.00**	1.14	7856	**1.00**	1.20	7856
Day/night (60° sun flag)	**0.99**	5.39	7864	**1.00**	1.36	7829	**1.03**	1.24	7829
1st half/2nd half	**1.01**	8.53	7864	**0.97**	1.33	7799	**1.01**	1.49	7799
Receiver channels⋆	**1.01**	2.03	7864	**0.91**	0.704	7864	**0.91**	0.305	7864

4.4.2.1 Even/odd numbered schedules

We want at least one jackknife test with no possible systematic differences between each data half. Such a jackknife could then be used as a baseline from which to compare other tests. One such split can be obtained by listing the reduced TOD FITS files alphanumerically by start date and separating the even and odd numbered lines into two separate lists. An example list of start dates split in this manner is shown in Table 4.3, where the dates take the form *DD-MMM-YYYY:hh:mm:ss*.

Any systematic differences between the data in these two file lists would have to manifest themselves on roughly 100 minute timescales and also somehow depend on the alphanumerical ordering of the reduced TOD FITS files. We can not identify any systematic effects that would exhibit this behaviour and therefore expect this to be the "best case jackknife" - free of systematics.

The maps made from each data half along with the difference maps and weighted difference maps are shown in Figure 4.12. Across much of the sky the difference maps appear noise like and the histograms in Figure 4.14 confirm this. There are excess residuals in the difference map in the brightest few pixels of the sky but these are small in comparison to the sky brightness temperature of the pixels.

The (weighted) difference maps show scan-synchronous structure radiating from the NCP. Whilst the magnitude of this excess is typically small compared to the sky signal, it is significant compared to the very high depth of observations in this region. The structure takes the form of residual correlated noise that was let through by the destriper as well as signal error caused by the small number of crossing angles (cf. Figure 4.9). It would likely be reduced by including scans at higher elevations to increase the number of crossing angles through each pixel and also by implementing covariant destriping in DESCART.

The angular power spectra (Figure 4.13) show that the jackknives are white-noise dominated in I on multipoles larger than $\ell \sim 50$ and in Q and U on multipoles larger than $\ell \sim 40$.

The scaled sky and white-noise spectra that give the best least-squares fit to the difference jackknife indicate that the white-noise power has been underestimated by around 10%, 19% and 15% in I, Q and U respectively. These are consistent with the χ^2_{red} values, which suggest that σ_{II} has been underestimated by $\sim 14\%$ and that σ_{QQ}, σ_{UU} and σ_{QU} have been under estimated by $\sim 12\%$.

At low ℓ, say $\ell \sim 10$, the residual power in I is approximately three orders of magnitude smaller than the sky signal. The residual power in Q is approximately

Table 4.3: An example list of schedule start dates, listed alphanumerically, split into two halves based on whether their line numbers are even or odd. The dates take the form *DD-MMM-YYYY:hh:mm:ss*.

Line number	Half 1 (odd lines)	Half 2 (even lines)
1	01-Aug-2013:00:28:49	
2		01-Aug-2013:02:14:08
3	01-Aug-2013:04:06:50	
4		01-Aug-2013:06:06:33
5	01-Aug-2013:07:52:00	
6		01-Aug-2013:09:52:55
7	01-Aug-2013:11:38:18	
8		01-Aug-2013:13:39:02
9	01-Dec-2012:02:08:29	
10		01-Dec-2012:03:49:40
11	01-Dec-2012:06:08:29	
12		01-Dec-2012:07:49:07
13	01-Dec-2012:10:08:27	
14		01-Dec-2012:11:49:10
15	01-Dec-2012:14:07:04	
16		01-Dec-2012:23:47:25
17	01-Dec-2013:01:22:35	
18		01-Dec-2013:03:23:26
19	01-Dec-2013:14:35:06	
20		01-Dec-2013:16:20:28

two orders of magnitude smaller than the sky signal. The residual power in U is just under two orders of magnitude smaller than the sky signal

Figure 4.15 shows $\chi^2_{\rm red}$ of the difference between the two halves in the brightest 1% of pixels as a function of the multiplicative gain factor applied to the map made from odd numbered files. The gain factors that minimise $\chi^2_{\rm red}$ are 1.01, 1.00 and 1.00 in I, Q and U respectively. Any gain difference between these two halves is at the 1% level. This is consistent with the angular power spectra where the jackknife power at low ℓ has a different slope to the undifferenced sky power.

There are no other obvious systematic differences between these two halves. The jackknife test is failed because DESCART underestimates the white noise level, lets through residual correlated noise and introduces destriping errors due to the small number of crossing angles in most pixels.

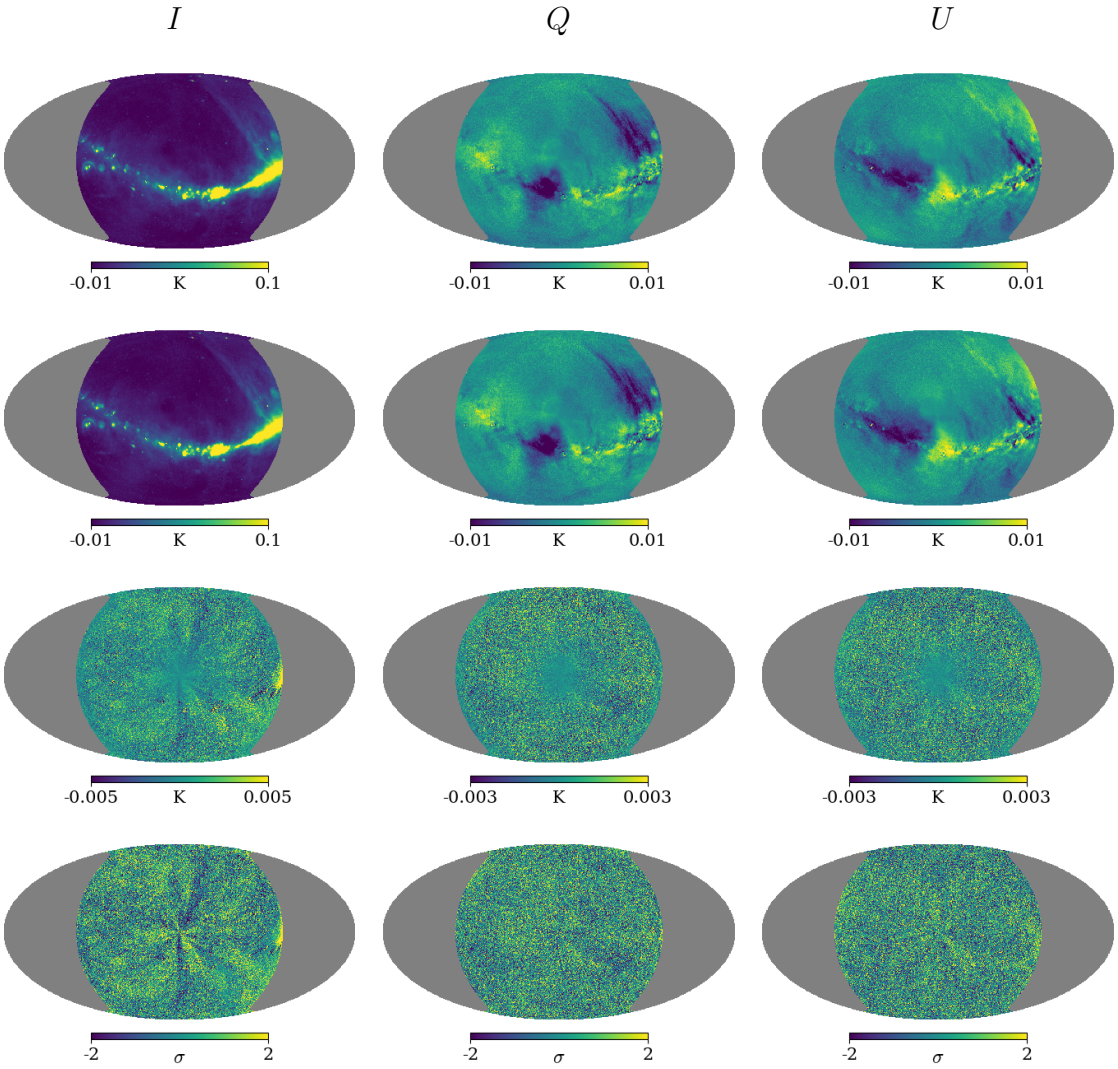

Figure 4.12: Even/odd numbered schedule jackknife maps. From *top* to *bottom* are the I, Q and U maps of odd numbered schedules (*first row*), maps of even numbered schedules (*second row*), difference maps (*third row*) and weighted difference maps (*fourth row*). From *left* to *right* the maps are of Stokes I, Q and U data.

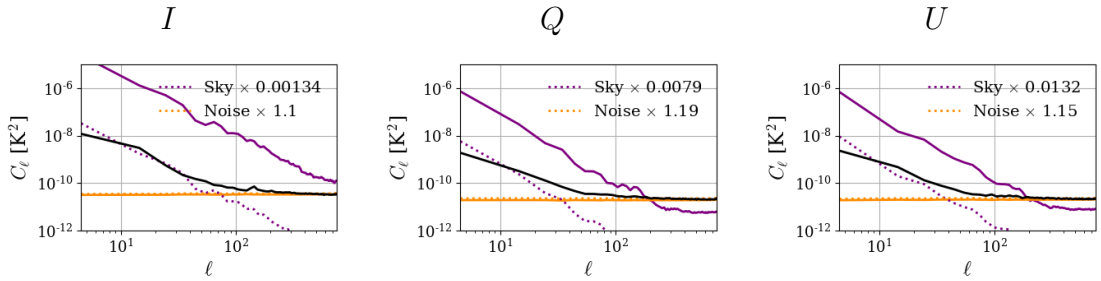

Figure 4.13: Even/odd numbered schedule jackknife angular power spectra. From *left* to *right*; Stokes I, Q and U. The *solid black line* is the jackknife, the *solid orange line* is the expected white noise and the *solid purple line* is the sky signal. The *dashed lines* show the scaled white noise and sky signal that give the best least-squares fit to the data.

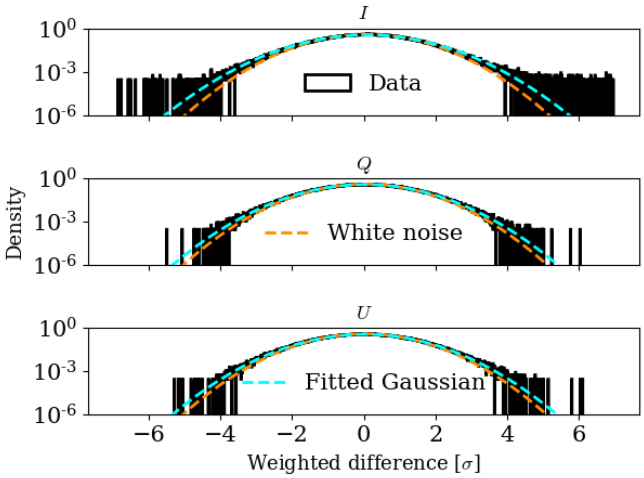

Figure 4.14: Even/odd numbered schedule jackknife weighted difference map histograms. The *dashed orange* line is the expected white noise from the covariance matrix and the *dashed cyan* line is the best fit gaussian to the data.

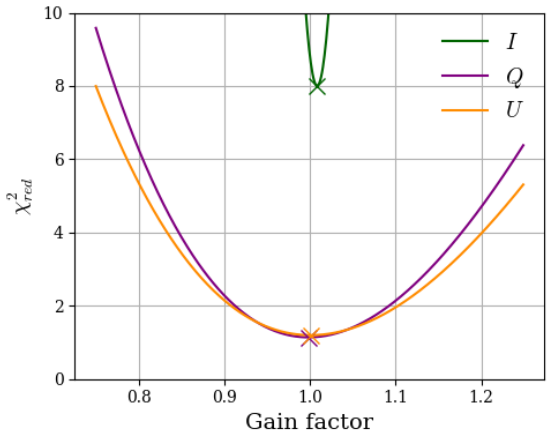

Figure 4.15: Even/odd numbered schedule jackknife gain factors. The curves show the $\chi^2_{\rm red}$ of the difference between the brightest 1% of pixels as a function of the gain factor applied to the map made from the odd numbered lines for I (*green*), Q (*purple*) and U (*orange*). The *crosses* show the factors that minimise $\chi^2_{\rm red}$.

4.4.2.2 Sun cuts

The Sun is the brightest astronomical source at 5 GHz and is also the only astronomical source that saturates the C-BASS receiver. Even when the Sun is in the telescope's far out sidelobes it is still bright enough to significantly contaminate the C-BASS maps. The RFI flagging subroutine flags TOD samples where the telescope main beam or first few sidelobes crossed through the sun but not when the Sun was in the far out sidelobes. We therefore apply a further level of flagging and remove TOD samples that fall within a predefined angular distance from the Sun. In the other jackknife tests we have used the E37S60 dataset (where samples within 60° of Sun are flagged). To test whether this level of sun flagging is sufficient we compared maps made from the E37N dataset (night-time only data) to maps made from the E37D, E37DS60 and E37DS90 datasets (daytime only data with 0°, 60° and 90° sun cuts respectively).

The difference maps are shown in Figure 4.16. In the case of no sun flagging (*top row*) there is obvious excess emission in daytime maps in a ring along the ecliptic plane. The ripple-like features either side of the brightest ring are individual sidelobes.

This obvious structure is removed by both the 60° and 90° sun cuts, shown in the *middle* and *bottom* rows respectively. The daytime map made from scans with a 90° sun cut has significantly less data than the 60° sun cut and so the residual noise level is higher and any contaminations from, for example, RFI average down less well.

The weighted difference maps in Figure 4.17 show that there is no significant improvement in using the 90° cut over the 60° cut. The $\chi^2_{\rm red}$ values in Table 4.1 and the angular power spectra in Figure 4.18 show that the daytime map made with a 90° sun cut is more discrepant from the night-time map than the daytime map made with a 60° sun cut. This is likely because there is far less TOD going into the map with the more restrictive 90° sun cut and so residual rogue RFI events are not averaged away as effectively.

As well as solar emission, the day/night jackknives test the ground removal subroutine. The pipeline constructs one ground template for each 24 hour period, which includes both daytime and night-time data. The thermal emission from the ground may change from day to night as the ground warms up and cools down. The ground could also reflect emission from the Sun. A feature at a constant azimuth is mapped to a line of constant declination and so residual ground emission would manifest as rings at constant declination centred on the NCP. In the U weighted difference maps there are faint rings at constant declinations with typical pixel levels of roughly $\sim \pm 1\sigma$. Whilst these are small in individual pixels, they are significant along the

angular scales they correspond to. We need to investigate this residual ground to try and reduce it further. However, the residual ground emission in Q and U is much smaller than the scan synchronous structure radiating from the NCP.

The angular power spectra in Figure 4.18 show that for daytime data with a 60° Sun cut, the I jackknife map is white noise dominated down to $\ell \sim 100$ and the Q and U jackknife maps are white noise dominated down to $\ell \sim 50$.

The scaled sky and white noise spectra that give the best least-squares fit to the difference jackknife indicate that the white noise power has been underestimated by around 10% in I and 15% in Q and U. The white noise levels would have to be scaled more than this to reduce the $\chi^2_{\rm red}$ values to one. The excess power on large angular scales is dominating these $\chi^2_{\rm red}$ values. For example, at $\ell = 10$ the residual power in the day/night jackknives (with either 60° or 90° sun flagging) is an order of magnitude larger than in the even/odd split. This extra power is likely from residual ground and solar emission.

Figure 4.19 shows $\chi^2_{\rm red}$ of the difference between the brightest 1% of pixels in the maps made from E37N and E37DS60 as a function of the multiplicative gain factor applied to the map made from night-time data. The gain factors that minimise $\chi^2_{\rm red}$ are 0.99, 1.00 and 1.03 in I, Q and U respectively. This suggests that any gain difference between the two halves is at the 1–3% level.

This test has shown that we should only use night-time data and that in future runs of the pipeline we should only use night-time data to construct the ground templates.

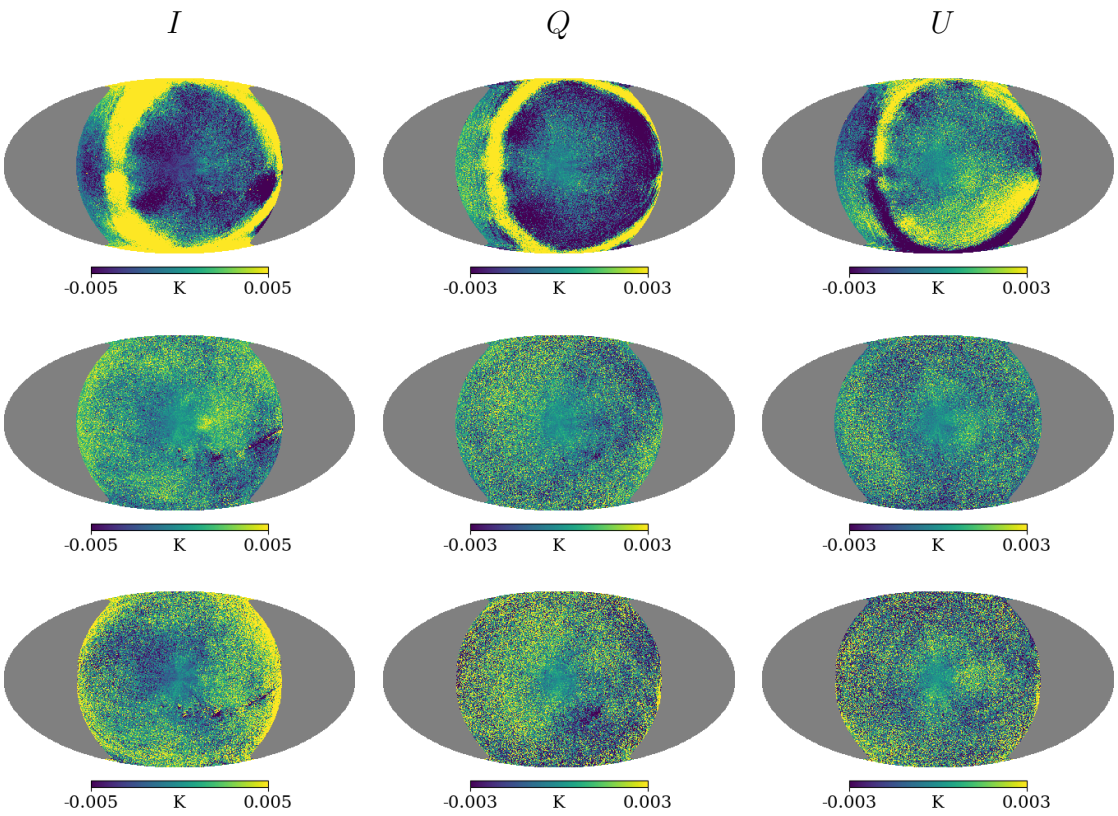

Figure 4.16: Day/night jackknife difference maps. From *top* to *bottom*; no sun cut (*first row*), 60° sun cut (*second row*) and 90° sun cut (*third row*). From *left* to *right*; Stokes I, Q and U.

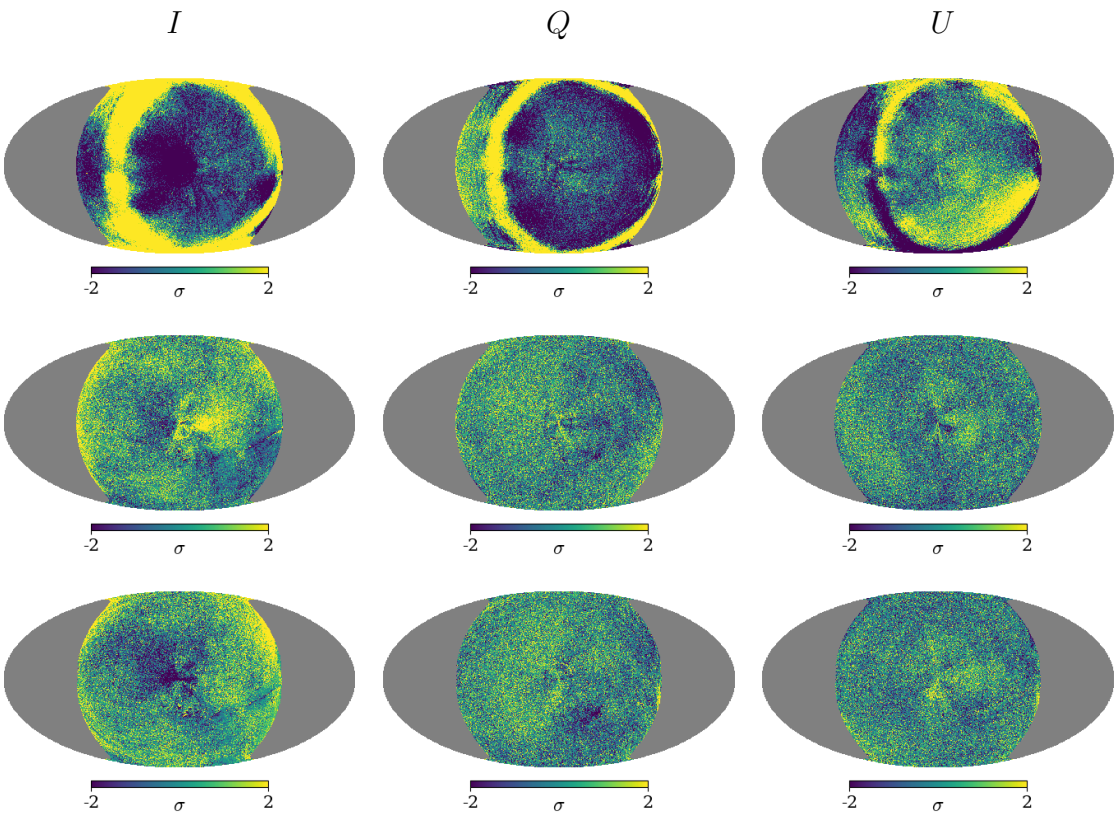

Figure 4.17: Day/night jackknife weighted difference maps. From *top* to *bottom*; no sun cut (*first row*), 60° sun cut (*second row*) and 90° sun cut (*third row*). From *left* to *right*; Stokes I, Q and U.

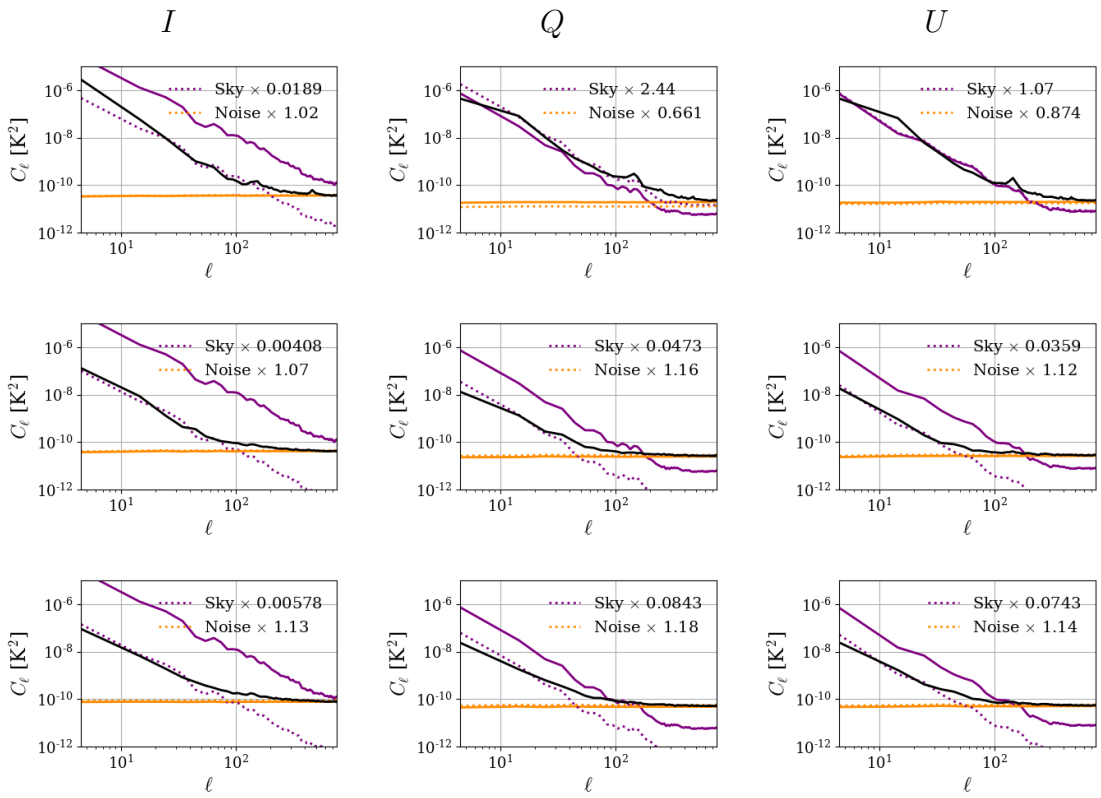

Figure 4.18: Day/night jackknife angular power spectra. From *top* to *bottom*; no sun cut (*first row*), 60° sun cut (*second row*) and 90° sun cut (*third row*). From *left* to *right*; Stokes I, Q and U. The *solid black line* is the jackknife, the *solid orange line* is the expected white noise and the *solid purple line* is the sky signal. The *dashed lines* show the scaled white noise and sky signal that give the best least-squares fit to the data.

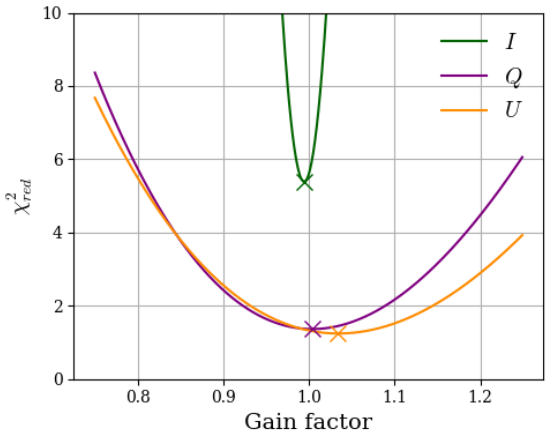

Figure 4.19: Day/night with a 60° sun cut jackknife gain factors. The curves show the $\chi^2_{\rm red}$ of the difference between the brightest 1% of pixels as a function of the gain factor applied to the map made from night-time data for I (*green*), Q (*purple*) and U (*orange*). The *crosses* show the factors that minimise $\chi^2_{\rm red}$.

4.4.2.3 First half/second half

We used different astronomical calibrators at different times of the year and so the C-BASS TOD could be inconsistently calibrated. The choice of calibrator was determined by what was observable at the time. When TauA was above the horizon it was used to calibrate the TOD, when it was below the horizon the secondary calibrator CasA was used. The reference flux values of each calibrator may not be on the exact same flux scale and this could introduce long timescale drifts into the TOD.

We tested for these differences by splitting the TOD into the first and second halves of the observing run. The first half of the elevation 37° scans fell between 01-Nov-2012 and 30-Nov-2013 inclusive, the second half of the elevation 37° scans fell between 01-Dec-2013 and 30-Sep-2014 inclusive.

For this jackknife $\chi^2_{\rm red} = 2.16$ for I and $\chi^2_{\rm red} = 1.3$ for polarization. These values are significantly worse than the even/odd test (1.14 in I and 1.12 in polarization) and the day/night test (1.72 in I and 1.24 in polarization). The histograms in Figure 4.22 show that far more pixels in this jackknife are inconsistent with white noise when compared to the even/odd test (Figure 4.14) and also that different global offsets between the two maps may be responsible for a large fraction of $\chi^2_{\rm red}$ in I.

The bright Galactic plane is clear in the difference and weighted difference maps shown in Figure 4.20. The large positive residuals in the brightest regions of the sky show that there is a gain difference between the two halves, with the first half brighter than the second. Figure 4.23 shows the gain factors that must be applied to the map made from the $2^{\rm nd}$ half of the data to minimise $\chi^2_{\rm red}$ in the difference maps of the brightest 1% of pixels. The best gain factors are 1.01, 0.97 and 1.01 in I, Q and U respectively.

Once again, the angular power spectra shown in Figure 4.21 show that in I the jackknife is white-noise dominated down to $\ell \sim 100$ and in Q and U to $\ell \sim 50$ and that the white-noise power is underestimated by 10–20%.

At the time of writing the calibration of the dataset is being tested in detail and we are developing the calibration subroutine to ensure all TOD is on the same, self-consistent, flux scale.

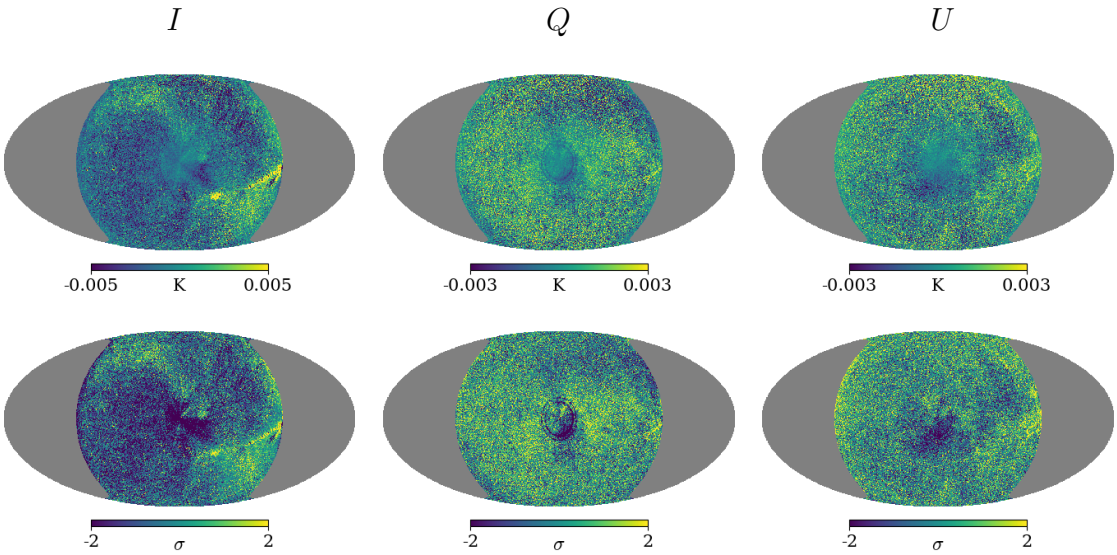

Figure 4.20: First and second halves of the data jackknife maps. The difference maps are in the *top row* and weighted difference maps are in the *bottom row*. From *left* to *right*; Stokes I, Q and U.

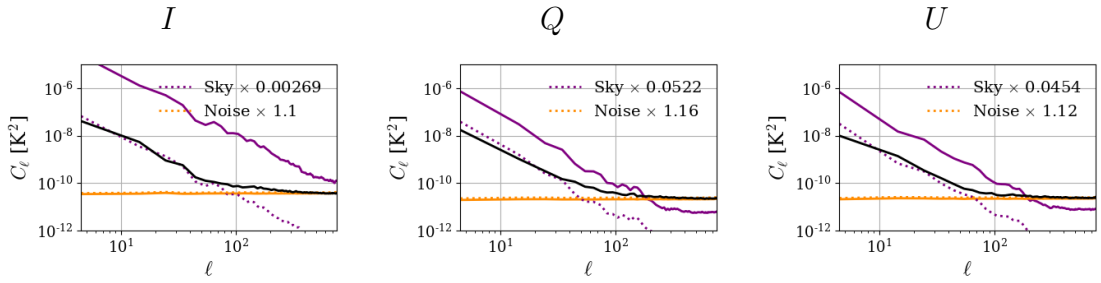

Figure 4.21: First and second halves of the data jackknife angular power spectra. From *left* to *right*; Stokes I, Q and U. The *solid black line* is the jackknife, the *solid orange line* is the expected white noise and the *solid purple line* is the sky signal. The *dashed lines* show the scaled white noise and sky signal that give the best least-squares fit to the data.

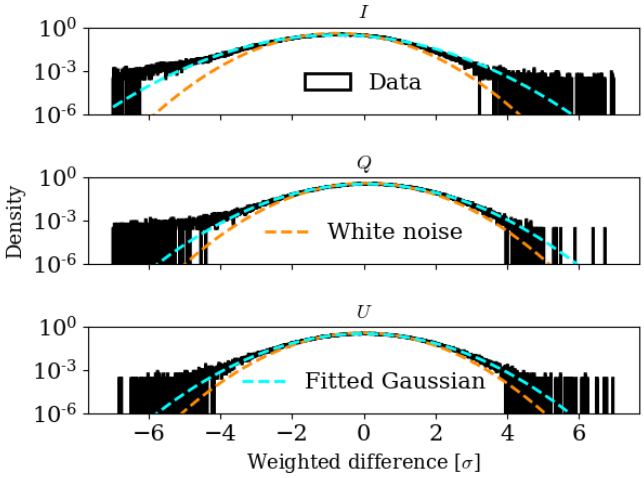

Figure 4.22: First and second halves of the data jackknife weighted difference histograms.

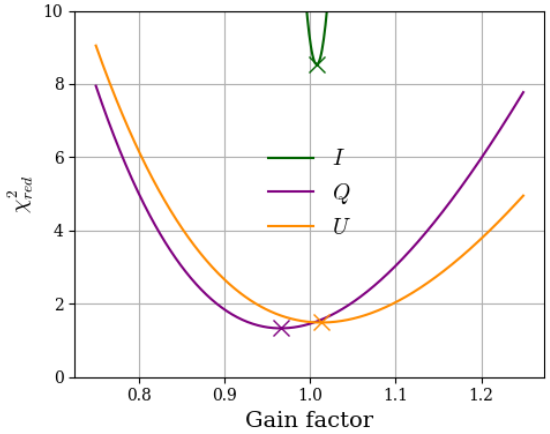

Figure 4.23: First and second halves of the data jackknife gain factors. The curves show the $\chi^2_{\rm red}$ of the difference between the brightest 1% of pixels as a function of the gain factor applied to the map made from the 2$^{\rm nd}$ half of the data for I (*green*), Q (*purple*) and U (*orange*). The *crosses* show the factors that minimise $\chi^2_{\rm red}$.

4.4.2.4 Seasons

Seasonal jackknives test the ground removal subroutine. The thermal emission from the ground will likely change with the ambient temperature, level of vegetation and moisture levels, all of which are strongly correlated with the time of year. Ground emission from a fixed azimuth maps to a ring of constant declination. If the residual ground emission changes between seasons this would manifest in the difference maps as rings centred on the NCP. We define the seasons as: March to May (spring), June to August (summer), September to November (autumn) and December to February (winter).

In the E37S60 dataset TOD within 60° of the Sun was flagged and excluded from the map. Over a three month period the Sun blocks a fraction of the sky and so there are missing patches in the seasonal maps. The residuals in the difference maps in Figure 4.24 are larger around these patches because they are observed less deeply. These features disappear in the weighted difference maps in Figure 4.25 once the relative weighting of the pixels is accounted for.

Because there is much less data in each season's map, RFI and other contaminating events do not get averaged away and these features dominate the residual maps.

There is no obvious residual ground in the weighted difference maps. There is, however, large scale structure, again consistent with residual correlated noise and destriping errors. The destriping errors are worse here compared to jackknives with more data as more of the pixels are crossed at fewer angles.

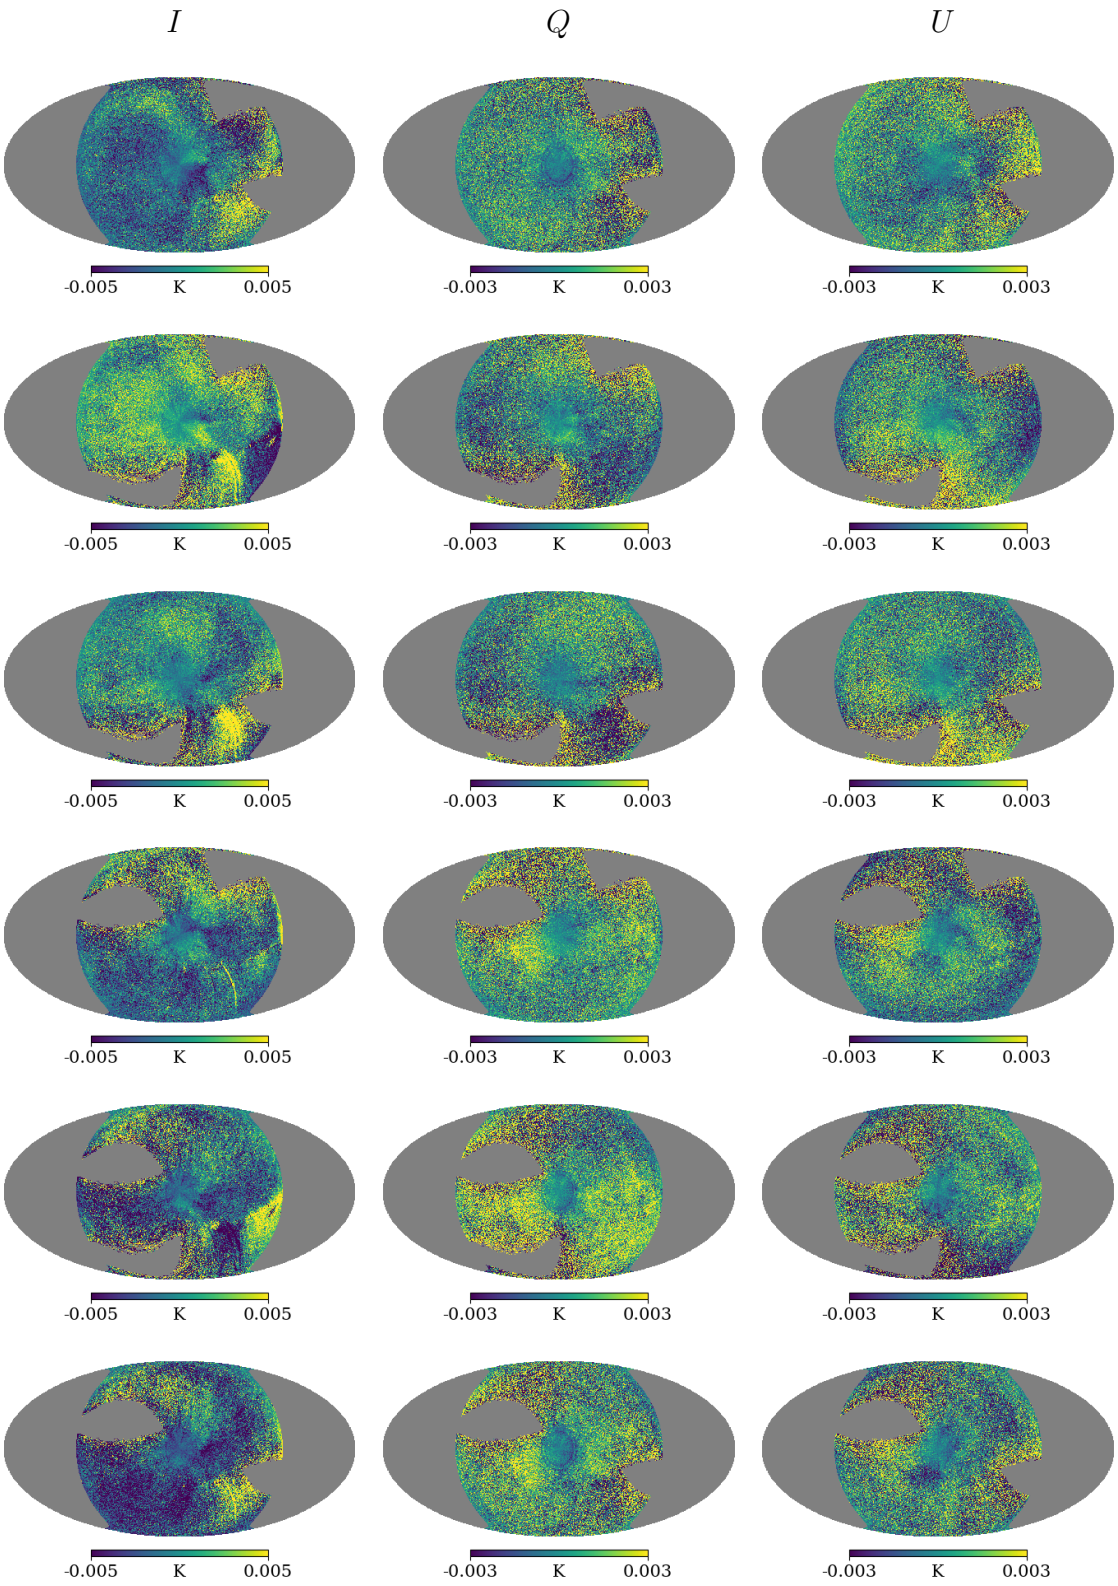

Figure 4.24: Seasonal jackknife difference maps. From *top* to *bottom*, autumn-winter (*first row*), spring-autumn (*second row*), spring-winter (*third row*), summer-autumn (*fourth row*), summer-spring (*fifth row*) and summer-winter (*sixth row*). From *left* to *right*; Stokes I, Q and U.

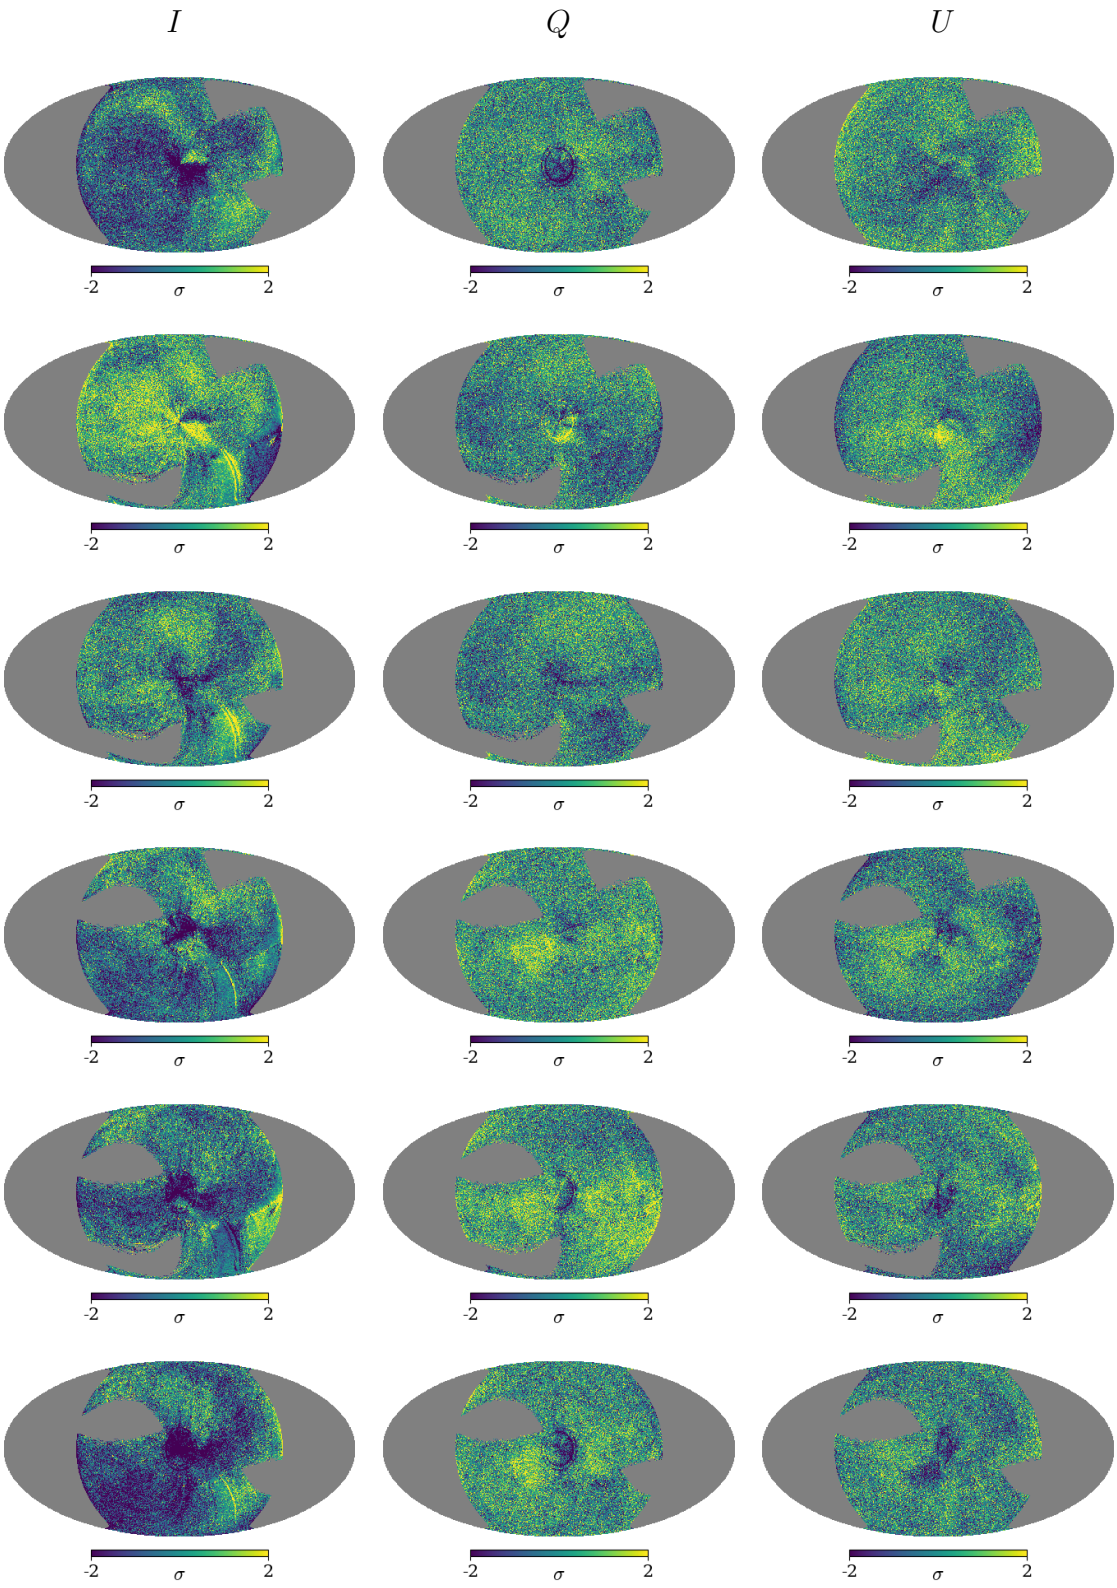

Figure 4.25: Seasonal jackknife weighted difference maps. From *top* to *bottom*, autumn-winter (*first row*), spring-autumn (*second row*), spring-winter (*third row*), summer-autumn (*fourth row*), summer-spring (*fifth row*) and summer-winter (*sixth row*). From *left* to *right*; Stokes I, Q and U.

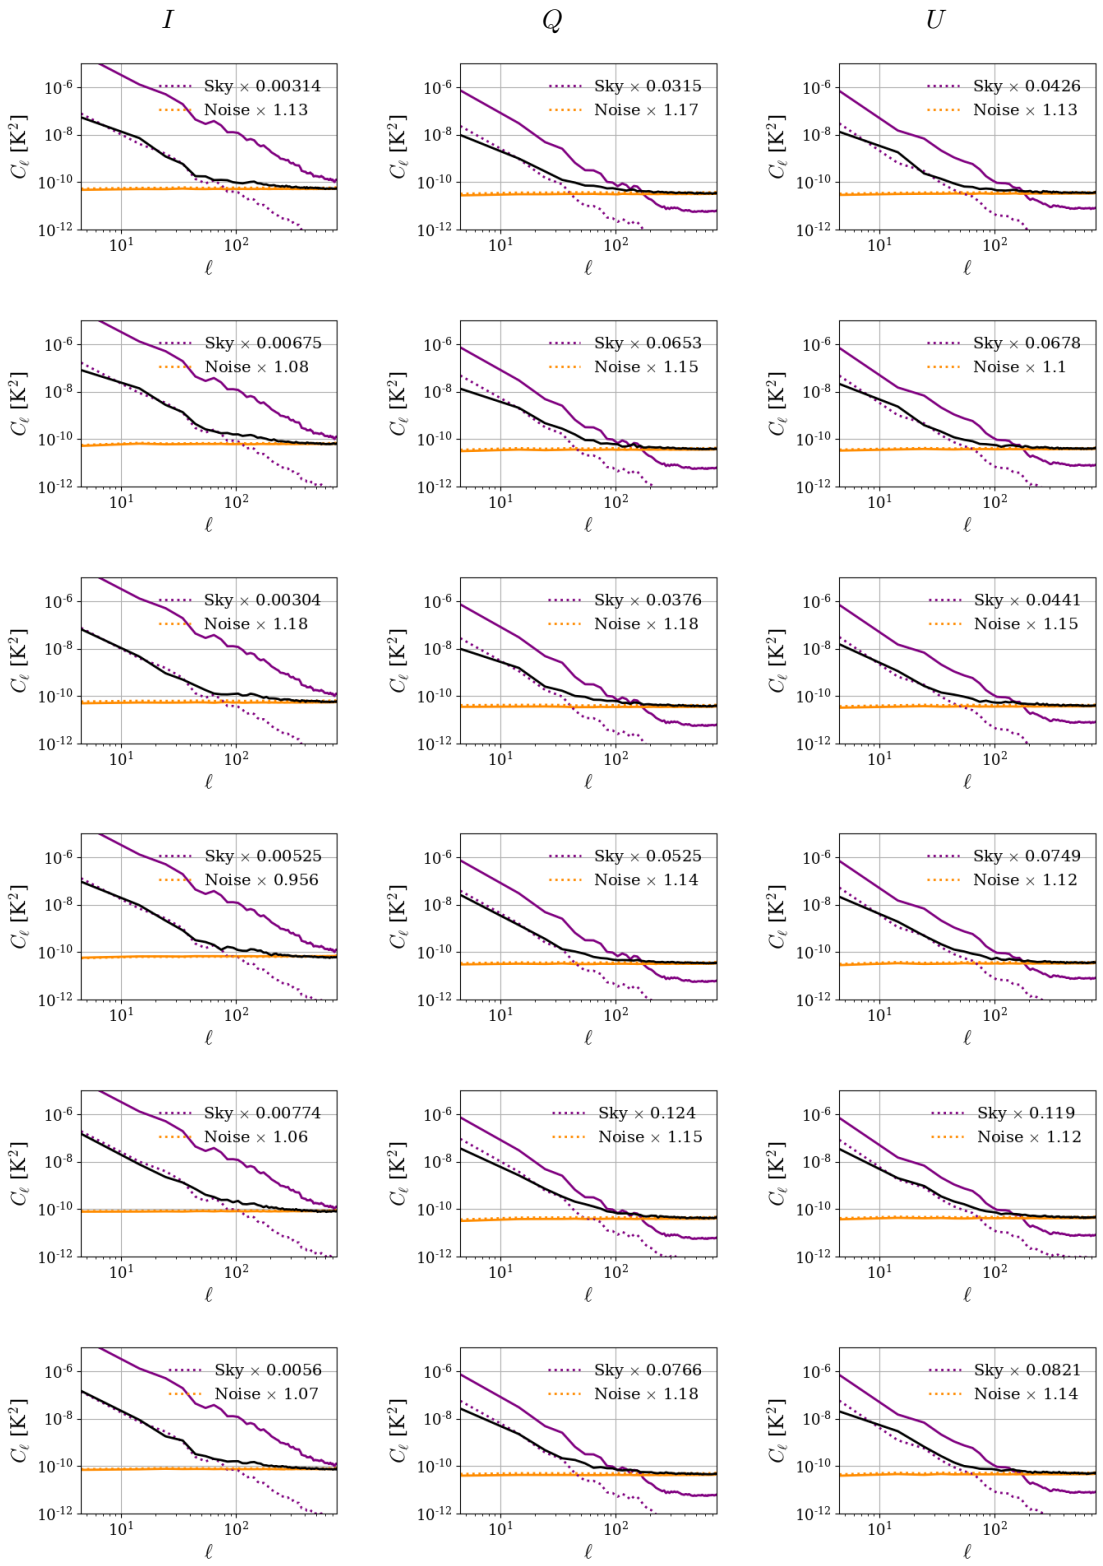

Figure 4.26: Seasonal jackknife angular power spectra. From *top* to *bottom*, autumn-winter (*first row*), spring-autumn (*second row*), spring-winter (*third row*), summer-autumn (*fourth row*), summer-spring (*fifth row*) and summer-winter (*sixth row*). From *left* to *right*; Stokes I, Q and U. The *solid black line* is the jackknife, the *solid orange line* is the expected white noise and the *solid purple line* is the sky signal. The *dashed lines* show the scaled white noise and sky signal that give the best least-squares fit to the data.

4.4.2.5 Receiver channels

Because the circularly polarized component of diffuse Galactic emission is vanishingly small ($V \simeq 0$), the C-BASS architecture provides two independent measurements of Stokes I, $I_1 = |E_l|^2 = I - V$ and $I_2 = |E_r|^2 = I + V$ (hereafter abbreviated $I_{1/2}$). The C-BASS architecture also provides two measurements of Q and U, although Q_1 and Q_2 ($Q_{1/2}$) and U_1 and U_2 ($U_{1/2}$) are not independent of one another as their receiver chains only diverge at the final 90° hybrids.

In the pipeline we calibrate the averages of $I_{1/2}$, $Q_{1/2}$ and $U_{1/2}$ and do not treat them independently. By averaging $I_{1/2}$ together we reduce the noise on the measurement of I and, provided both are calibrated accurately, remove any circularly polarized emission. Because they are not independent, averaging $Q_{1/2}$ and $U_{1/2}$ does not reduce the noise on measurements of Q or U. We still use their average values in order to hedge against errors that might be present in one of the channels.

There may be systematic differences between $I_{1/2}$ caused by differences between the two receiver chains as well as calibration errors. Any differences between $Q_{1/2}$ and $U_{1/2}$ can only be caused by errors in the final 90° hybrids or calibration errors.

Because $Q_{1/2}$ and $U_{1/2}$ are not independent it is incorrect to sum the covariance matrices and therefore the noise model against which we have compared the Q and U jackknife tests to is not appropriate. We still present the same metrics but, when applied to Q and U data, $\chi^2_{\rm red}$ should equal zero. For example, the total intensity jackknife has $\chi^2_{\rm red} = 1.11$ and the polarization jackknife has $\chi^2_{\rm red} = 0.11$.

The $I_{1/2}$, $Q_{1/2}$ and $U_{1/2}$ difference maps and weighted difference maps are shown in Figure 4.27. The jackknife angular power spectra are shown in Figure 4.28 and the histograms of the weighted difference maps are shown in Figure 4.29. Figure 4.30 shows the gain factors that must be applied to the I_2, Q_2 and U_2 maps to minimise the $\chi^2_{\rm red}$ in the difference maps of the brightest 1% of pixels. The gain factors are 1.01, 0.91 and 0.91 in I, Q and U respectively.

The extra power at large angular scales in the I jackknife is consistent with the residual noise in the even/odd test and the $I_{1/2}$ gain calibration is consistent to 1%. The best fitting scaling applied to the white noise power is less than one, perhaps indicating that I_1 and I_2 are not independent of one another.

The Q and U difference maps and angular power spectra are highly correlated with the sky signal. The angular power spectra show that across the whole sky $Q_{1/2}$ and $U_{1/2}$ differ from each other by $\sim 2 - 3\%$ in power, which equates to $\sim$ 15% in amplitude. The gain factor that minimises the $\chi^2_{\rm red}$ between brightest 1% of

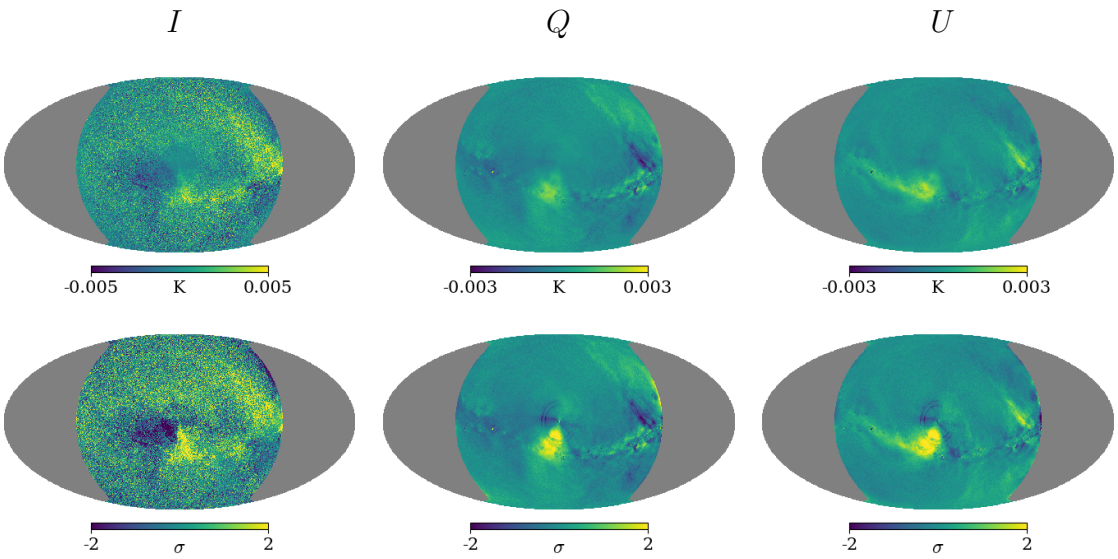

Figure 4.27: Receiver channel jackknife maps. From *top* to *bottom* are the difference maps (*top row*) and weighted difference maps (*bottom row*). From *left* to *right* the maps are of Stokes I, Q and U data.

pixels is 0.91. Given the limitations of both methods, these results are approximately consistent.

4.4.2.6 Summary of jackknife results

The results of the jackknife tests show that there is extra noise in the maps that is not modelled by DESCART. The $\chi^2_{\rm red}$ values of the even/odd and I_1/I_2 difference maps show that the covariance matrices output by DESCART underestimate the white noise variance by 10–15%, which corresponds to underestimating the white noise σ

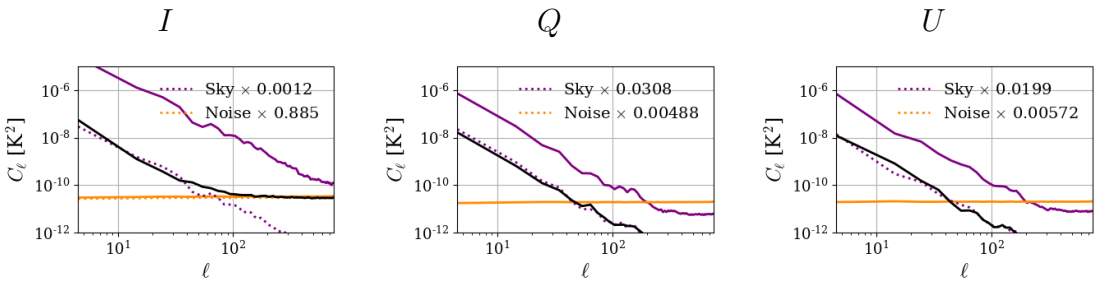

Figure 4.28: Receiver channel jackknife angular power spectra. From *left* to *right*; Stokes I, Q and U. The *solid black line* is the jackknife, the *solid orange line* is the expected white noise and the *solid purple line* is the sky signal. The *dashed lines* show the scaled white noise and sky signal that give the best least-squares fit to the data.

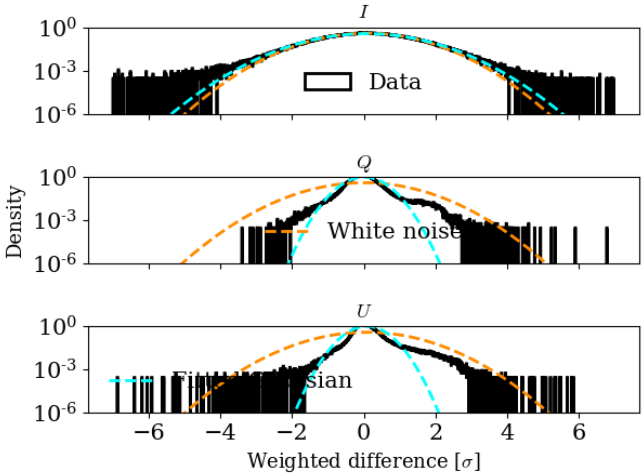

Figure 4.29: Receiver channel jackknife weighted difference map histograms. The *dashed orange* line is the expected white noise from the covariance matrix and the *dashed cyan* line is the best fit gaussian to the data.

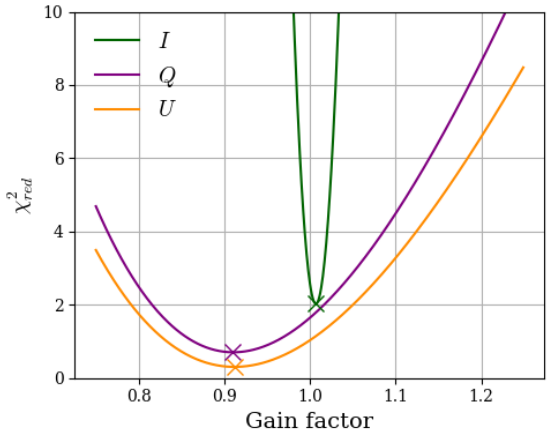

Figure 4.30: Receiver channel jackknife gain factors. The curves show the $\chi^2_{\rm red}$ of the difference between the brightest 1% of pixels as a function of the gain factor applied to the I_2 (*green*), Q_2 (*purple*) and U_2 (*orange*) maps. The *crosses* show the factors that minimise $\chi^2_{\rm red}$.

values by 5–7%. The mis-modelling is likely caused by the short timescale correlations in the TOD that are introduced by the finite width of the filters that decimate the digitised signals (Section 2.1.3). We therefore scale the pixel covariance matrices from DESCART by a factor of 1.12 before using them in scientific analysis.

The 1st half/2nd half jackknife shows that the telescope gain is currently inconsistently calibrated by 1–3%. We therefore add 2% of the I, Q and U values to their corresponding white noise maps.

Maps made from daytime and night-time data are discrepant from one another even when flagging data within a large angular distance of the sun. For this reason we only use night-time data for scientific analysis presented in this thesis.

On large angular scales ($\ell < 60$) the jackknife residuals are not consistent with white noise. The large scale excess power is however at least 1–2 orders of magnitude smaller than the sky power.

4.5 Northern survey maps

In this section we present maps made from the E37N dataset (night-time only, elevation 37° C-BASS North scans) that we have used for early science. Because the polarization angle is not currently well calibrated it cannot be used for scientific analysis. The polarized intensity P, on the other hand, is well calibrated and so we produce maps of P and its uncertainty.

To most easily analyse the astronomical emission in multiple maps at different frequencies they should all be in the same coordinate frame, smoothed to a common resolution and calibrated to the same temperature scale. In this section we describe how we produced C-BASS maps using DESCART, how we smoothed them to 1° resolution using a deconvolution, how we scaled them to brightness temperature and downgraded them to lower $N_{\rm side}$.

Section 4.5.1 contains a description of how we produced the I map and its white noise properties. Section 4.5.2 contains a description of how we made the P map and its noise properties.

4.5.1 Total intensity

In this section we present the C-BASS North I and σ_I maps. We describe the DESCART parameters, how we scaled the maps to brightness temperature, how we smoothed them using a deconvolution and how we downgraded them to lower $N_{\rm side}$.

Two correction factors need to be applied to the variance maps to correctly propagate the errors. One after smoothing, $\eta_{\rm smooth}$, and another after downgrading to lower $N_{\rm side}$, $\eta_{\rm dg}$.

We ran DESCART on the E37N dataset to produce a traditionally destriped $N_{\rm side} = 1024$ map using an offset length of $\lambda_b = 5\,{\rm s}$ and masking the brightest 5% of pixels during the baseline offset estimation (but including all of the pixels when binning to produce the final map). As a final extra cut on bad data we removed chunks of TOD that were further than 10-σ from the mapmaking solution at the $250^{\rm th}$ step of the PCG algorithm. We took the map solution after 1500 steps of the PCG algorithm. In this section we discuss the I map and the polarization maps are discussed in the following section.

The $N_{\rm side} = 1024$ I and σ_I^2 maps, that were produced by DESCART, are shown in Figure 4.31. The maps are in units of antenna temperature. Antenna temperatures are telescope specific and further scientific analysis is made easier by converting the maps to brightness temperature, measured in Kelvin. Because the telescope beam is non-Gaussian the conversion from Jy beam^{-1} to $\rm K_{RJ}$ is angular scale dependent.

Historically, antenna temperatures were scaled to so-called main-beam or full-beam scales, which are appropriate for analysing point sources or diffuse emission respectively (for examples see Reich (1982) or Jonas, Baart & Nicolson (1998)).

In previous analyses of C-BASS data, Irfan et al. (2015) scaled the C-BASS maps from antenna temperature to the full-beam scale with a multiplicative factor of 1.124 and used the scaled maps to analyse diffuse emission in the Galactic plane. Appendix A describes the conversion factors to these different temperature scales and the conditions under which they are appropriate.

In this work, by re-weighting the maps in spherical-harmonic space, we deconvolve the simulated C-BASS beam from the maps and smooth them to a Gaussian with $\theta_{\rm FWHM} = 1^\circ$. The deconvolution produces a map with a temperature scale that is not angular scale dependent and the process is described below.

The telescope beam filters the sky to remove small scale structure (at high ℓ). This filtering is encoded in the window function w_ℓ where, in the absence of noise, the observed angular power spectrum $C_\ell^{\rm obs}$ is related to the true sky angular power spectrum C_ℓ by

$$C_\ell^{\rm obs} = w_\ell C_\ell. \tag{4.14}$$

For an axially symmetric beam the window function is the square of the beam transfer function,

$$w_\ell = b_\ell^2. \tag{4.15}$$

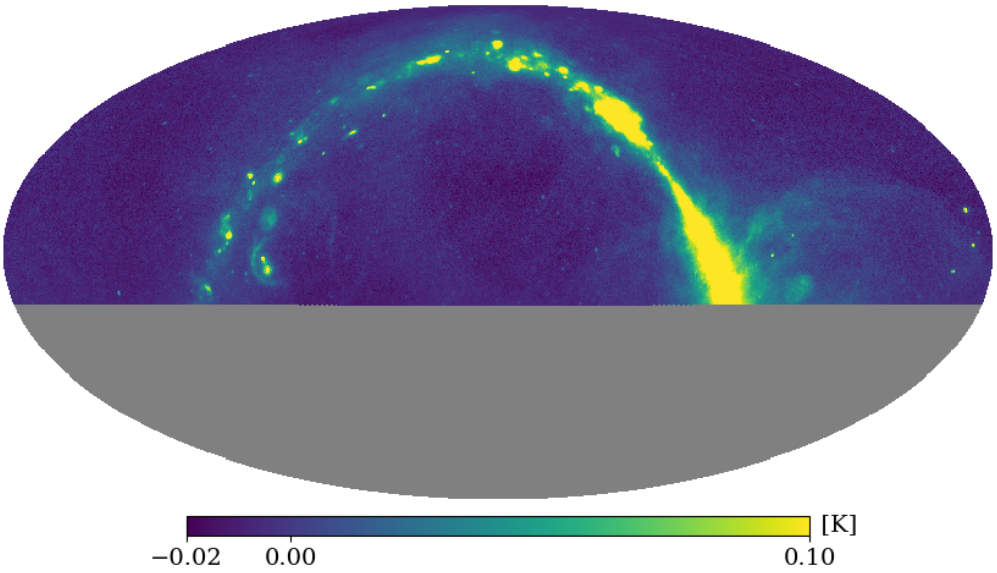

(a) I

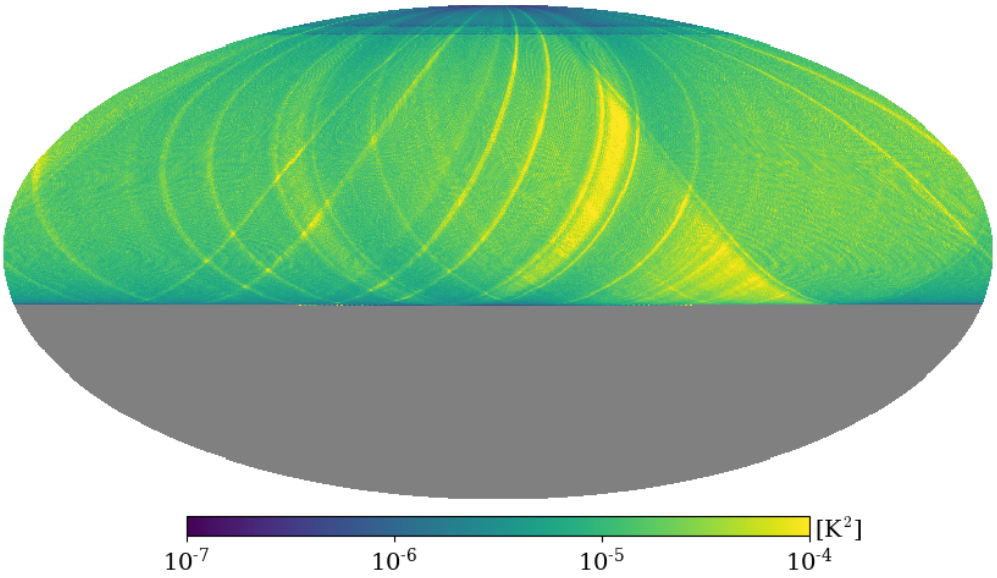

(b) σ_I^2

Figure 4.31: C-BASS I (*top*) and σ_I^2 (*bottom*) $N_{\rm side} = 1024$ maps in celestial coordinates. The I map is on a linear colour scale and the σ_I^2 map is on a logarithmic colour scale.

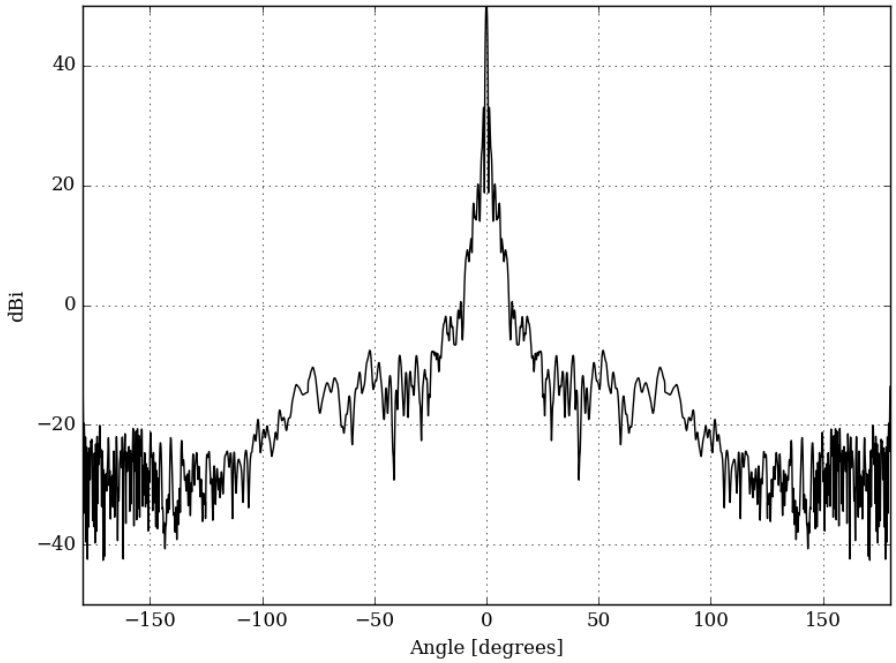

Figure 4.32: Simulated C-BASS North co-polar beam profile to 180°. Data from Holler et al. (2013).

The beam transfer function is the Legendre transform of the beam profile,

$$b_\ell = 2\pi \int \frac{b(\theta) P_\ell(\cos\theta)}{\Omega_{\rm beam}} d\cos\theta, \tag{4.16}$$

where $b(\theta)$ is the beam profile, P_ℓ are the Legendre polynomials and $\Omega_{\rm beam}$ is the beam solid angle (Page et al., 2003). The simulated C-BASS co-polar beam is plotted to 180° in Figure 4.32 (a zoom in of the first 10° was shown in Figure 2.2).

The window function of the C-BASS beam and a 1° Gaussian beam are shown in Figure 4.33. To smooth the C-BASS maps we multiplied their spherical harmonic transforms by $b_\ell^{\rm Gauss}/b_\ell^{\rm C\text{-}BASS}$ and inverse transformed back to the maps. The square of the ratio of the transfer functions is also shown in Figure 4.33.

We deconvolved the simulated C-BASS beam and smoothed the maps to a Gaussian beam with $\theta_{\rm FWHM} = 1°$ using this method. Before smoothing we scaled the maps from antenna temperature to brightness, measured in Jy beam^{-1}. The smoothed maps are then in units of Jy per 1° Gaussian beam, which we then scaled to brightness temperature. All of the conversion factors are set out in Appendix A.

We produced three sets of maps, each at a different effective frequency.[3] The effective frequencies were $\nu_{\rm eff} = 4.83$, 4.79 and 4.78 GHz, these are appropriately scaled

[3]Effective frequencies are defined in Appendix A.

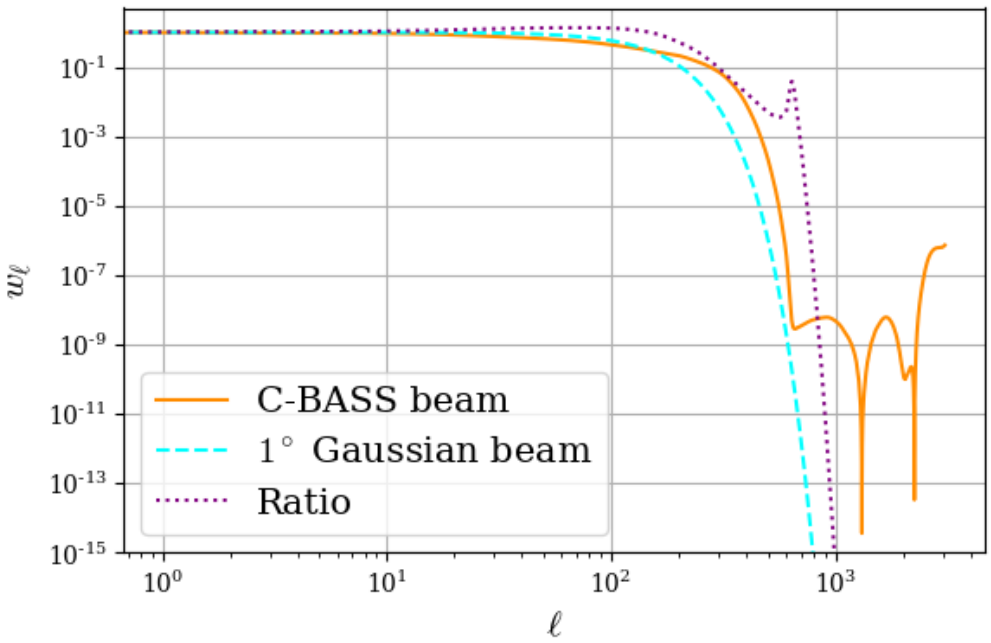

Figure 4.33: Window functions of the C-BASS beam (*solid orange line*), a 1° Gaussian beam (*dashed cyan line*) and the ratio between the Gaussian beam and the C-BASS beam window functions (*dotted purple line*).

for emission with spectral indices $\beta = 0.0$, 2.1 and 3.0 respectively (flat spectrum, typical free-free emission and typical synchrotron emission). In this section we only show the maps at $\nu_{\rm eff} = 4.83\,\mathrm{GHz}$ ($\beta = 0.0$).

Pixels near the lower declination bound are rendered useless due to edge effects in the deconvolution and so after smoothing we removed pixels within 3° of the edge of the observed region.

To correctly propagate the errors when convolving a variance map by a smoothing kernel requires a scaling factor $\eta_{\rm smooth}$. For the simple case of convolving a map with a 2D Gaussian beam the scaling factor is

$$\eta_{\rm smooth} = \frac{\Omega_{\rm pixel}}{2\Omega_{\rm Gaussian}}, \tag{4.17}$$

where $\Omega_{\rm pixel}$ is the solid angle of a map pixel and $\Omega_{\rm Gausian}$ is the solid angle of the Gaussian.

In the case of a deconvolution the scaling is more complex and we estimate it using simulations. To simulate the noise we multiplied each pixel in the σ_I map by a random number drawn from the standard normal distribution. The normalised deviation of the noise from zero therefore obeys the standard normal distribution and

its variance is unity.[4]

We deconvolved the C-BASS beam from the simulated noise and variance maps and smoothed them to a 1° Gaussian. The normalised deviation of the smoothed maps is only consistent with the standard normal distribution after the smoothed variance has been scaled by $\eta_{\rm smooth} = 0.0047$ and we therefore multiply the smoothed C-BASS σ_I^2 map by this factor.

This scaling factor does not depend on the level of noise or its spatial structure in the unsmoothed maps and we verified this with simulations. If smoothing to a different Gaussian beam then $\eta_{\rm smooth}$ would be different.

After smoothing we rotated the I and σ_I^2 maps from celestial to Galactic coordinates.

The smoothed $N_{\rm side} = 1024$ maps over sample the 1° Gaussian beam and we therefore downgrade the maps to lower resolutions for subsequent analyses. For maps with $\lesssim 2$ pixels per beam an appropriate map resolution is $N_{\rm side} = 64$, which has a pixel size of 55 arcmin. When downgrading the resolution of a HEALPix map the lower resolution superpixel fully contains $(N_{\rm side}^{\rm children}/N_{\rm side}^{\rm super})^2$ children pixels. Using the HEALPix UD_GRADE routine the superpixels are set to the mean of the children pixels.

A correction factor $\eta_{\rm dg}$ needs to be applied to the downgraded variance map. For the case of uncorrelated white noise,

$$\eta_{\rm dg} = \frac{N_{\rm side}^{\rm super}}{N_{\rm side}^{\rm child}}. \tag{4.18}$$

Because the C-BASS maps have been smoothed the noise is no longer un-correlated and the correction factor in this case is more complicated than the one in Equation 4.18. Instead of trying to calculate $\eta_{\rm dg}$ analytically we estimated it from simulations. We generated $N_{\rm side} = 1024$ maps of white noise from the C-BASS σ_I^2 map, deconvolved the C-BASS beam from the variance and simulated noise maps and smoothed them to 1°, and applied the $\eta_{\rm smooth}$ correction to the variance map. After downgrading them to $N_{\rm side} = 64$, the normalised deviation of the noise was only consistent with the standard normal distribution when the variance map was scaled by a factor of $\eta_{\rm dg} = 0.54$. This factor was insensitive to scalings in the noise level or the spatial structure of the un-smoothed variance map. We applied this scaling to the real deconvolved/smoothed C-BASS σ_I^2 map.

We then applied the corrections from the jackknives. The variance map was increased by 12% and the square of 2% of the I map was added.

[4]The normalised deviation of a vector $\mathbf{x}$ from $\hat{x}$ is $(\mathbf{x} - \hat{x})/\sigma_{\mathbf{x}}$.

The deconvolved, smoothed, rotated, downgraded and jackknife-corrected I and σ_I maps in units of brightness temperature are shown in Figure 4.34.

4.5.2 Polarization

In this section we present the C-BASS Q, U, σ_Q and σ_U maps. Produced the same way as the I and σ_I map. We then also produce $P = \sqrt{Q^2 + U^2}$ and σ_P maps.

The polarization maps were produced by DESCART and the destriping parameters were described at the start of the previous section. We treated the Q, U, σ_Q and σ_U maps the same as their total intensity counterparts but did not rotate them from celestial to Galactic coordinates.

The deconvolution and UD_GRADE routine do not account for the parallel transport of the polarization vectors over the surface of the celestial sphere. Over most of the sky neglecting the parallel transportation during smoothing or downgrading introduces a very small error into the Q and U maps. The errors are only considerable at the coordinate poles (Vidal Navarro, 2014). In this work we neglect the parallel transport of the polarization vector.

The $N_{\rm side} = 1024$ Q and U maps, in celestial coordinates and units of antenna temperature, along with their white noise levels are shown in Figures 4.35 and 4.36. The QU covariance map is shown in Figure 4.37. The band of low covariance is at the declination where elevation 37° scans cut through pixels at 90°. The QU covariance is much smaller than the Q and U variances and so we approximate Q and U as independent when estimating P below.

The polarization angle calibration is not currently applied and so there is a systematic error in the parallactic angle offset. Furthermore, QU mixing introduces errors in the polarization angles of a few degrees. The polarization angle in the C-BASS maps is too poorly defined to be useful until it is correctly calibrated but the polarized linear intensity P, on the other hand, is well constrained. P is defined as,

$$P^2 = Q^2 + U^2. \tag{4.19}$$

Unlike the I, Q and U maps, P maps are not subject to Gaussian errors. In a given pixel, the measured Q value is drawn from a Gaussian PDF,

$$p(Q|Q_0, \sigma_Q) = \frac{1}{\sqrt{2\pi\sigma_Q^2}} e^{-\frac{(Q-Q_0)^2}{2\sigma_Q^2}}, \tag{4.20}$$

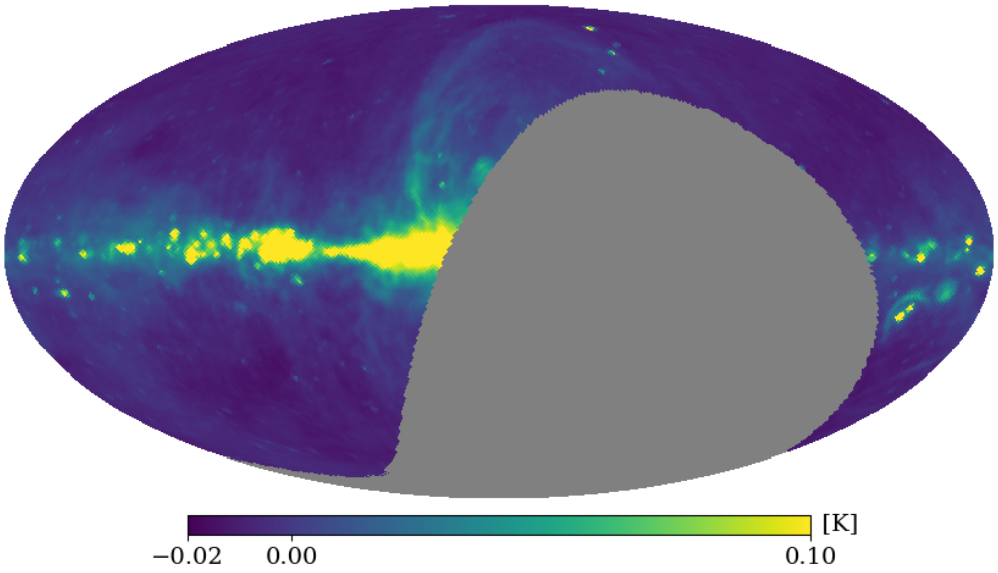

(a) I

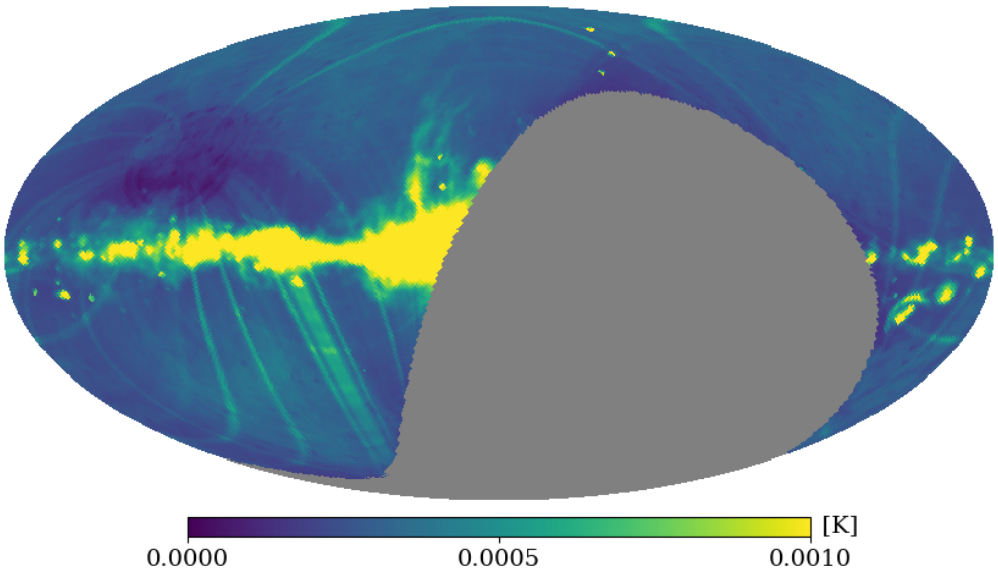

(b) σ_I

Figure 4.34: C-BASS I (*top*) and σ_I (*bottom*) $N_{\rm side} = 64$ maps in Galactic coordinates, smoothed to a Gaussian beam with 1° FWHM. The units are brightness temperature in Kelvin. Both maps are on linear colour scales and the jackknife corrections have been applied to σ_I.

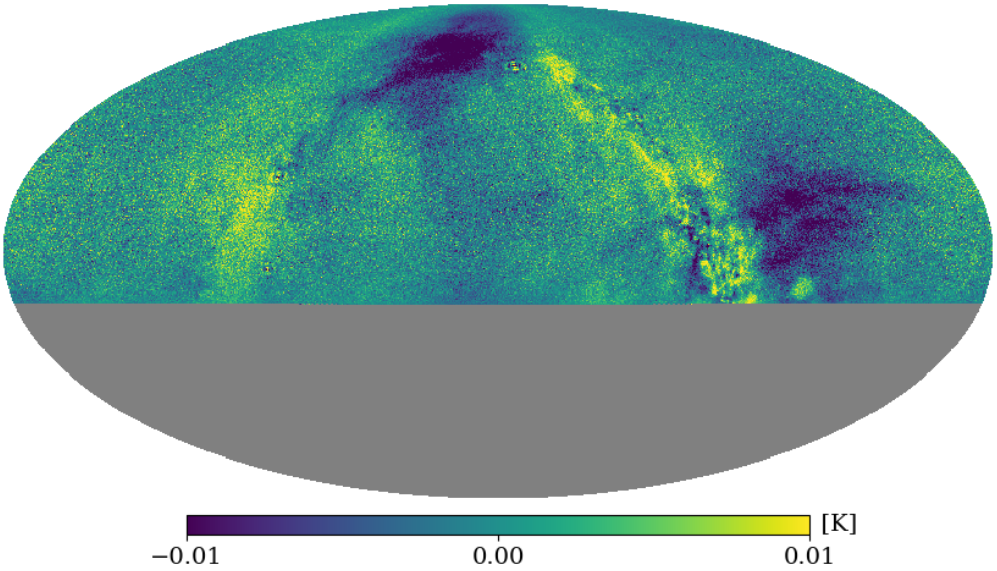

(a) Q

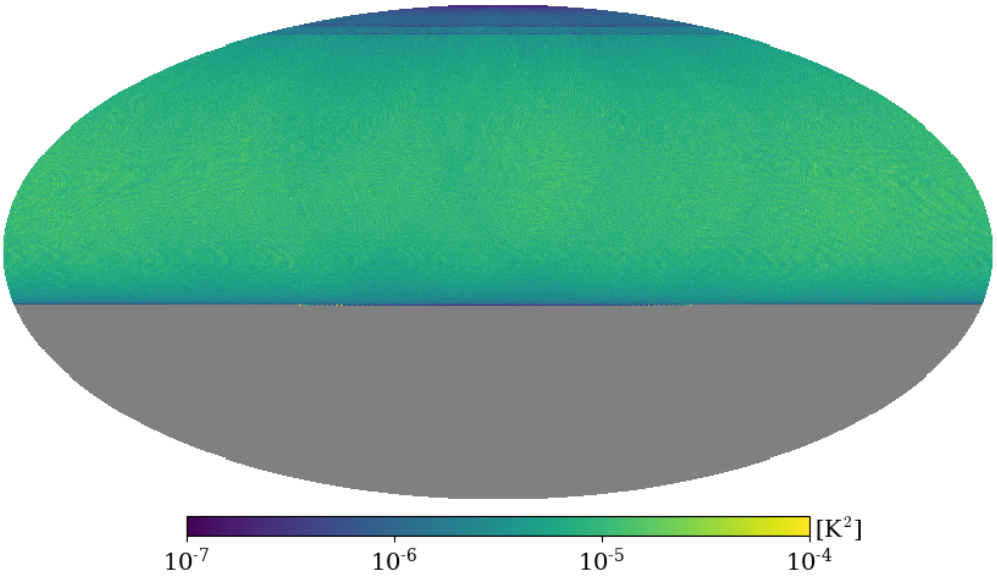

(b) σ_Q^2

Figure 4.35: C-BASS Q (*top*) and σ_Q^2 (*bottom*) $N_{\rm side} = 1024$ maps in celestial coordinates. The Q map is on a linear colour scale and the σ_Q^2 map is on a logarithmic colour scale.

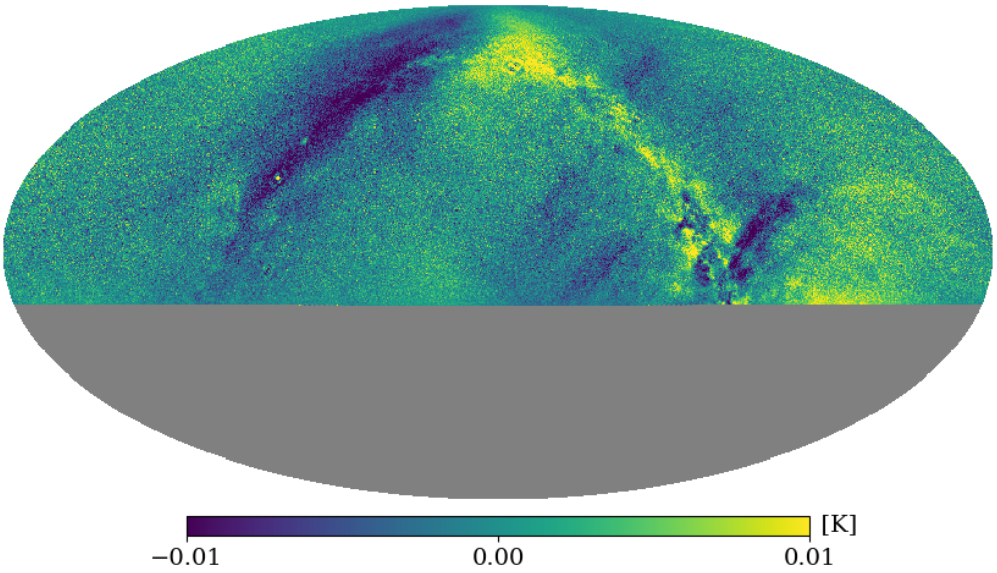

(a) U

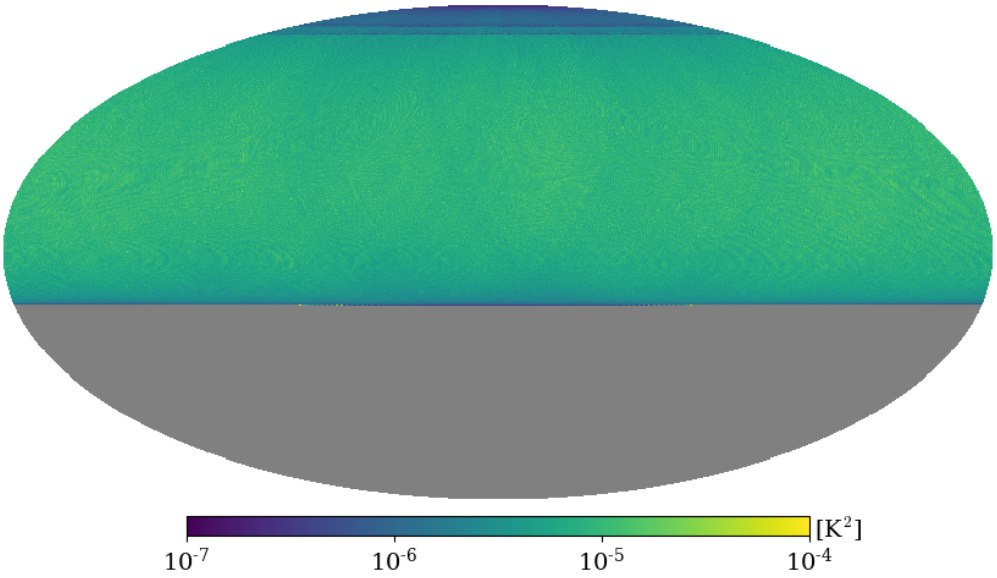

(b) σ_U^2

Figure 4.36: C-BASS U (*top*) and σ_U^2 (*bottom*) $N_{\rm side} = 1024$ maps in celestial coordinates. The U map is on a linear colour scale and the σ_U^2 map is on a logarithmic colour scale.

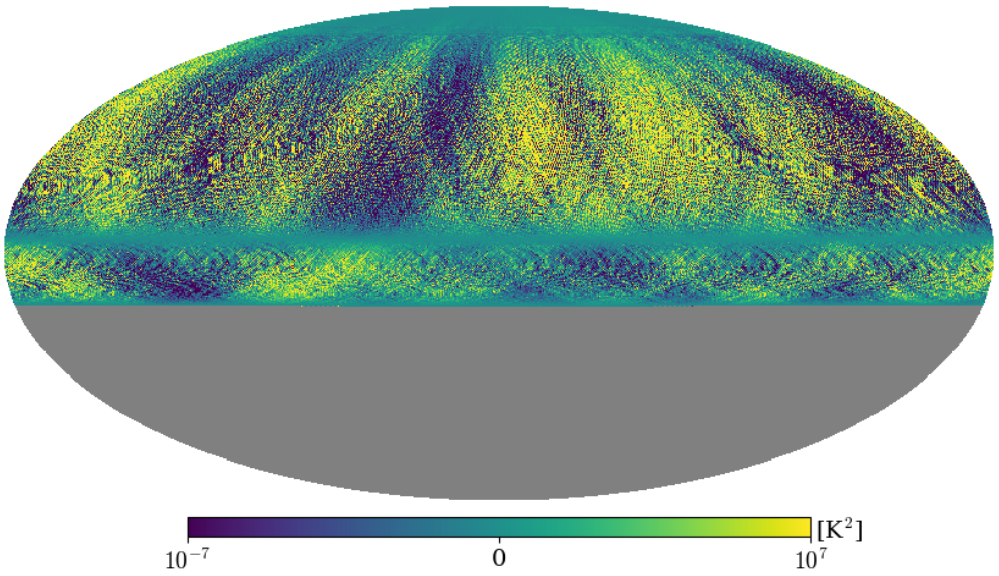

Figure 4.37: C-BASS Q and U covariance map. The map is on a linear colour scale, in celestial coordinates and has $N_{\rm side} = 1024$.

where Q_0 is the true underlying value and σ_Q is the pixel white noise level. Similarly for U,

$$p(U|U_0, \sigma_U) = \frac{1}{\sqrt{2\pi\sigma_U^2}} e^{-\frac{(U-U_0)^2}{2\sigma_U^2}}. \tag{4.21}$$

Because P is formed from the quadrature sum of two independent Gaussian random variables, measurements of P are not themselves drawn from a Gaussian distribution or even a symmetrical distribution. Making the simplifying assumption that $\sigma_Q = \sigma_U \, (\equiv \sigma_P)$ then measurements of P are drawn from a Rician distribution,

$$p(P|P_0, \sigma_P) = \frac{P}{\sigma_P^2} e^{-\frac{P^2+P_0^2}{2\sigma_P^2}} I_0\left(\frac{PP_0}{\sigma_P^2}\right), \tag{4.22}$$

where P_0 is the true linear polarization and I_0 is the zeroth order modified Bessel function of the first kind (not the true stokes I value in the pixel!).

A scatter plot of σ_U against σ_Q is plotted in Figure 4.38. It shows that across much of the C-BASS map $\sigma_U \simeq \sigma_Q$ and therefore we can approximate the distribution from which P values are drawn as Rician. We estimate σ_P from the average σ_Q and σ_U in each pixel. In the case of large P_0 and small σ_P then $p(P|P_0, \sigma_P)$ approaches a Gaussian distribution, but for weakly polarized signals with significant noise this is not the case.

We produced P from the smoothed and downgraded Q and U maps and estimated

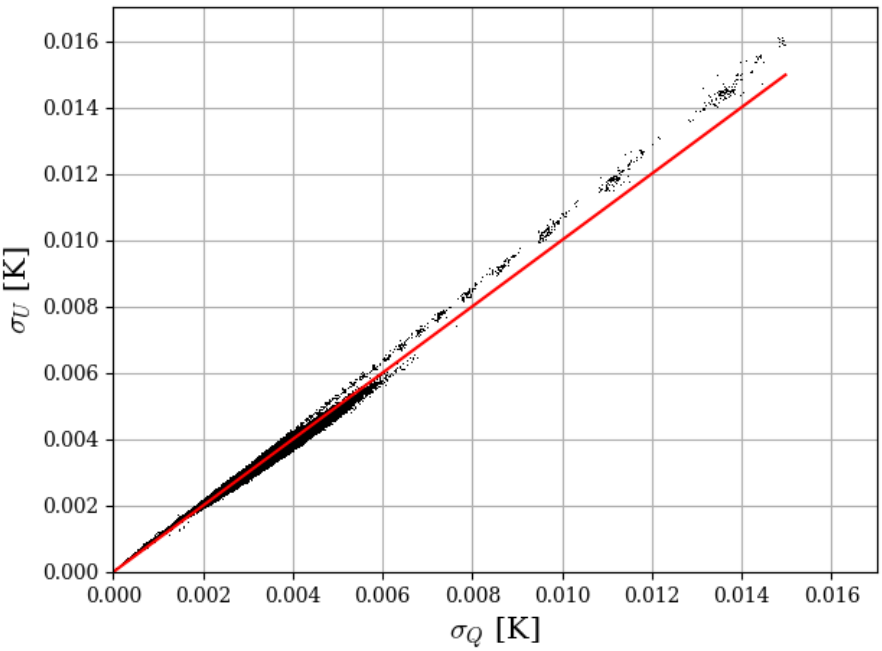

Figure 4.38: A scatter plot of σ_Q against σ_U for the $N_{\rm side} = 1024$ C-BASS map. The line $\sigma_U = \sigma_Q$ is shown in *red*.

σ_P from the average of the smoothed and downgraded σ_Q and σ_U maps.[5] We then rotated the P and σ_P maps from celestial to Galactic coordinates and applied the jackknife corrections (increased σ_P^2 by 12% and added 2% of the P map in quadrature to σ_P). The P and σ_P maps are shown in Figure 4.39. We created maps for the three effective frequencies (appropriate for emission with spectral indices $\beta = 0.0$, 2.1 and 3.0) and here we only show the map for $\nu_{\rm eff} = 4.83\,{\rm GHz}$ ($\beta = 0.0$).

[5]The averaging operations (deconvolution/smoothing and downgrading to lower $N_{\rm side}$) should be performed on the Q and U maps so that the white noise averages down.

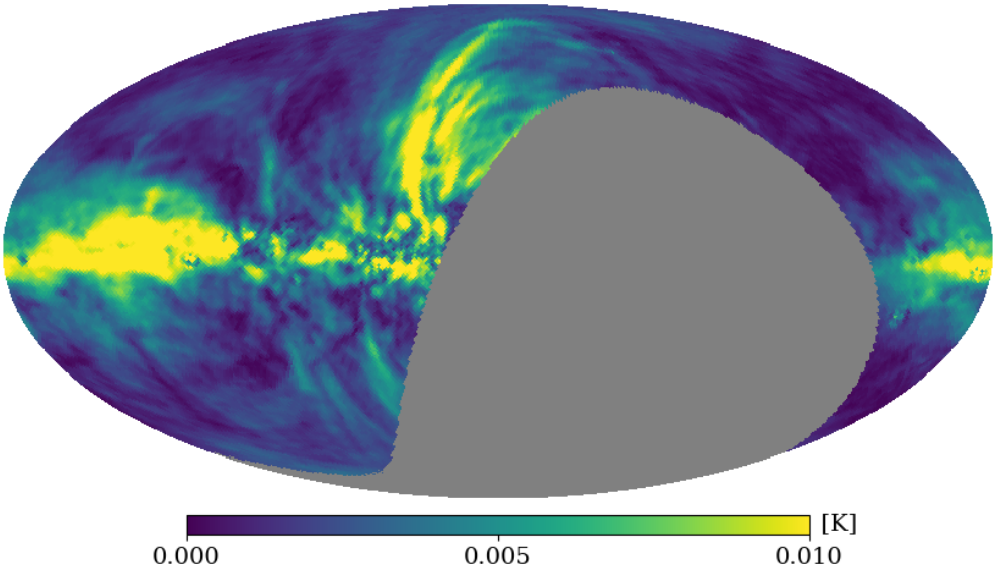

(a) P

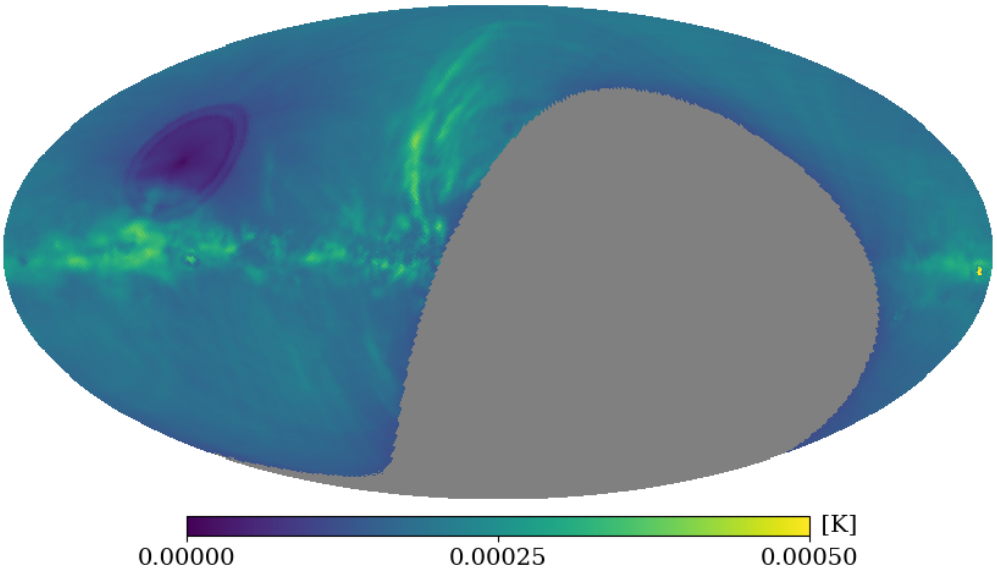

(b) σ_P

Figure 4.39: C-BASS P (*top*) and σ_P (*bottom*) $N_{\rm side} = 64$ maps in Galactic coordinates and smoothed to a Gaussian beam with 1° FWHM. The units are brightness temperature in Kelvin. Both maps are on linear colour scales.

4.6 Summary of preliminary C-BASS maps

In this chapter we have used simulations to test the DESCART mapmaker and now understand its current limitations.

We tested maps made from C-BASS North data using jackknife tests. The results from the jackknives should be considered before beginning any analysis using the C-BASS maps. The σ_I^2 and σ_P^2 maps should be increased by 12% to account for the short time-scale correlations in the TOD that DESCART does not account for. The inconsistent gain calibration should be accounted for by adding 2% of the I and P signal into the corresponding σ map. On scales larger than $\sim 3^\circ$ the jackknife residuals are not consistent with white noise, but these correlated residuals have power at least an order of magnitude smaller than the sky signal in P and two orders of magnitude smaller in I. There is residual ground in the P maps along constant declination bands at the single σ level in pixel space.

We produced I, P, σ_I and σ_P maps from the E37N dataset and smoothed them to 1° resolution.

Figures 4.40–4.46 show a series of zoomed in gnomonic projections of interesting features in the C-BASS I and P maps that are discussed further below.

Figure 4.40 shows a region centred on the supernova remnant TauA (Crab Nebula). TauA is a $\sim 5-10\%$ polarized source and the cross polar beam is evident in the P map. Surrounding the radio sources Orion A and B is Barnard's Loop, which is an HII region dominated by unpolarized free-free emission at 5 GHz. In between Barnard's Loop and TauA is λ Orionis, a patch approximately 8° across. Its emission is dominated by free-free radiation. At higher frequencies (20–30 GHz) it is surrounded by a ring of AME but this is negligible at 5 GHz (Planck Collaboration et al., 2015d).

Figure 4.41 is centred on the NPS, which is likely the expanding shell of a supernova remnant that is emitting synchrotron radiation (Vidal et al., 2015; Liu, Mertsch & Sarkar, 2014). It is a highly polarized spur with a bright counterpart in total intensity and has been observed across a wide range of frequencies. There are many loops and spurs across the radio sky that have been extensively studied, for examples see Quigley & Haslam (1965); Berkhuijsen, Haslam & Salter (1971); Carretti et al. (2013); Planck Collaboration et al. (2015d).

The Galactic plane is across the bottom of this region. The P map has canals of depolarization. The depolarization is caused by differential Faraday rotation, which occurs when there are both relativistic and non-relativistic populations of electrons

in the same magnetized medium (Burn, 1966; Sokoloff et al., 1998). In a uniform slab the polarization fraction f is given by

$$f = f_0 \mathrm{sinc}(|2R\lambda^2|), \tag{4.23}$$

where f_0 is the intrinsic polarization (for synchrotron up to 75%), R is the observed rotation measure and λ is the wavelength. Clearly when $|2R\lambda^2| = n\pi$ $(n = 1, 2, 3...)$ the polarization vanishes. The depolarization canals are infinitely thin structures convolved with the telescope beam. They have no observable counterparts in the I maps and because of this they are sometimes referred to as Faraday Ghosts (Shukurov & Berkhuijsen, 2003).

Figure 4.42 is centred on the quasar 3C273, to the bottom right of 3C273 is another quasar 3C279. Left of both sources is the unpolarized, supergiant eliptical galaxy VirgoA. When TauA was below the horizon, the radio sources Virgo A and 3C273 were used to calibrate the C-BASS North TOD. The highly polarized NPS enters the region at the bottom left.

Figure 4.43 is centred on the radio galaxy Cygnus A. Cygnus A is highly polarized and the cross polar beam is clearly visible in the P map. The Cygnus Loop supernova remnant is below Cygnus A. The Galactic plane lies across the middle of the image, and Faraday Ghosts are present at the right hand side (closest to the Galactic centre) and in the Cygnus complex. Spreading out from the Galactic plane are loops and spurs of polarized synchrotron radiation.

Figure 4.44 is centred on the NCP. This is the region of the C-BASS map with the deepest observations. There are no obvious scan synchronous features in the I map. 20° to the left and 10° above the NCP is a 15° wide patch of low emission in I. This patch is also in the Haslam map and is therefore a real feature and not an artefact. In P there is an obvious disc of pixels with very low noise levels. Loop II enters the region on the right hand side and is particularly clear in polarization. The NCP is the pole of the coordinate system and therefore P is not reliable within a few degrees of it due to us treating Q and U as scalar fields during the smoothing and downgrading.

Figure 4.45 is centred on the Pegasus region. This zoom in clearly shows lots of extra-galactic radio point sources in I that are not present in P. This region also contains polarized synchrotron spurs.

Figure 4.46 is centred on the Perseus Molecular Cloud, G159.6-18.5. It is one of the most well studied AME regions in the sky (Génova-Santos et al., 2015a). TauA is identified at the top left of the image. NGC 1499 (The California Nebula) depolarizes

the emission at 5 GHz and Planck Collaboration et al. (2013b) find AME here with a peak frequency of 50 ± 17 GHz, which is higher than typical. The Galactic plane is across the top of the image, there are no Faraday Ghosts in the P map because it is further from the Galactic centre.

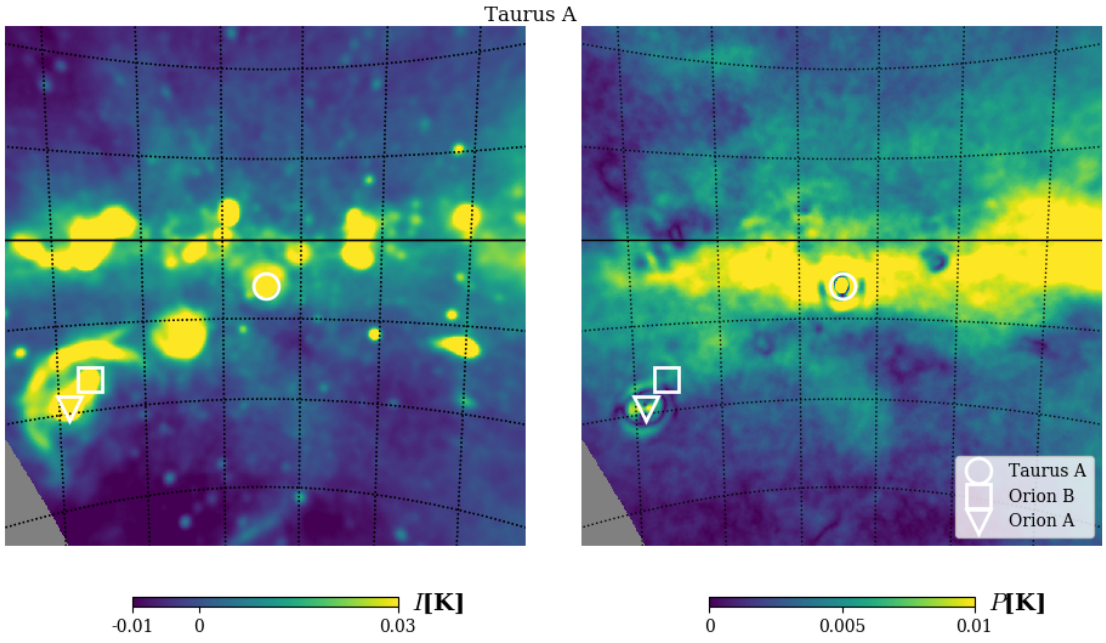

Figure 4.40: A gnomonic projection of the C-BASS I and P maps, centred on TauA. The images are $67° \times 67°$ and on linear colour scales. The graticules are lines of constant Galactic latitude and longitude in 10° intervals.

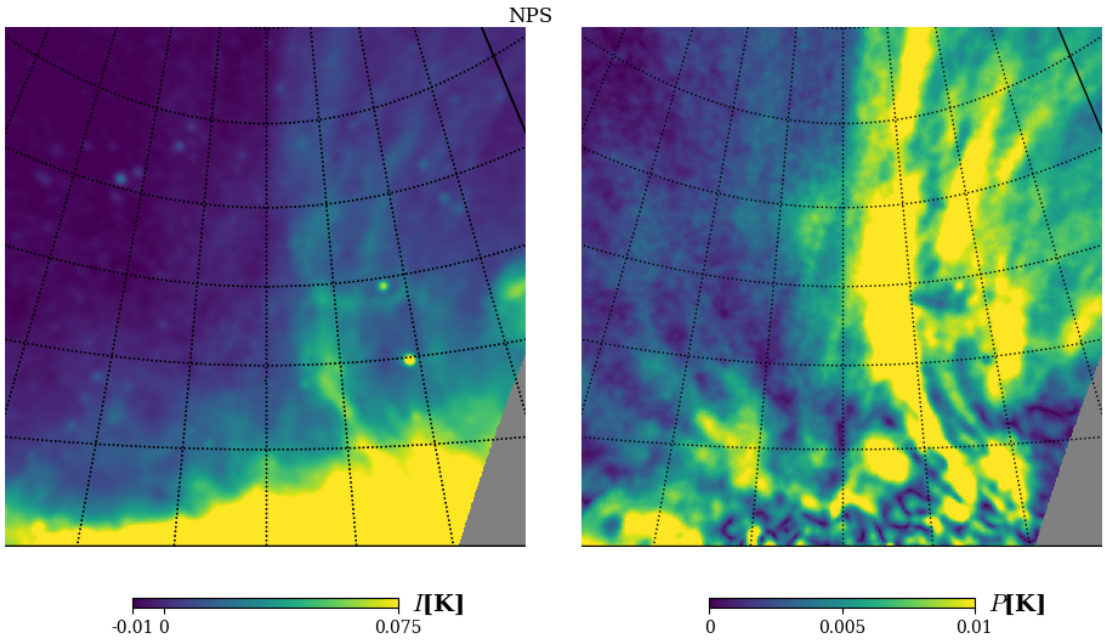

Figure 4.41: A gnomonic projection of the C-BASS I and P maps, centred on the NPS. The images are 67° × 67° and on linear colour scales. The graticules are lines of constant Galactic latitude and longitude in 10° intervals.

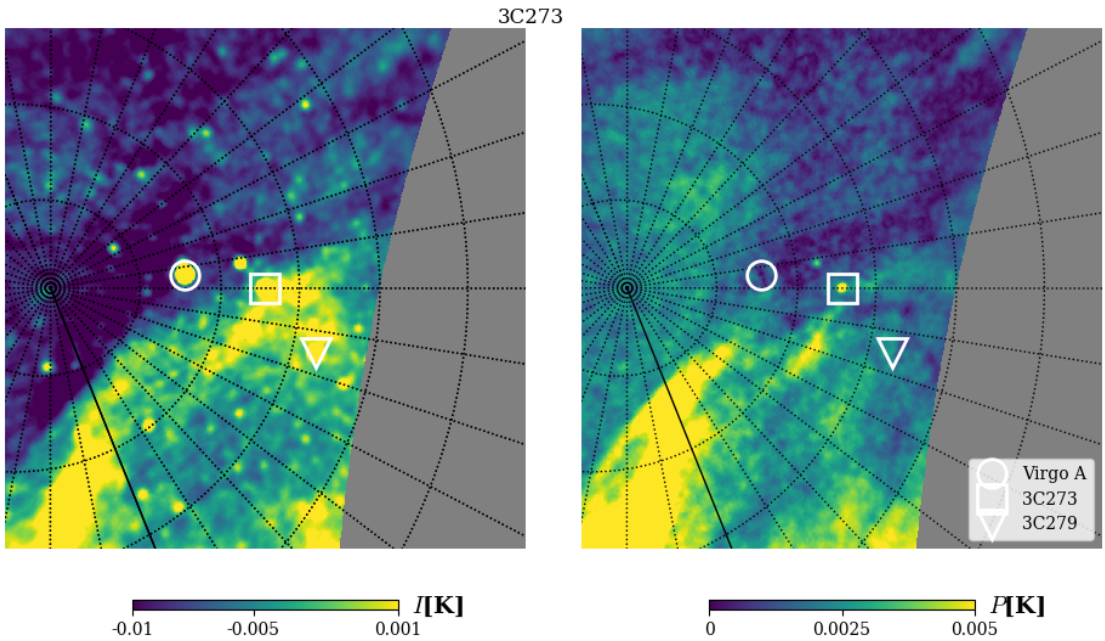

Figure 4.42: A gnomonic projection of the C-BASS I and P maps, centred on 3C273. The images are 67° × 67° and on linear colour scales. The graticules are lines of constant Galactic latitude and longitude in 10° intervals.

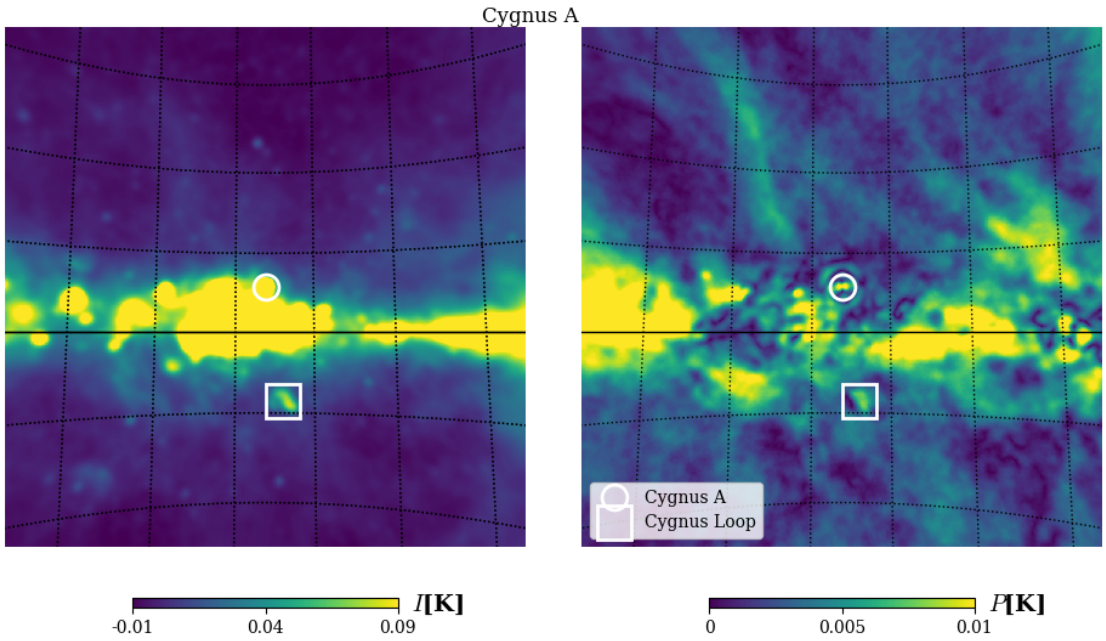

Figure 4.43: A gnomonic projection of the C-BASS I and P maps, centred on CygA. The images are 67° × 67° and on linear colour scales. The graticules are lines of constant Galactic latitude and longitude in 10° intervals.

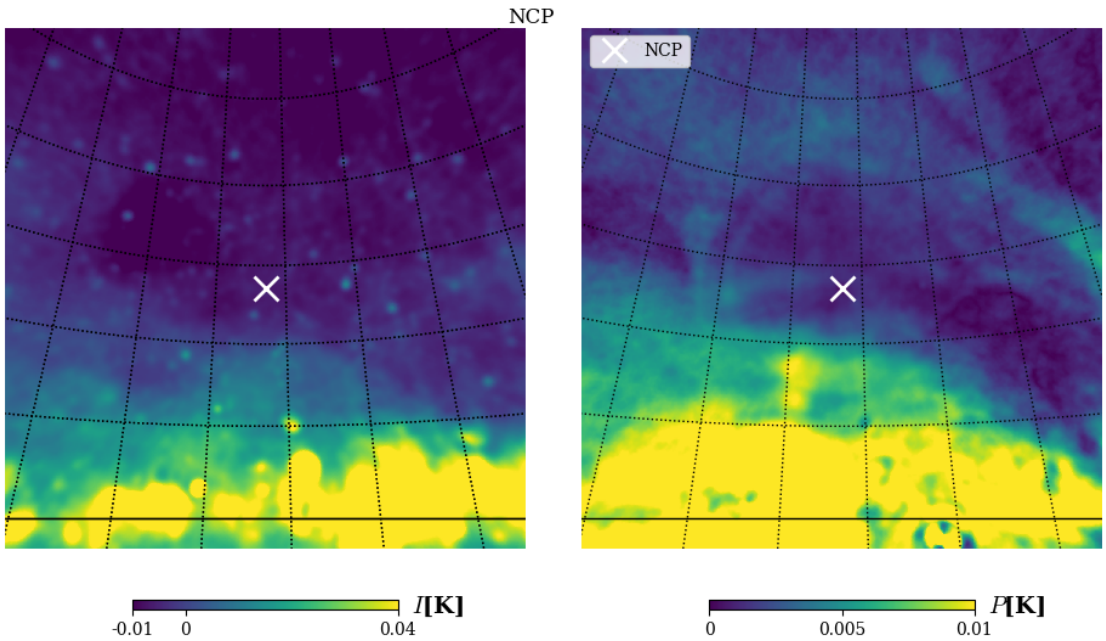

Figure 4.44: A gnomonic projection of the C-BASS I and P maps, centred on the NCP. The images are 67° × 67° and on linear colour scales. The graticules are lines of constant Galactic latitude and longitude in 10° intervals.

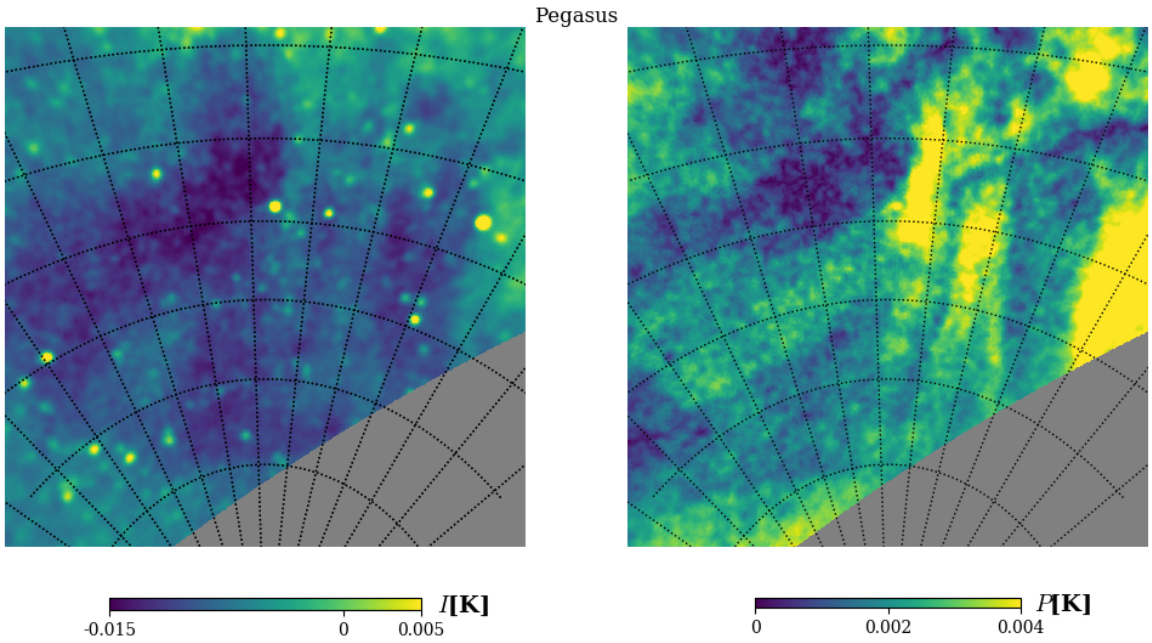

Figure 4.45: A gnomonic projection of the C-BASS I and P maps, centred on The Pegasus Cluster. The images are $67° \times 67°$ and on linear colour scales. The graticules are lines of constant Galactic latitude and longitude in $10°$ intervals.

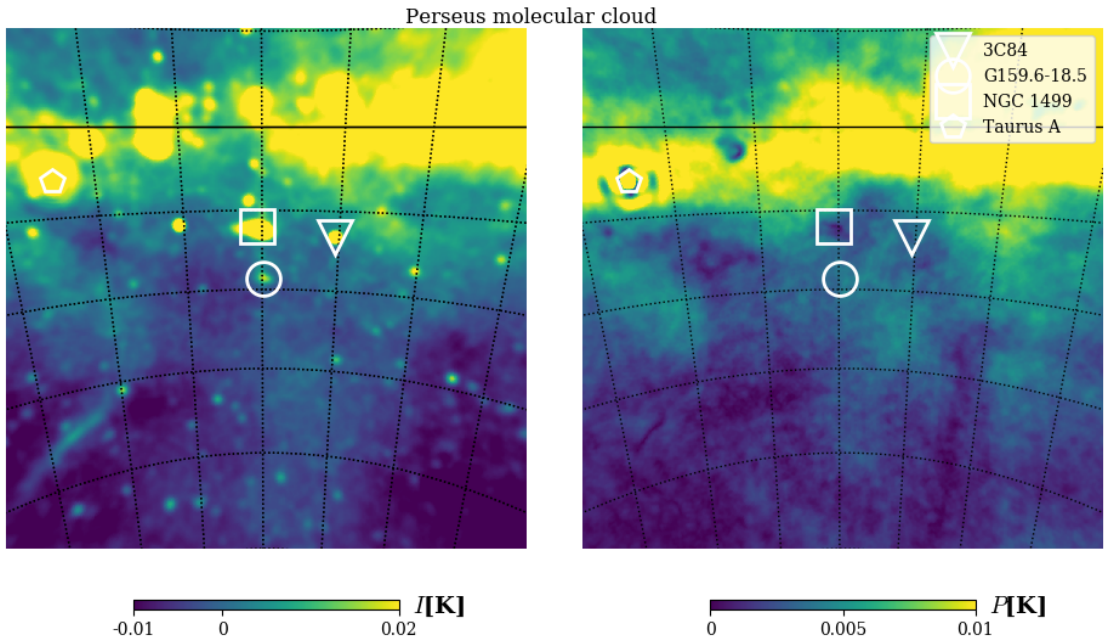

Figure 4.46: A gnomonic projection of the C-BASS I and P maps, centred on Perseus OB2. The images are $67° \times 67°$ and on linear colour scales. The graticules are lines of constant Galactic latitude and longitude in $10°$ intervals.

Chapter 5

Constraints on the Total Intensity Emission of Diffuse Galactic Foregrounds

In this chapter we use the C-BASS I map along with other all-sky surveys to demonstrate the complexity of low-frequency foregrounds to the CMB and to constrain the frequency spectrum of diffuse Galactic synchrotron radiation.

In Section 5.1 we make a 3-colour map from the C-BASS, Haslam and *WMAP* total intensity data. The colour of each pixel in the 3-colour map is determined by its frequency spectrum. The structure in this image demonstrates the complexity of low-frequency foregrounds.

In Section 5.2 we describe the method of T-T plots, which can be used to measure the spectral index between two maps at different frequencies. We extend the usual T-T plot method using a mixture model and test it on simulated data. The simulations show that the microwave sky is too complicated for this method to be used to estimate the spectral index systematically and blindly across the whole sky. We measure the spectral index between the C-BASS and Haslam maps using the extended T-T plot method and compare the results to simulations. In Barnard's Loop we also estimate the spectral index between the C-BASS map and the *WMAP* K-band and *Planck* 30 GHz maps.

5.1 A 3-colour map from Haslam, C-BASS and *WMAP* data

To qualitatively demonstrate the complexity of low-frequency foregrounds we produced a "3-colour" map. The *red*, *green* and *blue* channels of the 3-colour map are

Table 5.1: Summary of the $N_{\rm side} = 256$ I maps that we used to produce the 3-colour map.

Map	Frequency [GHz]	Typical $\sigma_I(1^\circ)$	Monopole	Unit	Reference
C-BASS	4.83	1.0	-15.0	$\mathrm{mK_{RJ}}$	This work
Haslam	0.408	1.1	8.9*	$\mathrm{K_{RJ}}$	Remazeilles et al. (2015)
WMAP K-band	23	6	-8.1*	$\mu\mathrm{K_{CMB}}$	Bennett et al. (2013)
WMAP V-band	61	6	1.0*	$\mu\mathrm{K_{CMB}}$	–"–

* Planck Collaboration et al. (2016b).

the Haslam, C-BASS and the difference between the *WMAP* K and V-band intensity maps respectively (subtracting the V-band map from the K-band map approximately removes the contribution from the CMB). The input maps are summarised in Table 5.1, they were smoothed to 1° resolution and downgraded to $N_{\rm side} = 256$. We did not apply any colour corrections to the input maps or adjust their effective frequencies to account for the frequency spectrum of the emission across each passband.[1]

The colour of each pixel in the 3-colour map is determined by its frequency spectrum. For the frequency spectrum to make sense, the monopoles of the maps have to be accurately set. Errors in the global offset subtraction will have the greatest effect in low-signal pixels. We subtracted the best-fit monopoles in Planck Collaboration et al. (2016b) from the Haslam and *WMAP* maps. We estimated the C-BASS monopole by assuming that there are some pixels that contain no signal from the Galaxy or from extra-galactic radio sources and they only contain realisations of the thermal and confusion noise. The temperatures of these pixels should be normally distributed about zero. The accuracy of this assumption needs to be tested in detail but is sufficient for the purposes of this qualitative demonstration.

The 3-colour map is shown in Figure 5.2. Pixels dominated by different foreground components appear as different colours. This rest of this section contains a description of how we scaled each input map, the spectra that give rise to different coloured pixels and a discussion of some of the interesting features in the 3-colour map.

We divided each input temperature map by $\nu^{-2.7}$. This scales the *red*, *green* and *blue* channels so that a pixel dominated by emission with a temperature spectral index $\beta = 2.7$ appears colourless. Any net colour in a pixel is caused by at least one of the maps having a temperature that is not consistent with emission obeying a power-law frequency spectrum with $\beta = 2.7$. The colour of a pixel therefore indicates the dominant emission mechanism within it. Figure 5.1 shows how the temperature in a pixel dominated by one of several typical frequency spectra is translated to the colour

[1] Appendix A contains a description of colour corrections and effective frequencies.

channels. The three main components of the diffuse Galactic emission at frequencies between 0.408 GHz and 23 GHz are synchrotron, free-free and AME.

Synchrotron emission typically has $\beta > 2.7$ and therefore the scaled *green* and *blue* channels will be low relative to the *red* channel. These pixels appear more orange and red as their spectra steepen.

AME has excess emission in the *blue* channel relative to the *red* and *green* channels, which makes pixels in AME dominated regions more blue.

Free-free emitting regions have $\beta < 2.7$ which causes a slight excess signal in the *green* and large excess in the *blue* channels relative to the *red* channel. Such pixels therefore appear cyan.

We applied an arctan scaling to the colour channels to give a broad dynamic range of brightness in the final image. This 3-colour map is not intended to be used to quantitatively analyse the foreground emission but demonstrates the complexity of diffuse low frequency foregrounds in total intensity. Gnomonic projections of some interesting areas of the 3-colour map are shown in Figure 5.3. Gnomonic projections of these regions in the C-BASS map were shown at the end of Section 4.6.

The *top left* projection is centred around the NPS. It is part of a radio loop and is dominated by synchrotron radiation and therefore appears *orange* in this image. On the far right hand side of this imageis the edge of the ζ Ophiuchi Nebula, also known as SH2-27. It is dominated by free-free emission and there have also been significant detections of AME. This complex appears *cyan* and *blue* in the 3-colour image, confirming those detections.

The *top right* projection is centred on the cloud G159.6-18.5 in Perseus. This AME emitting cloud appears *blue* in the 3-colour image. Above G159.6-18.5 is the California Nebula (NGC 1499), which is an ionised HII region and is dominated by free-free emission.

The *middle left* projection is centred on TauA, G159.6-18.5 is on the far right hand side of this image. To the left and below TauA are λ Orionis and Barnard's Loop, both of which are dominated by free-free emission and appear *cyan*. Around λ Orionis is a *blue* ring of AME.

The *middle right* projection is centred on the NCP. The subtle *blue/purple* cloud that passes vertically through the centre of the image is the Polaris Flare, which was the location of the first detected AME (Section 1.2.4).

The *bottom left* projection is centred on 3C273. Both it and 3C279 appear *cyan* suggesting frequency spectra close to that of free-free emission. The top of the synchrotron dominated NPS appears *orange*.

The *bottom right* projection is centred on CygA. There is lots of structure close to the Galactic plane with individual HII regions strong in free-free emission and synchrotron emission with a varying spectral index.

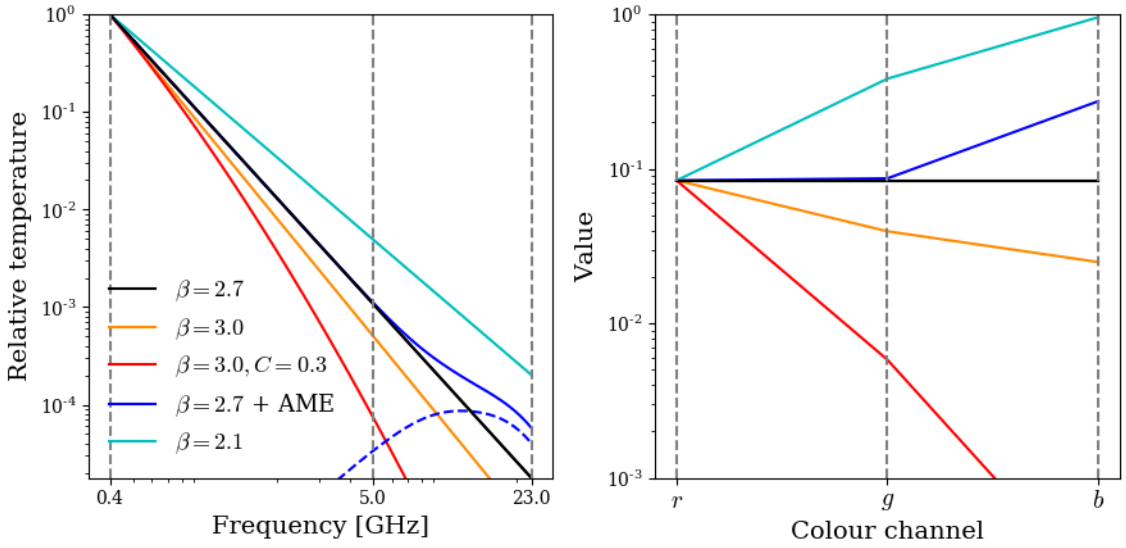

Figure 5.1: The *left hand* plot shows several representative temperature frequency spectra (relative to the temperature at 0.4 GHz). The *right hand* plot shows the resulting *rgb* colour channels when the maps are normalised so that $\beta = 2.7$ is colourless. The *dashed blue line* in the left hand plot is the frequency spectrum of typical AME. A pixel that only contained AME (and no other sources of emission) would have *red* and *green* values of zero and would therefore appear pure *blue*.

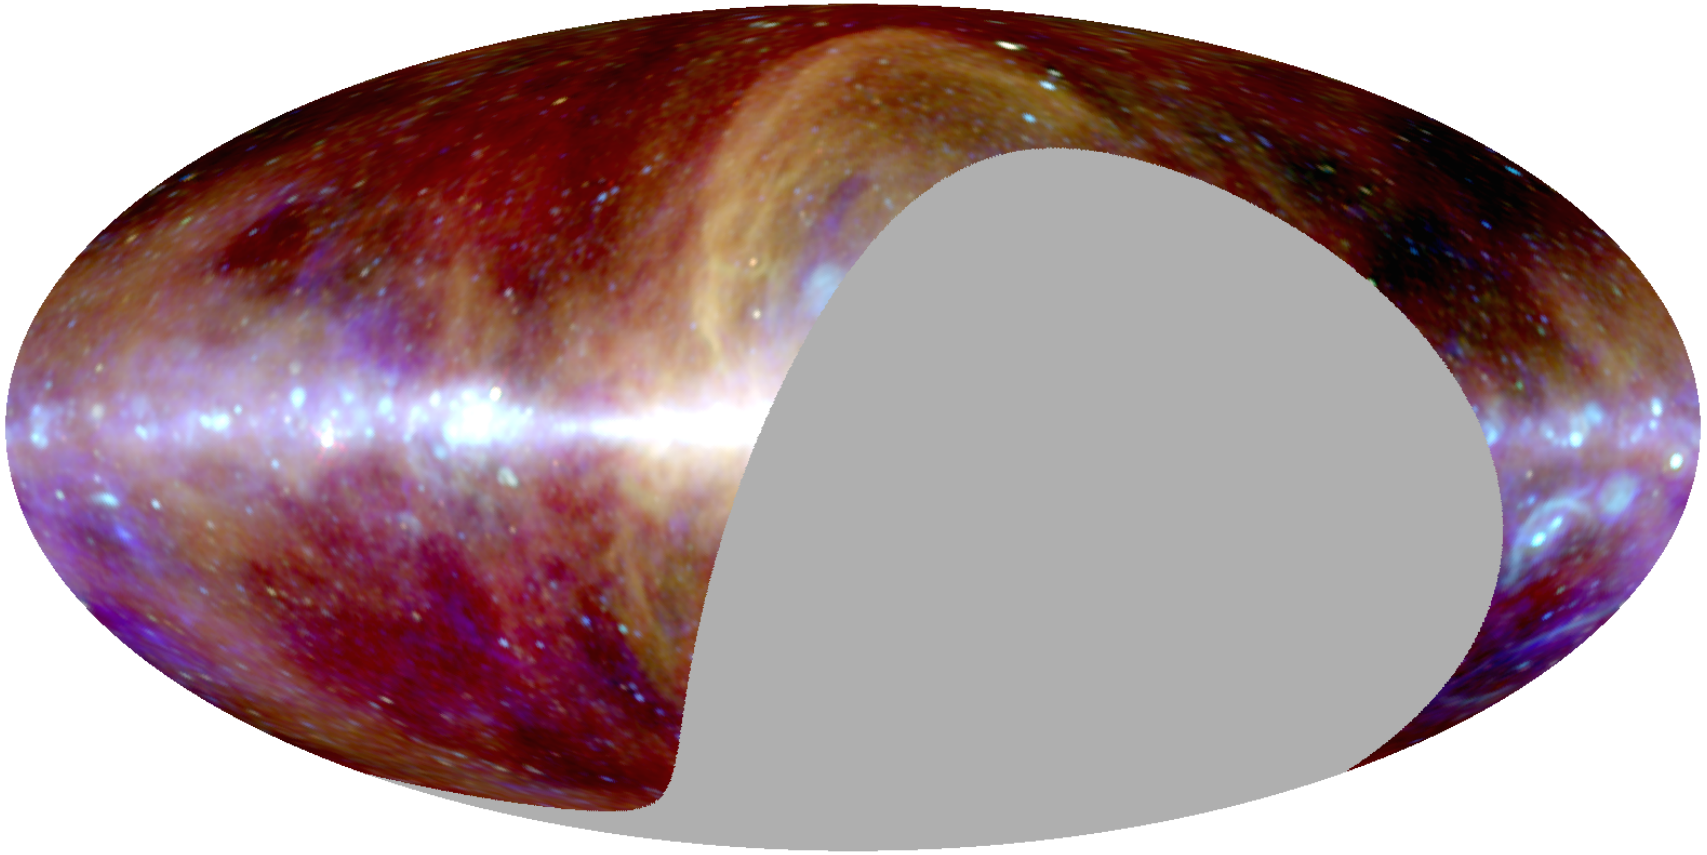

Figure 5.2: A 3-colour image in which the *red* channel is the Haslam 408 MHz map, the *green* channel is the C-BASS I map and the *blue* channel is the difference between *WMAP* K and V-band I maps. The colours are balanced such that a temperature spectrum with $\beta = 2.7$ is colourless. The intensity scale is highly non-linear.

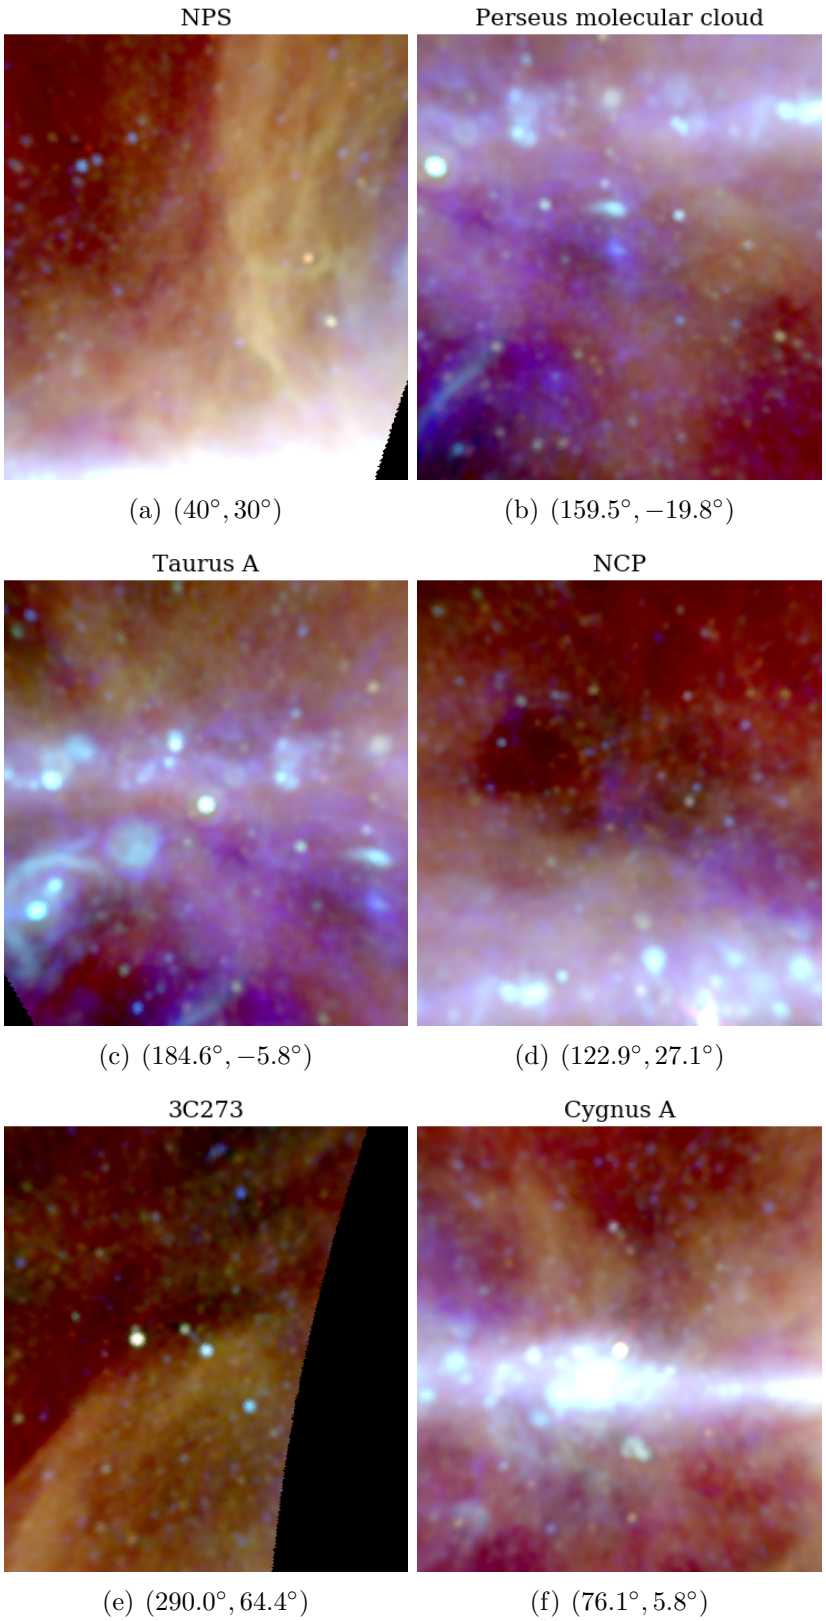

(a) $(40°, 30°)$ (b) $(159.5°, -19.8°)$

(c) $(184.6°, -5.8°)$ (d) $(122.9°, 27.1°)$

(e) $(290.0°, 64.4°)$ (f) $(76.1°, 5.8°)$

Figure 5.3: Gnomonic projections of the 3-colour image, see the text for details. The Galactic longitude and latitude of the centre of each projection is in degrees at the bottom of each image.

5.2 Constraining the synchrotron spectral index with T-T plots

In Section 1.2.1 we showed that diffuse Galactic synchrotron emission is expected to obey a power law over many decades of frequency and that the brightness temperature T_1 at a frequency ν_1 can therefore be parametrized as

$$T_1 = A_0 \left(\frac{\nu_1}{\nu_0} \right)^{-\beta}, \tag{5.1}$$

where A_0 is the temperature at reference frequency ν_0 and β is the synchrotron spectral index. If two maps at different frequencies are both dominated by synchrotron emission and have well constrained monopoles then β can be estimated pixel-by-pixel,

$$\beta = -\frac{\log(T_1/T_2)}{\log(\nu_1/\nu_2)}. \tag{5.2}$$

The presently available low frequency I maps have poorly constrained monopoles and so β can not be directly measured pixel-by pixel. The spectral index can instead be determined from so-called T-T plots, first used by Turtle et al. (1962). A plot of the temperature in one map against the temperature in another (a T-T plot) will obey a linear relationship if they are both dominated by the same emission mechanism with a power-law frequency spectrum. The slope of the linear relationship is determined by the spectral index and is not affected by global offsets in either map. In Section 5.2.1 we describe the method of T-T plots in more detail.

Because the spectral index varies across the sky and monopoles and dipoles in the maps introduce offsets that change over the sky, the maps should be broken into smaller regions over which the spectral index and offset is constant. T-T plots should then be used independently on each region. In Section 5.2.2 we describe how we used a Voronoi binning technique to divide the sky into the smallest possible regions that also contain a minimum total signal. Spectral index estimates found using the T-T plot method are incredibly robust in the correct circumstances. The insensitivity to global offsets in the map regions means that large scale residuals in the data have little effect on the recovered estimates of the spectral index.

To determine the spectral index from T-T plots requires unbiased estimates of the slope of the best-fit straight line to the data. The signal-to-noise ratios of C-BASS, Haslam, *Planck* and *WMAP* data are all comparable and in Section 5.2.3 we describe a fitting method that recovers unbiased estimates of the slope and intercept of a line in this regime.

If a map region contains multiple families of pixels that are each dominated by a different emission mechanism then the T-T plot will contain multiple families of points obeying different linear relationships. We extend the T-T plot method in an attempt to maintain its robustness but overcome its major limitation, that all of the pixels in a region must be dominated by emission obeying the same spectral parameters. In Section 5.2.4 we describe our extension to the standard T-T plot model that allows pixels to belong to one of several families using a mixture model.

In Section 5.2.5 we test the extended T-T plot model on simulated data and show that the total intensity foregrounds are too complex for this method to systematically and blindly determine β across the whole sky.

In Section 5.2.6 we use the method on the real C-BASS and Haslam maps, focussing our discussion on the NPS.

In Section 5.2.8 we focus on Barnard's Loop and estimate the spectral index between C-BASS and the Haslam, *WMAP* K-band and *Planck* 30 GHz maps in this region using T-T plots. Barnard's Loop is an HII region dominated by free-free emission across a broad range of frequencies.

5.2.1 Overview of the T-T plot method

If the sky is dominated by an emission mechanism with a power-law frequency spectrum then the brightness temperature of a map at a frequency ν_i can be parametrized,

$$\mathbf{T_i} = \nu_i^{-\beta}\mathbf{A} + B_i\mathbf{1}, \tag{5.3}$$

where $\mathbf{T_i}$ is the map vector, $\mathbf{A}$ is the reference amplitude map vector, B_i is the map monopole, β is the spectral index and $\mathbf{1}$ is the all-ones vector.

Stokes I maps are typically subject to arbitrary and unknown monopoles. Radio telescopes are often insensitive to global offsets in I and the destriping technique is insensitive to monopoles in I maps even if that information is present in the TOD. Map monopoles have to be determined from external datasets, for example measurements by ARCADE2 (Singal et al., 2011; Seiffert et al., 2011).

If the monopoles of two maps are unknown then β cannot be calculated pixel-by-pixel using Equation 5.2. It can, however, be determined using the method of T-T plots (Turtle et al., 1962). Eliminating $\mathbf{A}$ from Equation 5.3 for two maps, $\mathbf{T_1}$ and $\mathbf{T_2}$, gives

$$\mathbf{T_1} = \left(\frac{\nu_1}{\nu_2}\right)^{-\beta}\mathbf{T_2} + \left[B_1 - B_2\left(\frac{\nu_1}{\nu_2}\right)^{-\beta}\right]\mathbf{1}. \tag{5.4}$$

This is the equation of a straight line with $\mathbf{T_1}$ as the y ordinate and $\mathbf{T_2}$ as the x abscissa. The unknown monopoles B_1 and B_2 only impact upon the line's offset (the contents of the square brackets). The slope of the line, a, is determined by the spectral index and the frequencies of the two maps, which are known. It is trivial to show that

$$\beta = -\frac{\ln a}{\ln \frac{\nu_1}{\nu_2}}. \tag{5.5}$$

By fitting straight lines to T-T scatter plots, a and hence β can be estimated.

5.2.2 Voronoi binning

The spectral index varies spatially and monopoles and dipoles in the maps introduce offsets that change over the sky. The sky must therefore be split into regions over which the spectral parameters and map offsets are believed to be constant. T-T plots can then used to estimate β for each region. Fuskeland et al. (2014), for example, use T-T plots to measure the polarized synchrotron spectral index from Q and U maps using a physically motivated partitioning of the sky. They first divide the sky into two sections, a Galactic plane section and an off the plane section, before further subdividing these two sections into multiple regions. The Galactic plane section is divided into smaller regions than the off-plane section as the signal is stronger on the plane. As an alternative we partition the sky using Voronoi binning in order to create regions that are as small as possible whilst also having adequate signal to noise to form a meaningful T-T plot.

We adapted the Voronoi binning method of Cappellari & Copin (2003) to work on the surface of a sphere and used it to split the sky into the smallest regions that meet a target signal. This method divides the sky into regions and leaves no holes or overlaps. The regions are compact and have minimal variation in the total signal within them.

A two stage process is needed to construct the Voronoi regions. Stage one begins by assigning the highest signal pixel in a map to the first region before accreting neighbouring pixels, one at a time, until the threshold total signal in the region is reached. At each accretion step the closest unassigned pixel to the region centroid that also neighbours the region is accreted. Closeness is defined as the angular separation between the vectors to the region centroid and to the pixel centre. Once this region has either reached the threshold signal or there are no remaining neighbouring unassigned pixels, the next closest unassigned pixel to the region centroid is used

to start constructing a new region. This new starting pixel need not neighbour the previous region. This is repeated until there are no more unassigned pixels.

Stage two takes the centroids of each of the recovered regions that have reached the threshold signal and uses them as seeds to construct Voronoi cells over the surface of the sphere. The pixels in each cell are those that are closer to the cell's seed than any other seed. The resulting cells (or regions) then contain roughly equal signal whilst containing as few pixels as needed to reach the threshold.

When testing the T-T plot method on simulated data we calculated the angular separation, $\Delta\theta$, between two vectors $\mathbf{v_1}$ and $\mathbf{v_2}$ using $\Delta\theta = \arccos(\mathbf{v_1} \cdot \mathbf{v_2})$. When using the method on real data we used the more numerically accurate $\Delta\theta = \text{atan2}(\mathbf{v_2} \times \mathbf{v_2}, \mathbf{v_1} \cdot \mathbf{v_2})$.[2]

5.2.3 Straight line fitting with errors in both axes

Fitting a straight line parametrized with a slope a and an intercept b to data, $\{x, y\}$, is a common problem across all of science. When the errors on x are vanishingly small and the errors on y are Gaussian the method of least-squares can be used to find unbiased estimates of a and b. In many circumstances the required conditions for least-squares fitting to provide unbiased estimates of the parameters are not met. When using T-T plots, the uncertainties on $\mathbf{T}_1$ and $\mathbf{T}_2$ are often comparable to one another and using the method of least-squares is not appropriate. In this section we describe the correct posterior distribution for the problem, first identified by Gull (1989).

Let us assume that there is a true and exact underlying relationship between the unknown variables $\{\hat{x}, \hat{y}\}$ ($\hat{\mathbf{T}}_1$ and $\hat{\mathbf{T}}_2$) and that we have noisy measurements of them, $\{x, y\}$.

Gull's first insight was to treat each axis symmetrically by re-parametrizing the problem. The new unknown variables, X and Y, are related to x and y by

$$X = (x - x_0)/R_x \tag{5.6}$$

$$Y = (y - y_0)/R_y. \tag{5.7}$$

The exact underlying relationship between the new variables takes the simple form $\hat{X} = \pm\hat{Y}$. The two original parameters, a and b, are related to the four new parame-

[2] http://healpix.jpl.nasa.gov/html/subroutinesnode16.htm

ters, R_x, R_y, x_0 and y_0 by

$$a = R_y/R_x \tag{5.8}$$

$$b = y_0 - x_0 R_y/R_x. \tag{5.9}$$

This re-parametrization of the problem is a far more natural way of considering a line than through its slope and intercept. R_x and R_y determine the relative scalings of x and y values and x_0 and y_0 determine their offsets. Conceptually this is what we do when we plot a straight line on graph paper, we estimate these four parameters and use them to determine the appropriate axis ranges.

The posterior distribution to be maximised is,

$$p(\hat{x}, x_0, y_0, R_x, R_y|x, y) \propto p(x, y|\hat{x}, R_x, R_y, x_0, y_0)p(\hat{x}|R_x, R_y, x_0, y_0)p(R_x, R_y, x_0, y_0). \tag{5.10}$$

We can treat R_x and R_y as scale parameters and assign them $1/X$ priors and x_0 and y_0 as location parameters and assign them uniform priors. The final term in Equation 5.10 therefore takes the form $p(R_x, R_y, x_0, y_0) \propto \frac{1}{R_x R_y}$.

We don't know the true $\hat{x}$ (or equivalently $\hat{y}$) values and so must marginalise over them. Gull's second key insight was to assign $\hat{x}$ an informative prior, without which the resulting estimates of the parameters are biased. Conceptually this makes sense, if we have N measurements that all fall within a specific range then we do not expect the $N{+}1^{\text{th}}$ value to be too far away from the others. Maximum entropy arguments can be used to choose $p(\hat{x}|R_x, R_y, x_0, y_0)$. Obvious choices are either uniform or Gaussian distributions. For computational reasons Gull used a Gaussian prior on $\hat{x}$ and we use the same prior here.

If our noisy measurements of Stokes I $\{x, y\}$, have Gaussian errors, $\{\sigma_x, \sigma_y\}$, then the likelihood function is very simple. We can now write the terms in Equation 5.10, ignoring normalising constants, as functions of the data and the free parameters,

$$p(x, y|\hat{x}, R_x, R_y, x_0, y_0) \propto \prod_i \frac{1}{\sigma_{xi}\sigma_{yi}} e^{-\frac{(x_i-\hat{x}_i)^2}{2\sigma_{x_i}^2} - \frac{(y_i-\hat{y}_i)^2}{2\sigma_{y_i}^2}} \tag{5.11}$$

$$p(\hat{x}|R_x, R_y, x_0, y_0) \propto \prod_i \frac{1}{R_x} e^{-\frac{(\hat{x}_i-x_0)^2}{2R_x^2}} \tag{5.12}$$

$$p(R_x, R_y, x_0, y_0) \propto \frac{1}{R_x R_y}. \tag{5.13}$$

5.2.4 Extension to standard T-T plot method

When using T-T plots, great care has to be taken to define regions that are dominated by a single emission mechanism with uniform spectral parameters across each region. We want to estimate the synchrotron spectral index across the sky but it is impossible to divide the entire sky, a priori, into regions that are dominated by a single emission mechanism. We extended the usual T-T plot model and use a mixture model to allow for three different populations of pixels. This mixture model is appropriate if each data point comes from one of the three separate populations and we are not interested in which pixel belongs to which population. The mixture model likelihood function is

$$p(x, y|...) = \prod_i (1 - f^{\mathrm{I}} - f^{\mathrm{II}})p^{\mathrm{I}}(x_i, y_i|...) + f^{\mathrm{II}}p^{\mathrm{II}}(x_i, y_i|...) + f^{\mathrm{III}}p^{\mathrm{III}}(x_i, y_i|...), \quad (5.14)$$

where $p^j(x_i, y_i|...)$ is the likelihood of observing (x_i, y_i) if it belongs to population j and f^j is the fraction of data points that belong to population j. All of the uninteresting parameters can be marginalised over to give the posterior distributions of the spectral indices.

Component I in the mixture model is used to model pixels that are dominated by synchrotron emission. We use an uninformative prior on the slope of this line.

Component II models pixels that are dominated by free-free emission, which also has a power-law frequency spectrum. The free-free spectral index is close to 2.1 over much of the sky and we can therefore use an informative prior on the slope of this component. The informative prior helps to remove the degeneracy between this component and the first power law.

Component III models outlying pixels that are not dominated by a single emission mechanism or are dominated by emission from radio point sources. We modelled the outliers as a diffuse 2D Gaussian with no covariance between x and y. Because we do not know the frequency spectrum of outlying points we cannot model them accurately, but when modelling outliers the exact form of p^{outlier} is often not particularly important (Hogg, Bovy & Lang, 2010).

The priors on the free parameters in the model are given in Table 5.2.

By explicitly defining the parameter $\hat{y}_i^{\mathrm{I}}$ as

$$\hat{y}_i^{\mathrm{I}} = x_i \frac{R_y^{\mathrm{I}}}{R_x^{\mathrm{I}}} + y_0^{\mathrm{I}} - \frac{R_y^{\mathrm{I}}}{R_x^{\mathrm{I}}} x_0^{\mathrm{I}}, \quad (5.15)$$

Table 5.2: A table of the priors for the free parameters in the T-T plot mixture model.

Parameter	Prior
Component I (line with uninformative prior on slope)	
x_0^{I}	$\propto$ constant
y_0^{I}	$\propto$ constant
R_x^{I}	$\propto 1/R_x^{\mathrm{I}}$
R_y^{I}	$\propto 1/R_y^{\mathrm{I}}$
Component II (line with informative prior on slope)	
x_0^{II}	$\propto$ constant
y_0^{II}	$\propto$ constant
R_x^{II}	$\propto 1/R_x^{\mathrm{II}}$
R_y^{II}	$\propto \frac{1}{\sigma_{\mathrm{ff}}} e^{-\frac{(R_y^{\mathrm{II}}-\mu_{\mathrm{ff}})^2}{2\sigma_{\mathrm{ff}}^2}}$ where $\mu_{\mathrm{ff}} = R_x^{\mathrm{II}} \left(\frac{\nu_1}{\nu_2}\right)^{2.1}$ and $\sigma_{\mathrm{ff}} = 0.05$
f^{II}	$\propto$ constant, $0 < f^{\mathrm{II}} < 1$
Component III (outliers)	
x_0^{III}	$\propto$ constant
y_0^{III}	$\propto$ constant
R_x^{III}	$\propto 1/R_x^{\mathrm{III}}$
R_y^{III}	$\propto 1/R_y^{\mathrm{III}}$
f^{III}	$\propto$ constant, $0 < f^{\mathrm{III}} < 1 - f^{\mathrm{II}}$
Other	
$\hat{x}$	$\propto \prod_i (1 - f^{\mathrm{II}} - f^{\mathrm{III}}) \frac{1}{R_x^{\mathrm{I}}} e^{-\frac{(x_i - x_0^{\mathrm{I}})^2}{2R_x^{\mathrm{I}2}}} + f^{\mathrm{II}} \frac{1}{R_x^{\mathrm{II}}} e^{-\frac{(x_i - x_0^{\mathrm{II}})^2}{2R_x^{\mathrm{II}2}}} + f^{\mathrm{III}} \frac{1}{R_x^{\mathrm{III}}} e^{-\frac{(x_i - x_0^{\mathrm{III}})^2}{2R_x^{\mathrm{III}2}}}$

with a similar relation for $\hat{y}_i^{\mathrm{II}}$ and defining $f^{\mathrm{I}} = 1 - f^{\mathrm{II}} - f^{\mathrm{III}}$, we can write the likelihood function

$$
\begin{aligned}
p(x, y | \hat{x}, R_x^{\mathrm{I}}, R_y^{\mathrm{I}}, x_0^{\mathrm{I}}, y_0^{\mathrm{I}}, R_x^{\mathrm{II}}, R_y^{\mathrm{II}}, x_0^{\mathrm{II}}, y_0^{\mathrm{II}}, R_x^{\mathrm{III}}, R_y^{\mathrm{III}}, x_0^{\mathrm{III}}, y_0^{\mathrm{III}}, f^{\mathrm{II}}, f^{\mathrm{III}}) \propto \\
\prod_i \frac{1}{\sigma_{xi}} e^{-\frac{(x_i - \hat{x}_i)^2}{2\sigma_{x_i}^2}} \left(\frac{f^{\mathrm{I}}}{\sigma_{y_i}} e^{-\frac{(y_i - y_i^{\mathrm{I}})^2}{2\sigma_{y_i}^2}} + \right. \\
\frac{f^{\mathrm{II}}}{\sigma_{y_i}} e^{-\frac{(y_i - \hat{y}_i^{\mathrm{II}})^2}{2\sigma_{y_i}^2}} + \\
\left. \frac{f^{\mathrm{III}}}{\sqrt{\sigma_{y_i}^2 + R_y^{\mathrm{III}2}}} e^{-\frac{(y_i - y_0^{\mathrm{III}})^2}{2(\sigma_{y_i}^2 + R_y^{\mathrm{III}2})}} \right).
\end{aligned}
\tag{5.16}
$$

The complicated and highly multidimensional posterior distribution, found by multiplying priors in Table 5.2 and the likelihood in Equation 5.16 together, can be sampled using MCMC methods. We sample the posterior distribution using the Metropolis-Hastings algorithm, implemented with PyMC.

The spectral index is calculated for each step in the chains using,

$$\beta = -\frac{\ln(R_y^{\rm I}/R_x^{\rm I})}{\ln(\nu_1/\nu_2)}. \tag{5.17}$$

These chains of β are then our estimate of its marginalised posterior distribution.

5.2.5 Tests on simulated data

We tested our extended T-T plot method on simulated C-BASS and Haslam maps. The method recovered unbiased estimates of the synchrotron spectral index when the simulation only contained a synchrotron component and realisations of the noise. When the simulation also included a free-free component the method was unable to recover unbiased estimates of the synchrotron spectral index. The recovered estimates of the synchrotron spectral index were biased low by the free-free component. Modelling pixels as either dominated by synchrotron emission, free-free emission or as outliers is a poor approximation in many pixels. In this section we present the results from those simulations.

We simulated diffuse Galactic total intensity emission at the nominal frequencies of the C-BASS (5 GHz) and Haslam (0.408 GHz) surveys. At the higher frequencies of the *WMAP* K-band and *Planck* 30 GHz surveys AME is also important and would further complicate the analysis. We used PySM (Thorne et al., 2016) to generate maps of the synchrotron and free-free total intensity emission at 5 GHz that were smoothed to 1° and pixelised at $N_{\rm side} = 64$. We extrapolated both components to 0.408 GHz using a synchrotron spectral index of 3 and a free-free spectral index of 2.1.

We added realisations of the noise to the maps. The noise maps were generated by multiplying each pixel in the sensitivity maps by a random number drawn from the standard normal distribution. We used the sensitivity map of a C-BASS map produced from the E37S60 dataset and did not convert the units from antenna temperature to brightness temperature. The $N_{\rm side} = 64$ Haslam sensitivity map was set to 1 K plus 1% of the total intensity map. The simulated C-BASS map is therefore subject to slightly less noise than the current real C-BASS map and the simulated Haslam map is subject to more noise than the real map. These differences in the exact noise levels are not significant for the purposes of testing the method.

We generated the sky regions from the simulated C-BASS map containing only synchrotron emission using the Voronoi binning technique described in Section 5.2.2. The target signal per region was 2.5 K with a minimum of 20 pixels per region. This

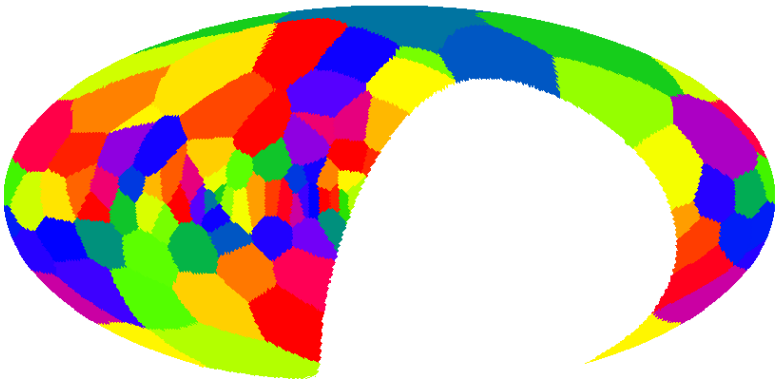

Figure 5.4: T-T plot regions generated by using the Voronoi binning technique on a simulated C-BASS map containing only synchrotron emission and noise. The target signal in each region was 2.5 K with a minimum of 20 pixels per region.

target was chosen as it typically gave sufficient signal-to-noise for the T-T plots to be meaningful whilst also dividing bright and complex regions of the sky into small areas. A map of the regions is shown in Figure 5.4. We used the less numerically accurate arccos function to calculate the angles between vectors.

5.2.5.1 Only synchrotron emission

First we tested the fitting procedure and the appropriateness of the mixture model on maps that only contained a synchrotron component and realisations of the noise. We fit a mixture model including components I and III (one line with an uninformative prior on its slope and a population of outliers) and also a mixture model including components I, II and III (2 lines, one of which has an informative prior on its spectral index to keep it near 2.1, and a population of outliers). Maps of the recovered estimates of β and the uncertainty on those estimates are shown in Figure 5.5.

Scatter plots of the estimated values of β with 1-σ error bars are shown in Figure 5.6 and histograms of the normalised deviation of the estimates are shown in Figure 5.7. The normalised deviation is the difference between the estimated value of β and the true value used to simulate the data ($\beta = 3$) divided by the uncertainty of the estimate.

There are no significant biases in the estimates of β. There is no reason for the posterior distribution of β to be exactly Gaussian but the histograms of the normalised deviations are fairly consistent with the standard normal distribution and

so the 1-σ error bars can approximately be considered as such. Importantly, when no free-free emission is present in the data, including the extra power-law component in the mixture model (component II) has not biased the estimates of the spectral index of the synchrotron component.

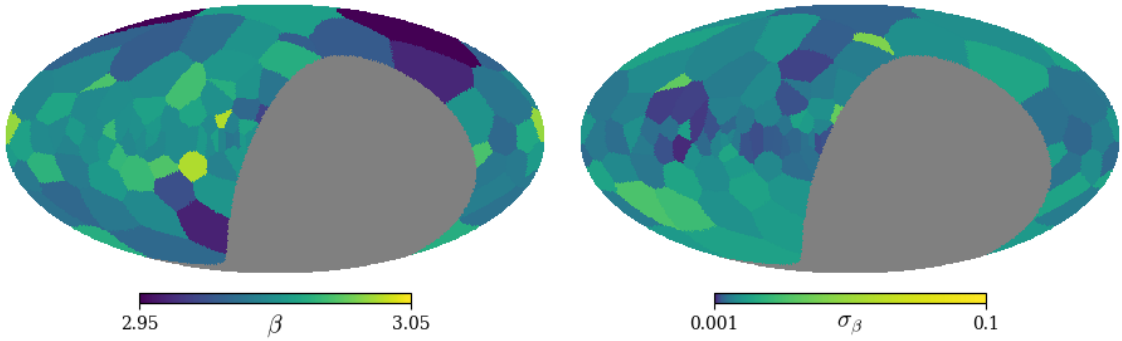

(a) 1 line and outliers

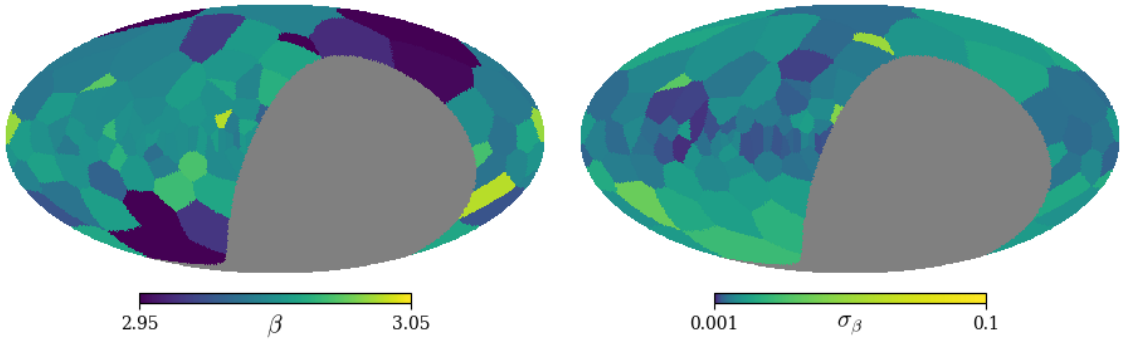

(b) 2 lines and outliers

Figure 5.5: Maps of the recovered estimates of β and σ_β from simulated C-BASS and Haslam maps that only contain a synchrotron component and realisations of the noise. The *top* maps show the estimates when the mixture model fitted to each T-T plot included components I and III, i.e. not the second straight line. The *bottom* maps show the estimates when the mixture model included all three components. The true synchrotron spectral index used to simulate the data was $\beta = 3$ across the entire sky. The spectral index maps are on linear colour scales and the uncertainty maps are on logarithmic colour scales.

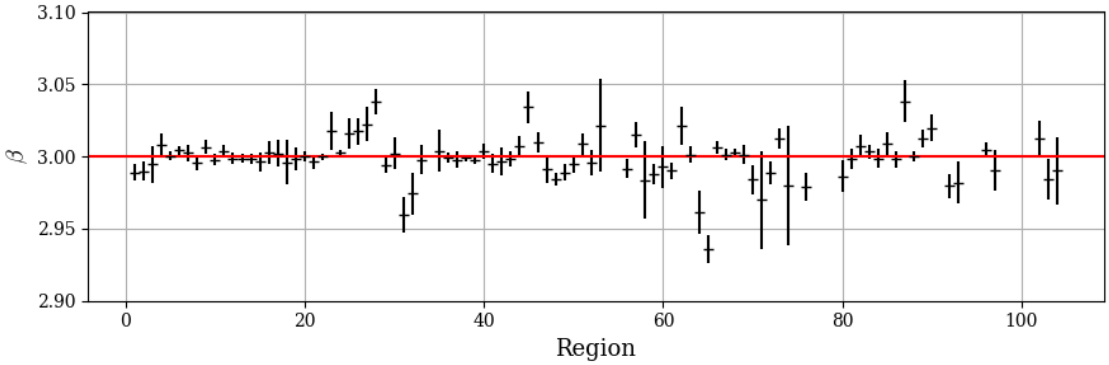

(a) 1 line and outliers

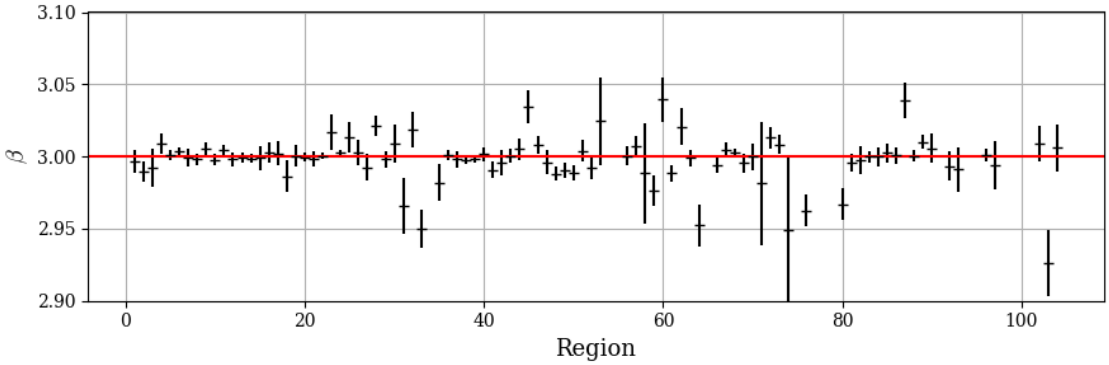

(b) 2 lines and outliers

Figure 5.6: Scatter plots of the estimated values of β with 1-σ errorbars from simulated C-BASS and Haslam maps that contain a synchrotron component and realisations of the noise. The *top* panel shows the estimates when the mixture model fitted to each T-T plot included components I and III, i.e. not the second straight line. The *bottom* panel shows the estimates when the mixture model included all three components. The *red* line is at the true synchrotron spectral index that was used to simulate the data.

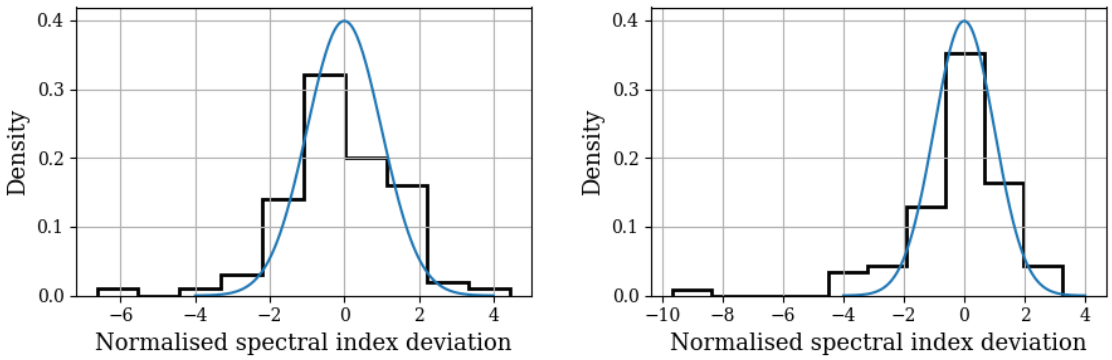

(a) 1 line and outliers

(b) 2 lines and outliers

Figure 5.7: Histograms of the normalised deviation of the estimated spectral index from simulated C-BASS and Haslam maps that contain a synchrotron component and realisations of the noise. The *left* panel shows the estimates when the mixture model fitted to each T-T plot included components I and III, i.e. not the second straight line. The *right* panel shows the estimates when the mixture model included all three components.

5.2.5.2 Synchrotron and free-free components

Secondly we tested the procedure on simulated maps that also included a free-free component. Figure 5.8 shows maps of the means and standard deviations of the posterior distributions of β for each region and Figure 5.9 shows the estimates as scatter plots. From this second dataset we could not recover reliable estimates of the synchrotron spectral index, β.

The T-T plots of a representative sample of the regions are shown in Figure 5.10.

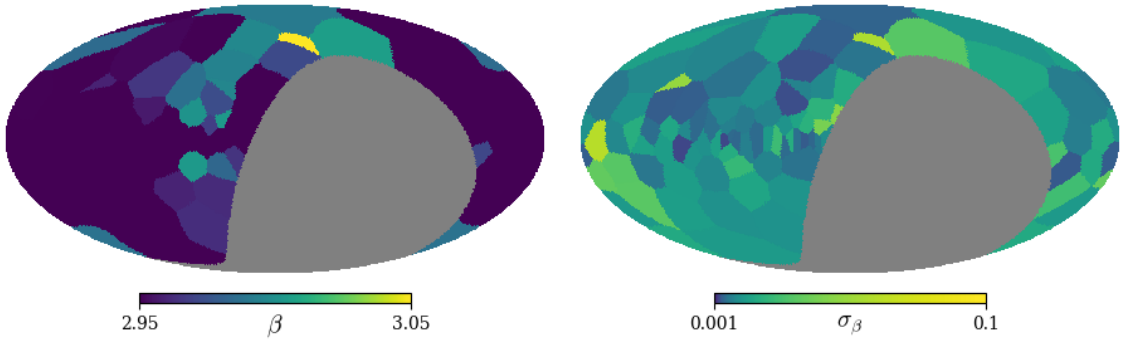

(a) 1 line and outliers

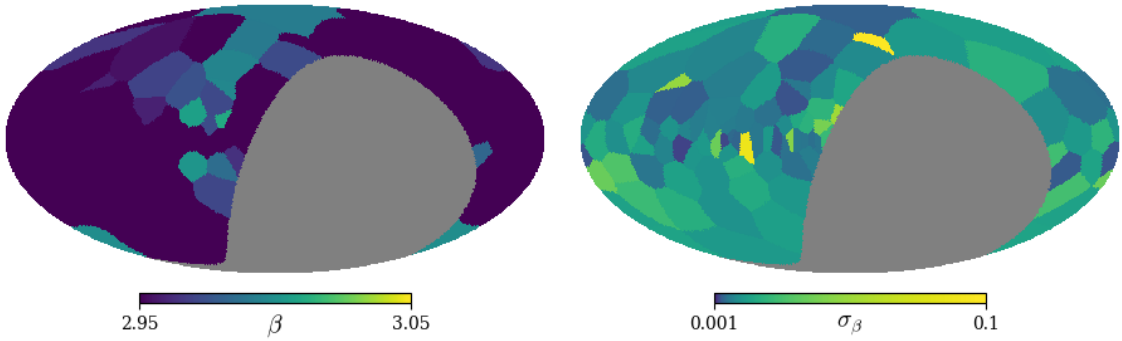

(b) 2 lines and outliers

Figure 5.8: Maps of the recovered estimates of β and σ_β from simulated C-BASS and Haslam maps that contain synchrotron and free-free components and realisations of the noise. The *top* maps show the estimates when the mixture model fitted to each T-T plot included components I and III, i.e. not the second straight line. The *bottom* maps show the estimates when the mixture model included all three components, i.e. 2 lines and outliers. The spectral index maps are on linear colour scales and the uncertainty maps are on logarithmic colour scales. These maps show that neither of the mixture models could recover unbiased estimates of the synchrotron spectral index at low Galactic latitude. The true synchrotron spectral index used to simulated the data was $\beta = 3$, the free-free emission was extrapolated with a spectral index of 2.1.

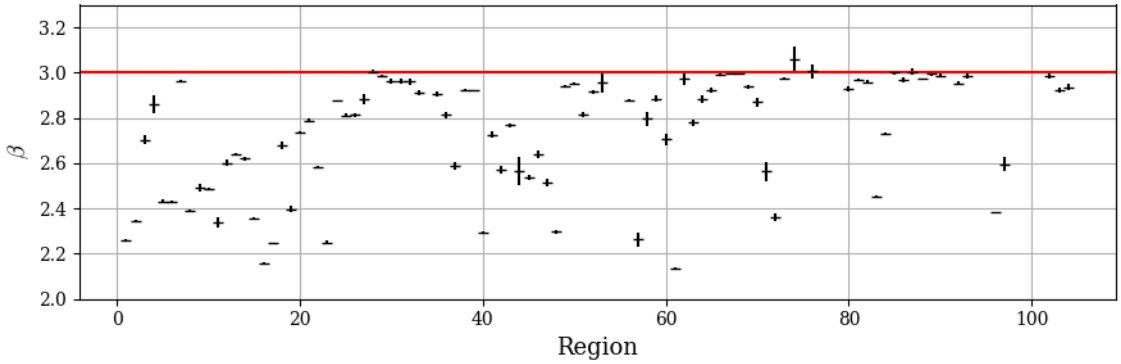

(a) 1 line and outliers

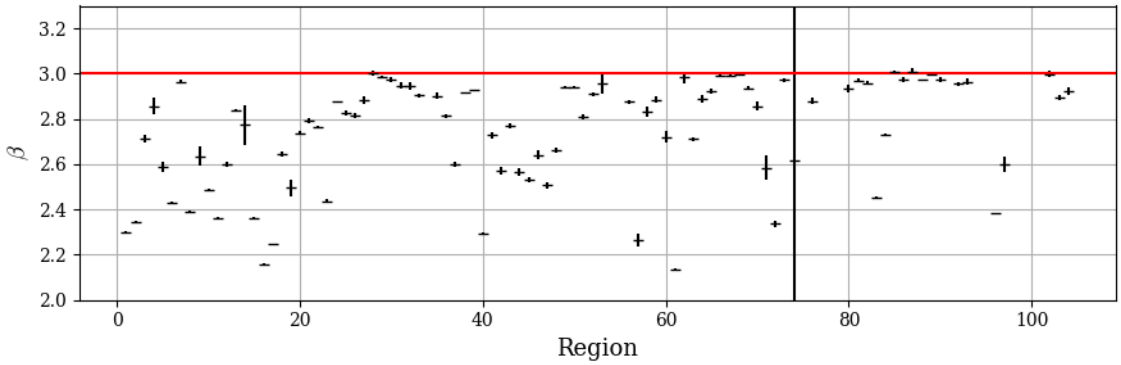

(b) 2 lines and outliers

Figure 5.9: Scatter plots of the estimated values of β with 1-σ error bars from the simulated C-BASS and Haslam maps that contain synchrotron and free-free components and realisations of the noise. The *red line* shows the true synchrotron spectral index used to generate the data ($\beta = 3$). The free-free spectral index was 2.1. The *top* plot shows the recovered estimates when the mixture model only included components I and III, i.e. not the second straight line. The *bottom* plot shows the recovered estimates when the mixture model included all three components, i.e. 2 lines and outliers. These scatter plots show that neither mixture model could consistently recover estimates of the synchrotron spectral index across the sky.

The *first row* show regions that were dominated by emission with a spectral index of 3. The *second row* show regions where none of the pixels are dominated by a single component and the combination of free-free and synchrotron emission results in an effective spectral index shallower than 3. The *third row* shows regions with two distinct populations of pixels, each dominated by a different spectral index. Regions like this are often misfitted and the *fourth row* shows examples regions where this misfitting has occurred.

To avoid the sampler getting stuck in local maxima we could try using a Gibbs

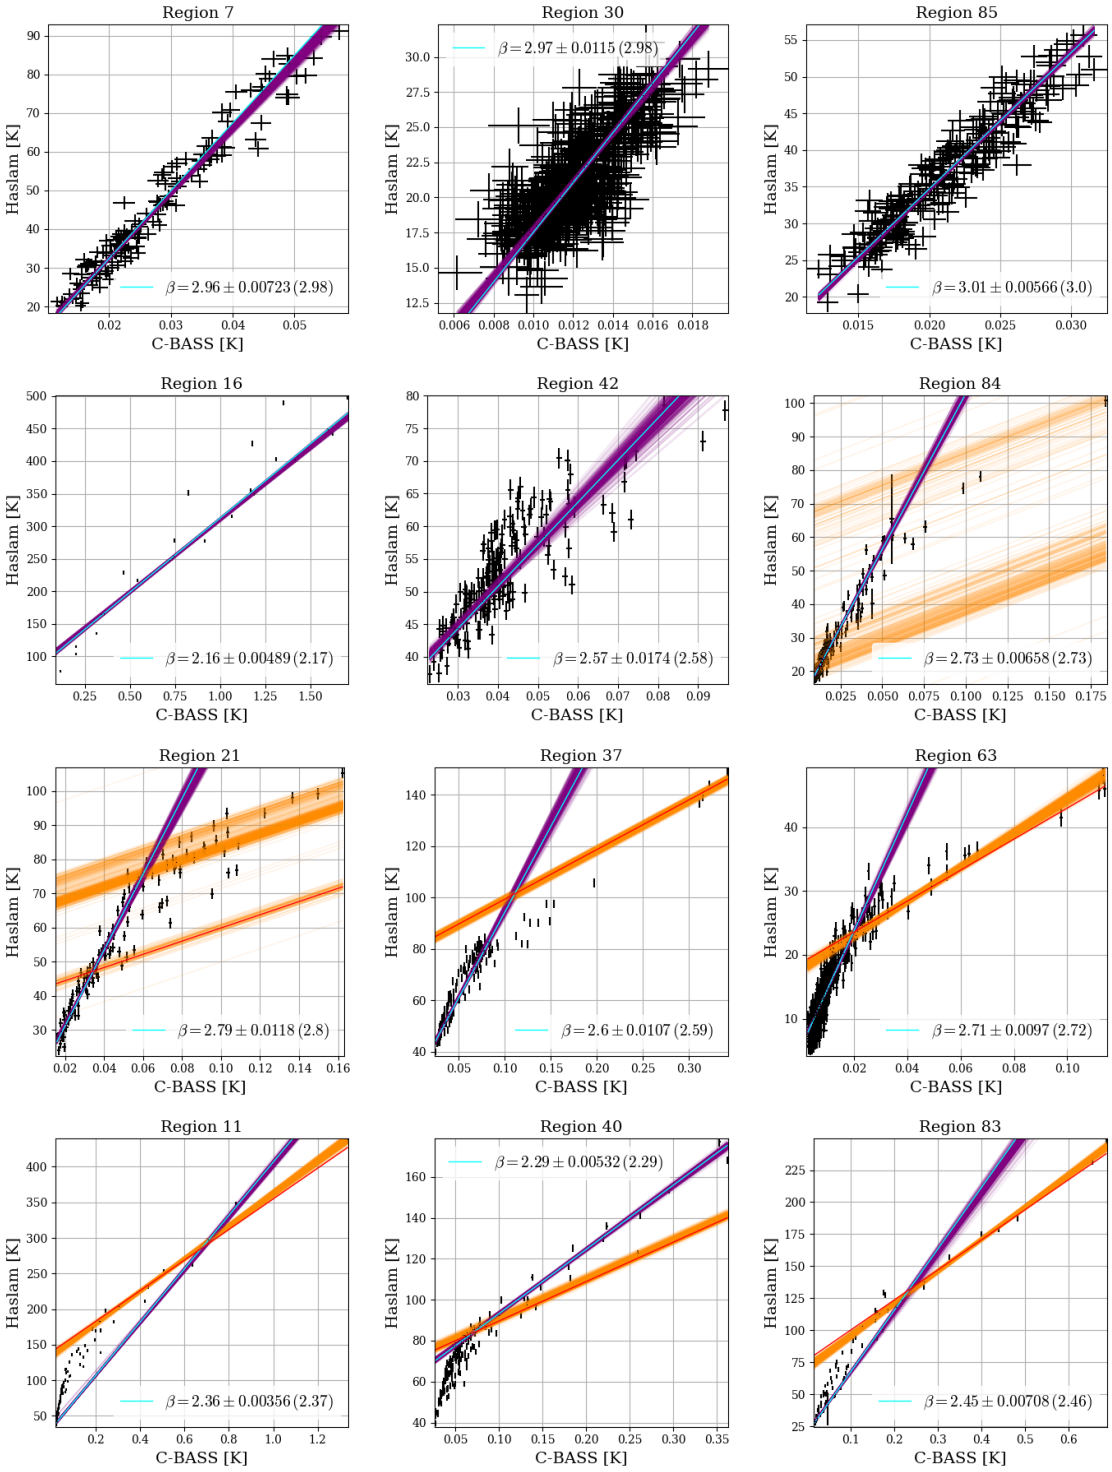

Figure 5.10: T-T plots in a representative sample of regions from the simulated C-BASS (x-axes) and Haslam (y-axes) data that include synchrotron and free-free components and realisations of the noise. A thinned subset of the converged chains lines are plotted. The *purple* lines are component I and the *orange* lines are component II, component III is not plotted. The estimated mean and standard deviation of the posterior distribution of the spectral index for component I are labelled along with the maximum posterior value in brackets.

sampling algorithm to explore the parameter space more efficiently and also try to improve the starting values of the parameters. A more informed dividing of the sky into regions might also help. The simulations have shown that this method does not lend itself to an automated blind-analysis of the whole sky.

This set of simulations did not include an AME component or radio point sources. These extra components make the true sky even more complicated than these simulations, particularly at higher frequencies.

5.2.6 Results from real data

We estimated the spectral index between the real C-BASS and Haslam maps using T-T plots and the method outlined above, noting that the estimated spectral index in each region is likely caused by a superposition of synchrotron and free-free emission. Maps at frequencies above 5 GHz, such as those of *WMAP* and *Planck*, will be increasingly contaminated by AME and other foregrounds and the T-T plot method will be even less appropriate across large areas of the sky.

In this section we describe the real C-BASS and Haslam datasets, use the Voronoi binning technique to divide the sky into regions and estimate the spectral index in each region by fitting a mixture model to the data.

We used the destriped Haslam map produced by Remazeilles et al. (2015). The noise in each $N_{\rm side} = 256$ pixel was set to 1 K. We did not apply any further smoothing to the map as its precise FWHM is not known but is approximately 1°. We then downgraded the map to $N_{\rm side} = 64$ using the HEALPix UD_GRADE routine, scaled the downgraded map by the factor $\eta_{\rm dg}$ calculated using Equation 4.18 and added 1% of the total intensity measurement to the white noise map in quadrature (as in Planck Collaboration et al. (2016b)). We did not apply any colour corrections to the Haslam map as they are unknown.

We used the C-BASS I map described in Section 4.5.1. When converting to brightness temperature we used the effective frequency $\nu_{\rm eff} = 4.83$ GHz, this is appropriate for emission with spectral index $\beta = 0.0$.

By not accounting for the spectrum of emission across the bands of either of the surveys we will have introduced gain errors into the maps at the $\sim 1\%$ level (Irfan et al., 2015).

The C-BASS and Haslam I and σ_I maps are shown in Figure 5.11. Table 5.3 lists their centre frequencies and typical sensitivities.

We used the Voronoi binning technique described in Section 5.2.2 to divide the sky into regions, using the more numerically accurate atan2 method of calculating

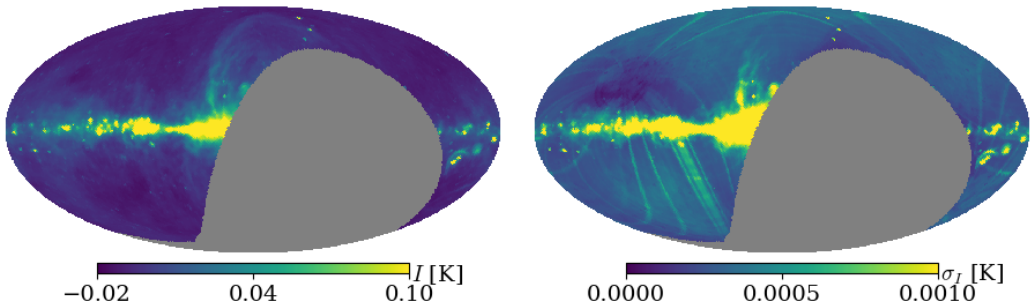

(a) C-BASS I map

(b) C-BASS σ_I map

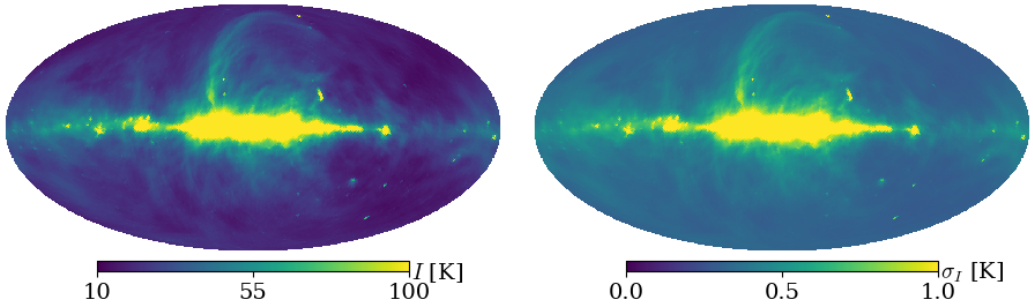

(c) Haslam I map

(d) Haslam σ_I map

Figure 5.11: C-BASS and Haslam I and σ_I maps downgraded to $N_{\text{side}} = 64$. All maps are in brightness temperature and on linear colour scales.

Table 5.3: The C-BASS and Haslam centre frequencies and the typical sensitivities of the $N_{\text{side}} = 64$ I maps.

Map	Frequency [GHz]	Typical $\sigma_I(1°)$
C-BASS	4.83	0.31 mK_{RJ}
Haslam	0.408	0.36 K_{RJ}

the angle between vectors. We created the bins from the C-BASS map, the target signal per region was 2.5 K and the minimum number of pixels was 20. A map of the regions is shown in Figure 5.12.

In the simulations, including a second line did not provide better estimates of the synchrotron spectral index and so we do not include a second line here. The mixture model included components I and III, i.e. a single line with an uninformative prior on its slope and a population of outliers. This model was fit to each region.

Maps of the mean and standard deviation of the spectral index from the MCMC chains are shown in Figure 5.13. The maps show that the effective spectral index is

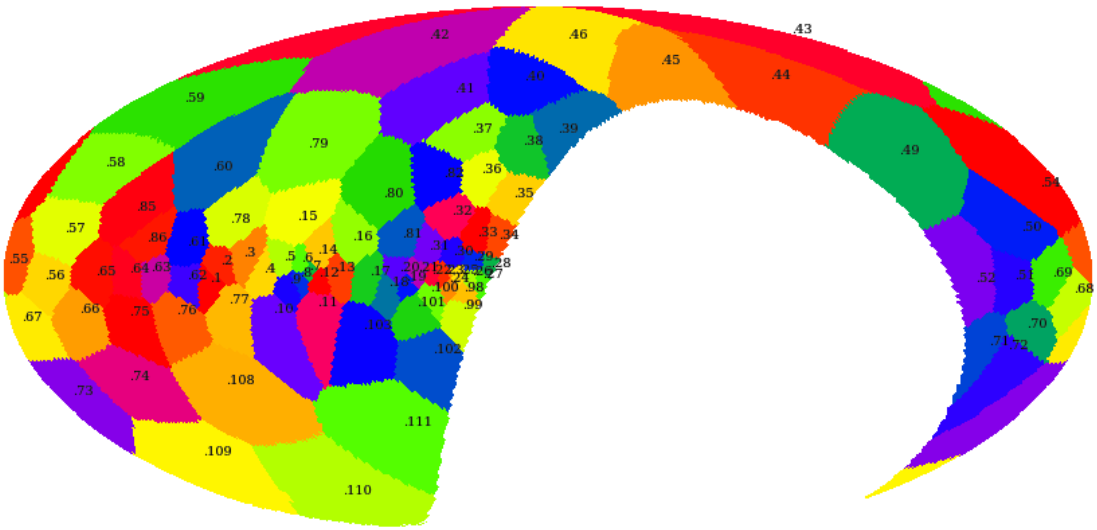

Figure 5.12: Voronoi regions produced from the real C-BASS map with a target signal of 2.5 K per region.

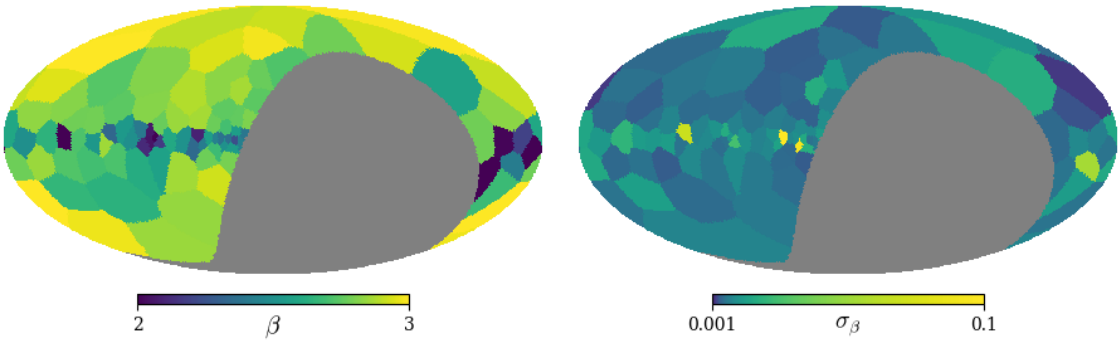

Figure 5.13: Map of the estimated spectral index (*left*) and its uncertainty (*right*) in each region between the real C-BASS and Haslam maps. The spectral index map is on a linear colour scale and the uncertainty map is on a logarithmic colour scale.

shallowest along the Galactic plane. The spectral index is only close to 3 in a small number of regions that are far from the plane.

We show the T-T plots for a set of representative regions that follow the NPS. A gnomonic projection of the region boundaries is shown in Figure 5.14. We find that the spectral index between the C-BASS and Haslam maps increases along the NPS but can not determine whether this is caused by a changing synchrotron spectral

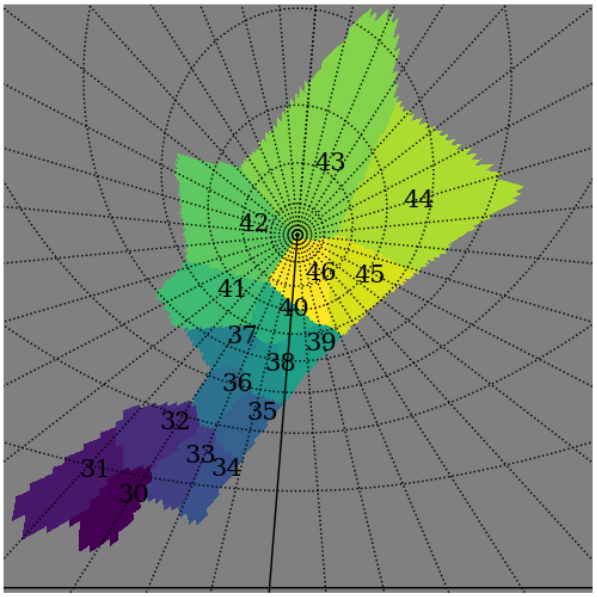

Figure 5.14: Regions along the NPS, starting close to the Galactic plane and extending up to the north Galactic pole. The graticules show lines of constant Galactic latitude and longitude, separated by 10° intervals.

index or varying mixtures of free-free and synchrotron emission in the pixels.

The NPS passes through or close to regions 30–46. It starts close to the Galactic plane in regions 30 and 31. Most of the pixels in these regions are within 10° of the Galactic plane. Travelling along the NPS away from the plane the region numbers increase. Close to the north Galactic pole the NPS ends and the signal level drops. The T-T plots for these regions are shown in Figures 5.15, 5.16 and 5.17. Underneath the T-T plot for each region we also show the region in the 3-colour map.

The T-T plots for regions 30–35 are shown in Figure 5.15. The sky here is bright and so the Voronoi binning method produces small regions. The spectral indices in these regions are typically 2.7–2.8. Close to the plane pixels are not likely to be dominated by a single emission mechanism and will contain both free-free and synchrotron radiation. Region 35 contains part of SH2-27, which is strongly dominated by free-free emission (the cyan pixels at the top right of the 3-colour map projection). The T-T plot for this region demonstrates the success of the fitting method, there are a collection of pixels in the plot that do not obey the straight-line relationship that is expected for a power law. The fitting algorithm has successfully identified the

straight line component.

Figure 5.16 shows the T-T plots and 3-colour maps for regions 36–41. These regions are further from the Galactic plane and the 3-colour maps have more red and orange pixels than regions 30–35, which are closer to the plane. The 3-colour maps suggest that spectral indices in these regions are steeper than $\beta = 2.7$ (the 3-colour map is normalised such that emission with a spectral index of $\beta = 2.7$ is colourless). The typical spectral indices found by the fits to the T-T plots of these regions are 2.8–3. The small number of outlying pixels in the T-T plots of regions 36 and 41 are caused by radio point source emission. The outlying pixels in region 38 are from the top of SH2-27.

Figure 5.17 shows the regions 42–46. Regions 42, 43 and 44 are close to the north Galactic pole and the NPS does not pass through them. The 3-colour maps are largely red, indicating steep spectra, but also not bright. In the low signal regime, small errors in the map monopoles cause large errors in the 3-colour map but not in the T-T plots. The T-T plots for these low signal regions exhibit lots of scatter but the estimated spectral indices are 2.9–3.0. The T-T plot for region 44 has two populations of pixels, one with a slope shallower than that fitted. The 3-colour map of region 44 has two vertical patches of orange pixels that are responsible for the shallower slope. The NPS does pass through regions 45 and 46. The bright radio point sources Virgo A, 3C283 and 3C279 are responsible for the outlying pixels in the T-T plot in region 45. The fitted line has not been significantly biased by their presence.

Two further T-T plots are shown in Figure 5.18. These are for regions 63 and 64, which both lie close to the Galactic plane. The T-T plots have two populations of pixels, one obeying a shallow spectral index, of around 2, and one obeying a steeper spectral index, around 2.7. In region 63 the fitting routine has recovered the spectral index of the steeper component and in region 64 it has recovered the spectral index of the shallower component. Regions such as these would be better fit if we could reliably include a second straight line component in the mixture model.

The scatter plot in Figure 5.19 shows the estimated spectral indices for all regions and their 1-σ error bars. The distribution of points for real data is similar to the distribution in Figure 5.9, which shows the estimated spectral indices when fitting the mixture model to simulated maps that include synchrotron and free-free components. The simulations showed that many pixels contain a mixture of synchrotron and free-free emission and that the T-T plots can still appear to obey a power law but with a spectral index between that of free-free and synchrotron, even far from the Galactic

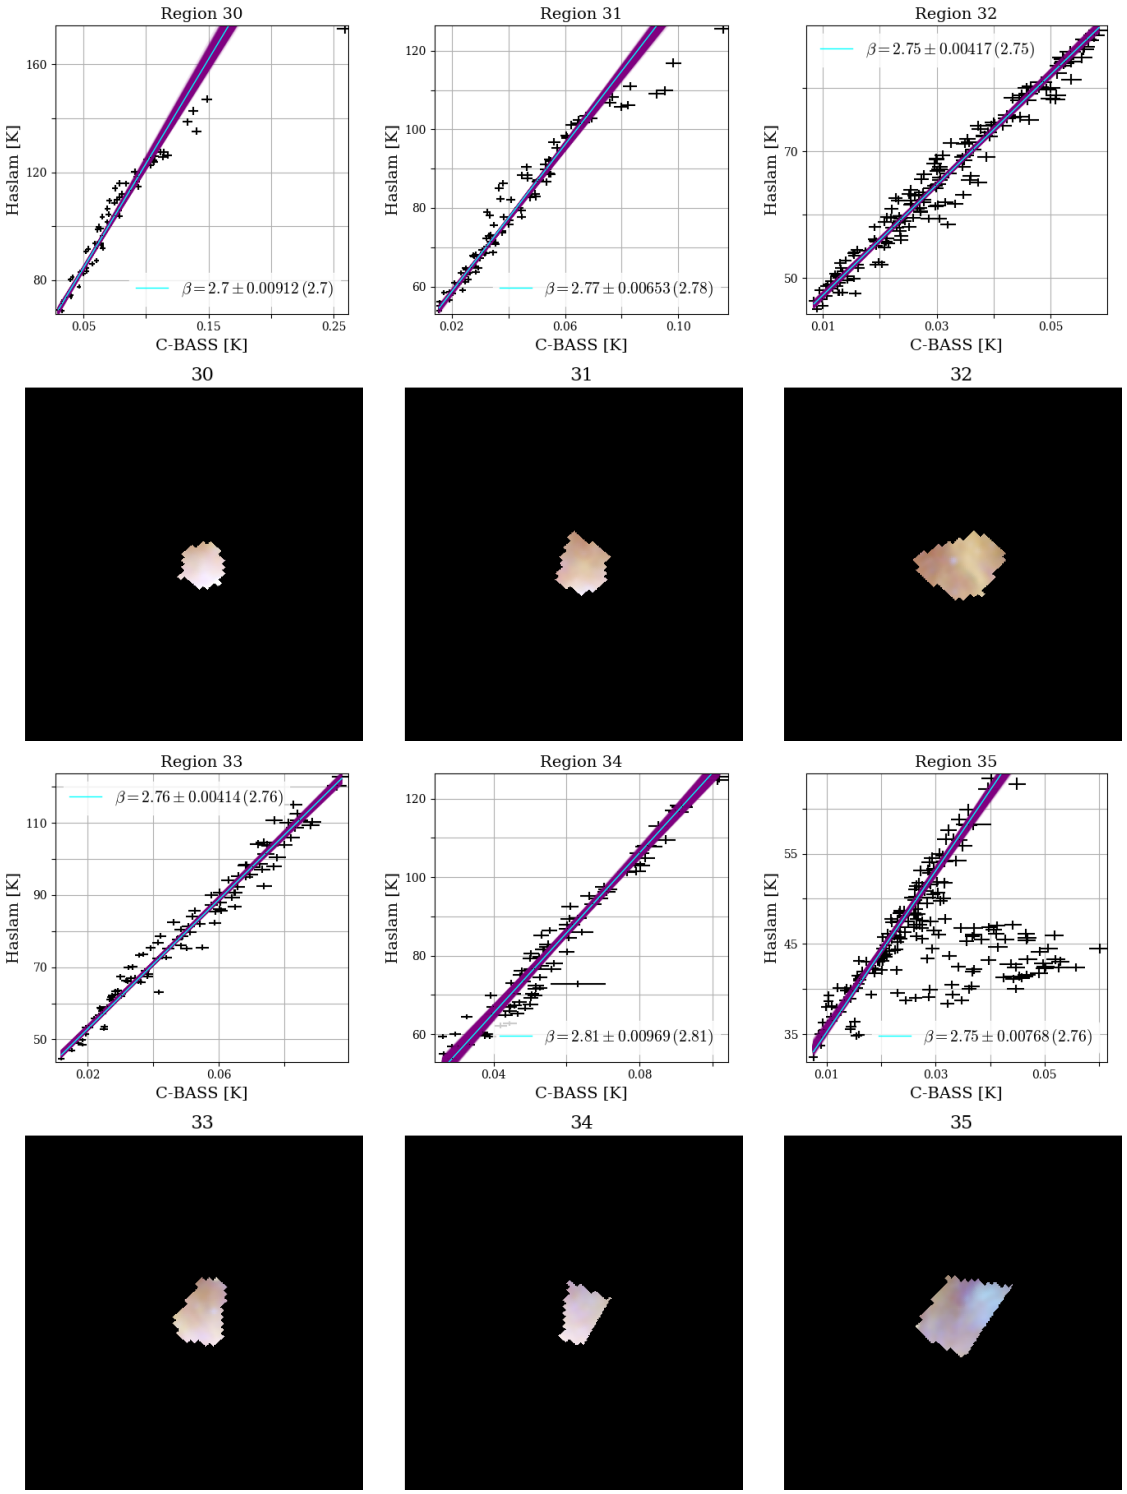

Figure 5.15: T-T plots between the C-BASS and Haslam maps from regions 30–35. A thinned subset of the converged MCMC chains are plotted. The *purple* lines are component I. The mean and standard deviation of the posterior distribution of the spectral index are labelled along with the maximum posterior value in brackets. Beneath the T-T plots are gnomonic projections of the region in the 3-colour map.

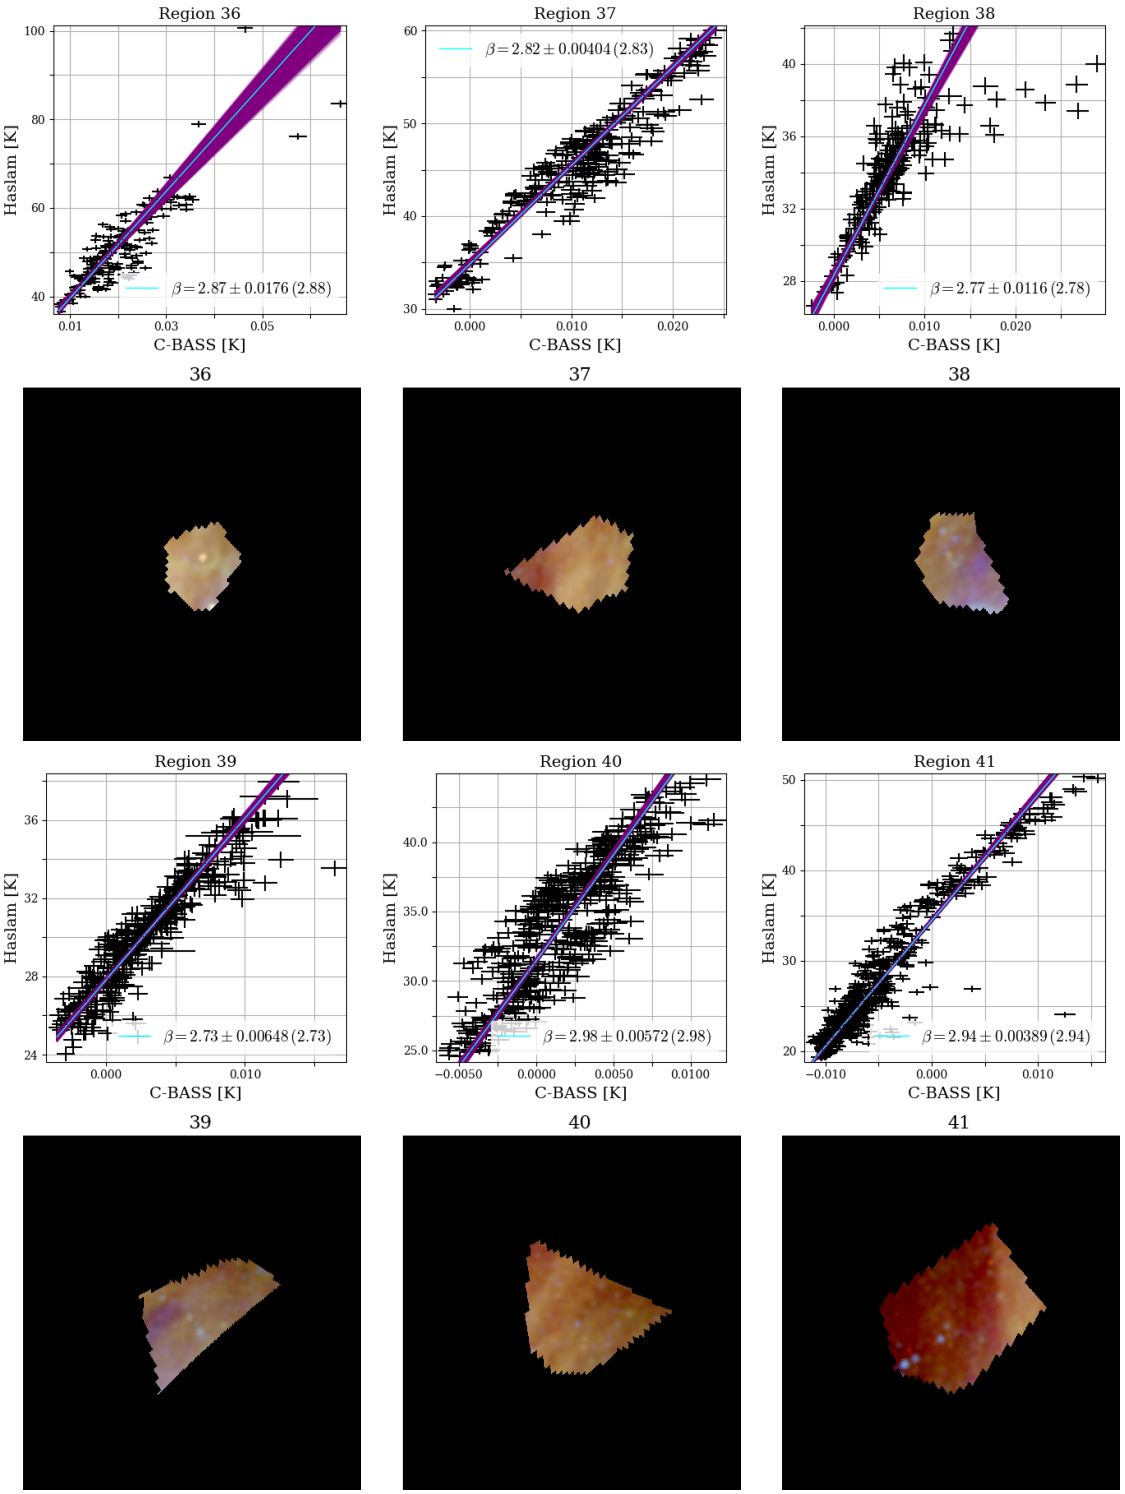

Figure 5.16: T-T plots between the C-BASS and Haslam maps from regions 36–41. A thinned subset of the converged MCMC chains are plotted. The *purple* lines are component I. The mean and standard deviation of the posterior distribution of the spectral index are labelled along with the maximum posterior value in brackets. Beneath the T-T plots are gnomonic projections of the region in the 3-colour map.

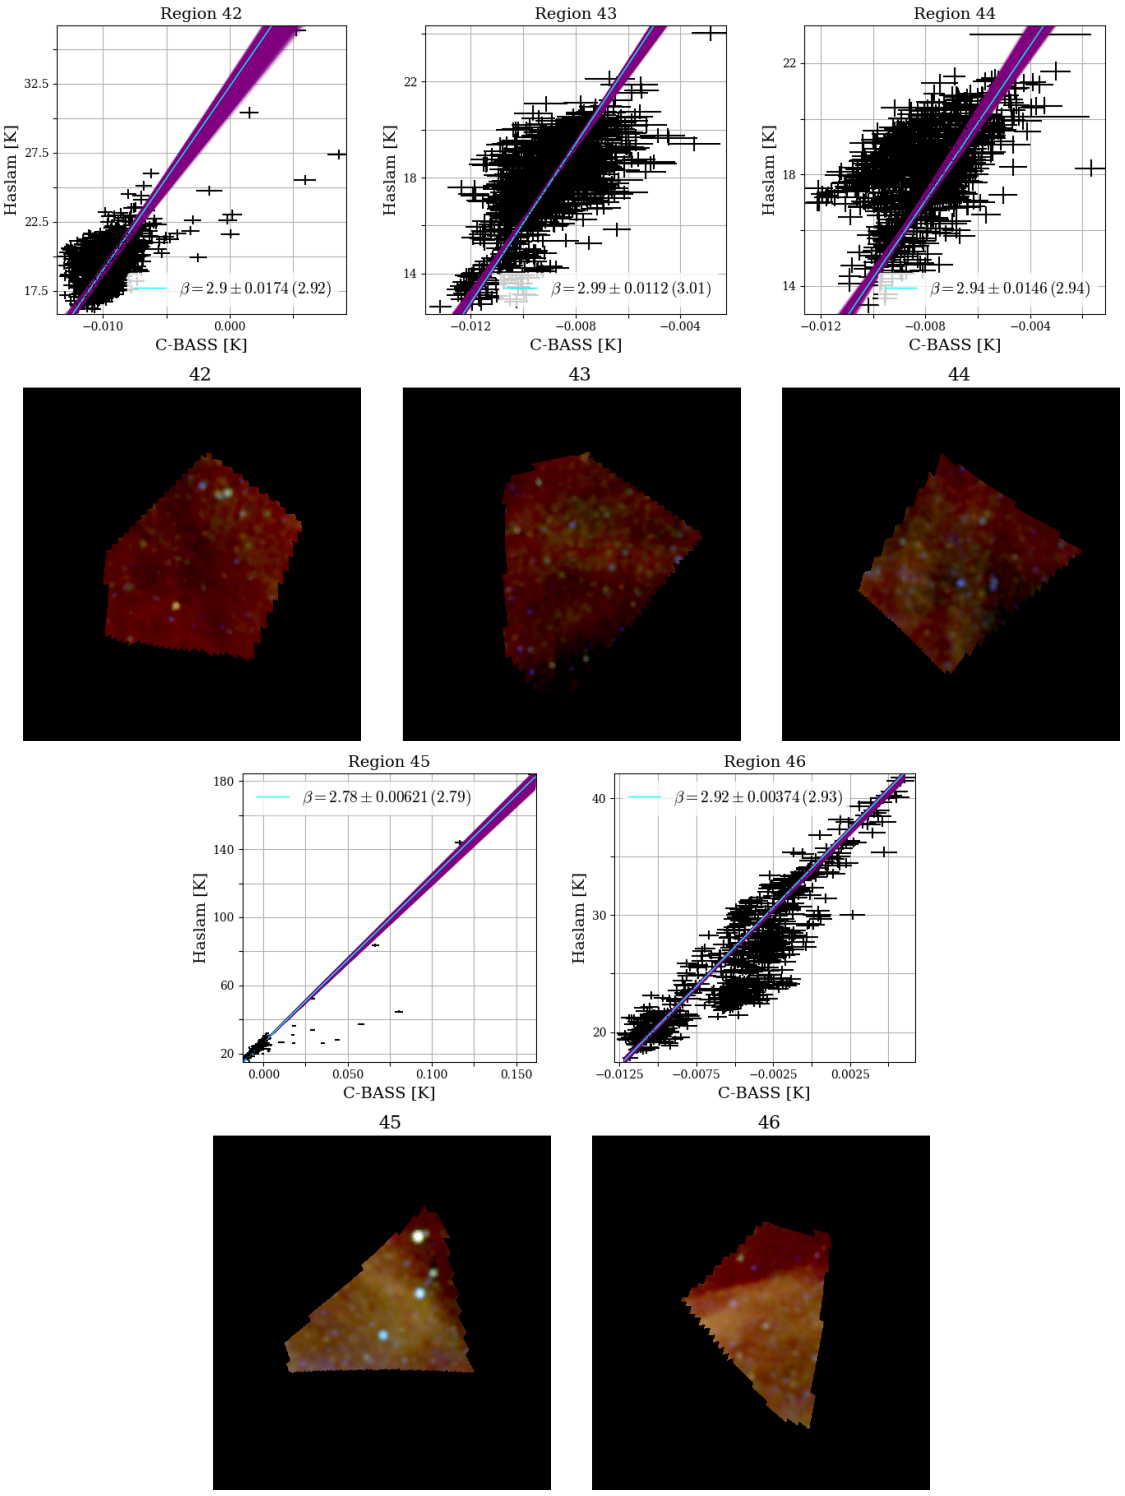

Figure 5.17: T-T plots between the C-BASS and Haslam maps from regions 42–46. A thinned subset of the converged MCMC chains are plotted. The *purple* lines are component I. The mean and standard deviation of the posterior distribution of the spectral index are labelled along with the maximum posterior value in brackets. Beneath the T-T plots are gnomonic projections of the region in the 3-colour map.

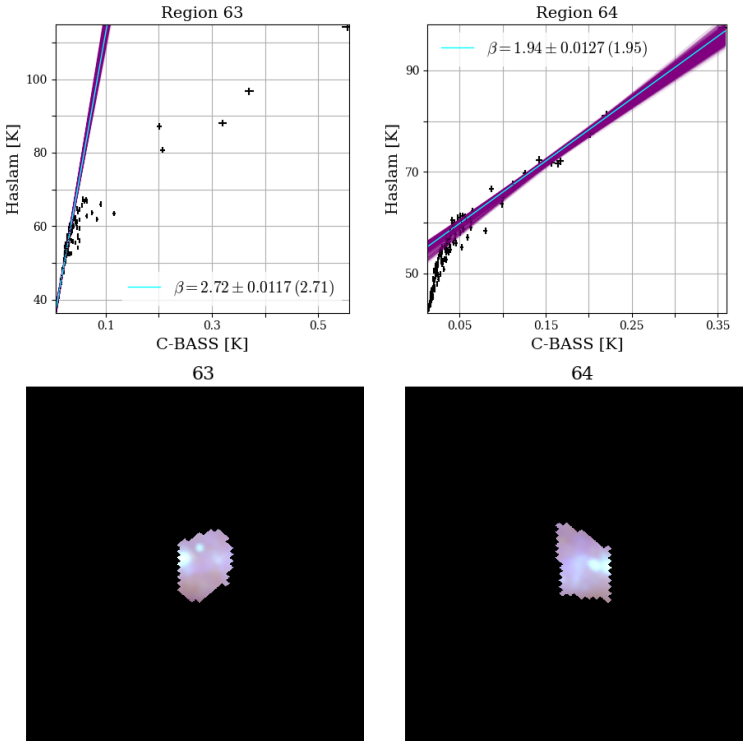

Figure 5.18: T-T plots between the C-BASS and Haslam maps from regions 63 and 64. A thinned subset of the converged MCMC chains are plotted. The *purple* lines are component I. The mean and standard deviation of the posterior distribution of the spectral index are labelled along with the maximum posterior value in brackets. Beneath the T-T plots are gnomonic projections of the region in the 3-colour map.

plane. The spectral indices in many regions of the simulation were between 2–3 and this is also the case for the real data. The estimates of β in each region are not estimates of either the synchrotron or free-free spectral indices but rather a mixture of the two.

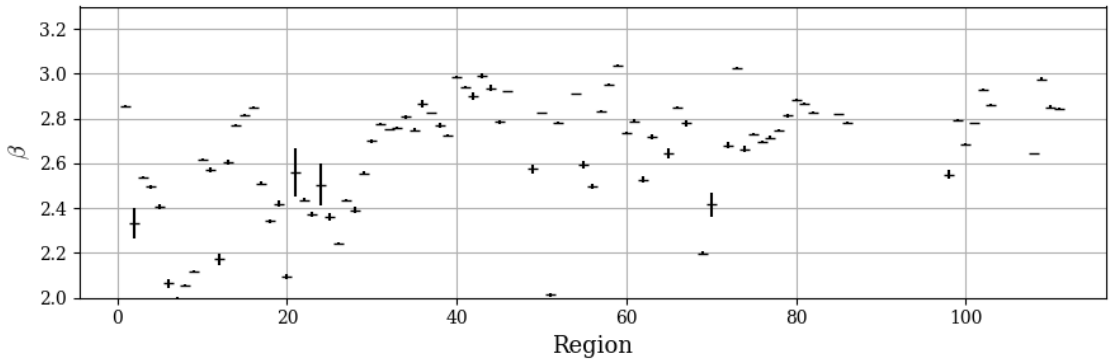

Figure 5.19: Scatter plot of the estimated spectral index in each region between the real C-BASS and Haslam maps.

5.2.7 Discussion of results from real data

The simulations showed that the spectral index of synchrotron and free-free emission could not be unambiguously measured using this method, even when dividing the sky into the smallest regions that meet a required signal threshold and allowing the pixels to be dominated by one of several components.

The real data show a broad scatter of spectral indices between pure free-free and $\beta \sim 3$ synchrotron emission. This tells us little about the true spatial variation of the synchrotron spectral index across the sky and instead we need to use maps at more frequencies to fit a more complicated model pixel-by-pixel, allowing for pixels to contain a mixture of emission mechanisms across a wide frequency range and not dominated by a single component.

Commander is the Gibbs sampling code used by the Planck Collaboration to fit a parametric model of the sky to a range of maps at different frequencies, pixel-by-pixel. Commander can also fit for the monopoles, dipoles and gain factors of the input maps. Commander has been run on *WMAP*, *Planck* and Haslam maps (Planck Collaboration et al., 2016b). Including C-BASS gives another free parameter in the fit, which allows for the synchrotron spectral index to be determined across the sky.

We ran Commander with an early C-BASS map. The signal-to-noise ratio in the C-BASS map was high enough that artefacts in the map clearly impacted on the results and showed that we had to get a better grasp of the systematics. All of the work in Chapter 4 was to ensure that the C-BASS map is good enough for use in analysis with Commander. We are running Commander with the latest C-BASS map to obtain more accurate estimates of the spectral parameters of low-frequency foregrounds. This work is ongoing and is not presented here.

The T-T plot method still has use when a region of the sky can be tightly defined and is known to be dominated by a single component over a range of frequencies. In particular for C-BASS we can use this to cross-check the calibration against other maps. In the following section we look at Barnard's Loop, known to be free-free dominated and use T-T plots to cross-check the C-BASS calibration.

5.2.8 Barnard's Loop

In this section we use the extended T-T plot method developed above to check that the C-BASS I map is on the same temperature scale as the Haslam, *WMAP* K-band and *Planck* 30 GHz maps by measuring the spectral index of Barnard's Loop.

Barnard's Loop is an HII region in Orion and at frequencies between several and tens of GHz its emission is dominated by optically thin free-free emission (Barnard, 1894; Reich, 1978). Free-free emission has a typical temperature spectral index of 2.12 ± 0.02 (Draine, 2011). In Barnard's Loop, the spectral index between 85, 240 and 408 MHz was measured by Haslam & Salter (1971) to be 2.05 ± 0.3 and Heiles et al. (2000) have confirmed that the spectral index is consistent with 2.1 up to 2.3 GHz.

The *WMAP* seven- and nine-year data releases included MCMC fits of several parametric foreground models to their data. In all of their analyses the spectral shape of the free-free emission was fixed but its amplitude was allowed to vary. They tested several free-free frequency spectra including a power-law spectrum with $\beta_{\text{free-free}} = 2.16$ and a spectrum from Oster (1961) with fixed electron temperature $T_e = 8000\,\text{K}$ where spectral index steepens from 2.09 at 0.408 GHz to 2.14 at K-band and to 2.17 at W-band (Gold et al., 2011; Bennett et al., 2013). They find that all of their models gave good fits to the data.

To confirm the calibration of an earlier version of the C-BASS maps, Irfan et al. (2015) estimated the spectral index between the C-BASS, Haslam, *WMAP* K-band and the 1.4 GHz Reich maps in the rectangular region defined by Galactic longitude $208 < l < 213$ and latitude $-13.5 < b < -12$ using T-T plots. Since Irfan et al. (2015) we have improved the C-BASS pipeline. In particular the 1.2 Hz removal, ground removal, calibration, sun-flagging, RFI flagging, map making and noise model have all been improved. Irfan et al. (2015) scaled the C-BASS map to the full-beam temperature scale and no attempt was made to deconvolve the telescope beam. Because Barnard's Loop does not fill the telescope beam and neither is it a point source, neither the full-beam nor main-beam temperature scales are strictly appropriate.

The top part of Table 5.4 contains typical free-free spectral indices and a summary of measured values of the spectral index in Barnard's Loop. From these results we expect the emission in Barnard's Loop to obey a power law with a spectral index of ~ 2.1–2.2 between 0.408 and 30 GHz.

We use T-T plots to estimate the spectral index of Barnard's Loop in two regions. Firstly we use the same region as Irfan et al. (2015), hereafter called region 1. Secondly we use a larger region that traces the length of Barnard's Loop, hereafter called region 2. A gnomonic projection of Barnard's Loop in the 3-colour map is shown in Figure 5.20. Region 1 is outlined in *solid red* and region 2 is outlined in *dotted white*.

We estimate the spectral index between the C-BASS map and the Haslam, *WMAP* K-band and *Planck* 30 GHz maps. Projections of Barnard's loop from each map are shown in Figure 5.21, regions 1 and 2 are again outlined in *solid red* and *dotted white*

Table 5.4: Summary of the typical free-free spectral index across the sky and measured values in Barnard's Loop from both this work and the work of others.

Result	β	Comment
Typical free-free spectral indices		
Draine (2011)	2.12 ± 0.02	Typical free-free spectrum
Bennett et al. (2013)	2.09–2.17	0.408–22–90 GHz
Spectral index of Barnard's Loop		
Irfan et al. (2015)	2.16 ± 0.05	C-BASS & Haslam (region 1)
–"–	2.13 ± 0.11	C-BASS & Reich (region 1)
–"–	2.15 ± 0.05	C-BASS & *WMAP* K-band (region 1)
Haslam & Salter (1971)	2.05 ± 0.30	85.5 MHz, 240 MHz, 408 MHz
Heiles et al. (2000)	~ 2.1	Up to 2.3 GHz
This work, region 1		
C-BASS & Haslam	1.99 ± 0.03	
C-BASS & *WMAP* K-band	2.25 ± 0.02	
C-BASS & *Planck* 30 GHz	2.23 ± 0.02	
This work, region 2		
C-BASS & *WMAP* K-band	2.15 ± 0.01	
C-BASS & *Planck* 30 GHz	2.08 ± 0.01	

Table 5.5: The effective frequencies and typical sensitivities of the C-BASS, *WMAP* K-band and *Planck* 30 GHz $N_{\rm side} = 64$ maps that have been smoothed to 1°. The effective frequencies of the C-BASS and *WMAP* maps are for emission with temperature spectral index of 3, the *Planck* map has had the appropriate colour correction.

Map	Frequency [GHz]	Typical σ_I	Unit
Haslam	0.408	0.36	K
C-BASS	4.79	0.30	mK
WMAP K-band	22.5	7.7	μK
Planck 30 GHz	28.4	1.1	μK

respectively. The centre frequencies and typical sensitivities of the maps are shown in Table 5.3.

The C-BASS map was the same as in Section 5.2.8 except that we used the effective frequency for emission with a spectral index of 2.1 when converting to brightness temperature.

We used the $N_{\rm side} = 512$ *WMAP* K-band map from the 9-year data release that was smoothed to 1° (Bennett et al., 2013). We assumed that the pixel noise is Gaussian and spatially uncorrelated and calculated it using $\sigma = \sigma_0/\sqrt{N_{\rm obs}}$ where $\sigma_0 = 1.429\,{\rm mK_{CMB}}$.[3] We scaled the maps from thermodynamic temperature units to

[3] https://lambda.gsfc.nasa.gov/product/map/current/

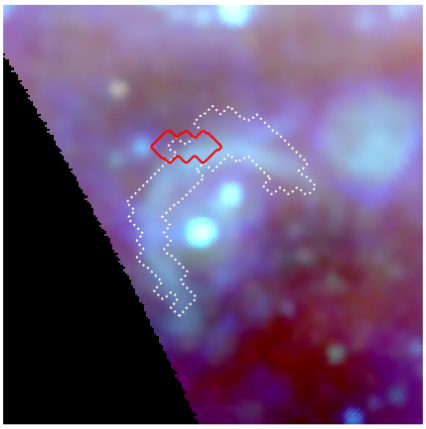

Figure 5.20: Gnomonic projection of Barnard's Loop in the 3-colour map, the image is 30° across. The *solid red line* outlines region 1 and the *dotted white line* outlines region 2.

brightness temperature, downgraded the map to $N_{\rm side} = 64$. After downgrading to lower $N_{\rm side}$ we scaled the white noise variance by the ratio of the pixel solid angles (Equation 4.18). As in Collaboration Planck et al. (2011) we added a 3% gain error to the downgraded *WMAP* K-band map. We used the effective frequency for emission with a spectral index of 2.1.

We used the *Planck* 30 GHz map from the 2015 data release (Planck Collaboration et al., 2016a), converted it from thermodynamic temperature units to brightness temperature, approximately smoothed the map to 1° and downgraded to $N_{\rm side} = 64$. The map was smoothed to a Gaussian with FWHM = 1° by assuming that the telescope beam was equal to the equivalent Gaussian, with effective FWHM 32.4″ (Planck Collaboration et al., 2015f). The smoothing kernel was then a Gaussian with $\theta_{\rm FWHM}^{\rm smoothing} = \sqrt{\theta_{\rm FWHM}^{\rm target\ 2} - \theta_{\rm FWHM}^{\rm equivalent\ 2}} = 50.3'$. The equivalent Gaussian is not a particularly well defined quantity but the exact shape of telescope beam should not make much difference to very diffuse emission. The exact beam shape is most important when considering radio point sources. We calculated $\eta_{\rm smooth}$ using Equation 4.17 and estimated $\eta_{\rm dg}$ from simulations. After downgrading the *Planck* 30 GHz map to $N_{\rm side} = 64$ we added $0.19\,\mu{\rm K_{CMB}}$ and 0.35% of the I map to each pixel in the sensitivity map to account for systematic effects (Planck Collaboration et al., 2016c). We applied the colour correction for emission with spectral index 2.1, which is listed in Appendix A.

We used the same fitting technique as in the previous section, including components I and III in the mixture model (i.e. a single straight line with uninformative

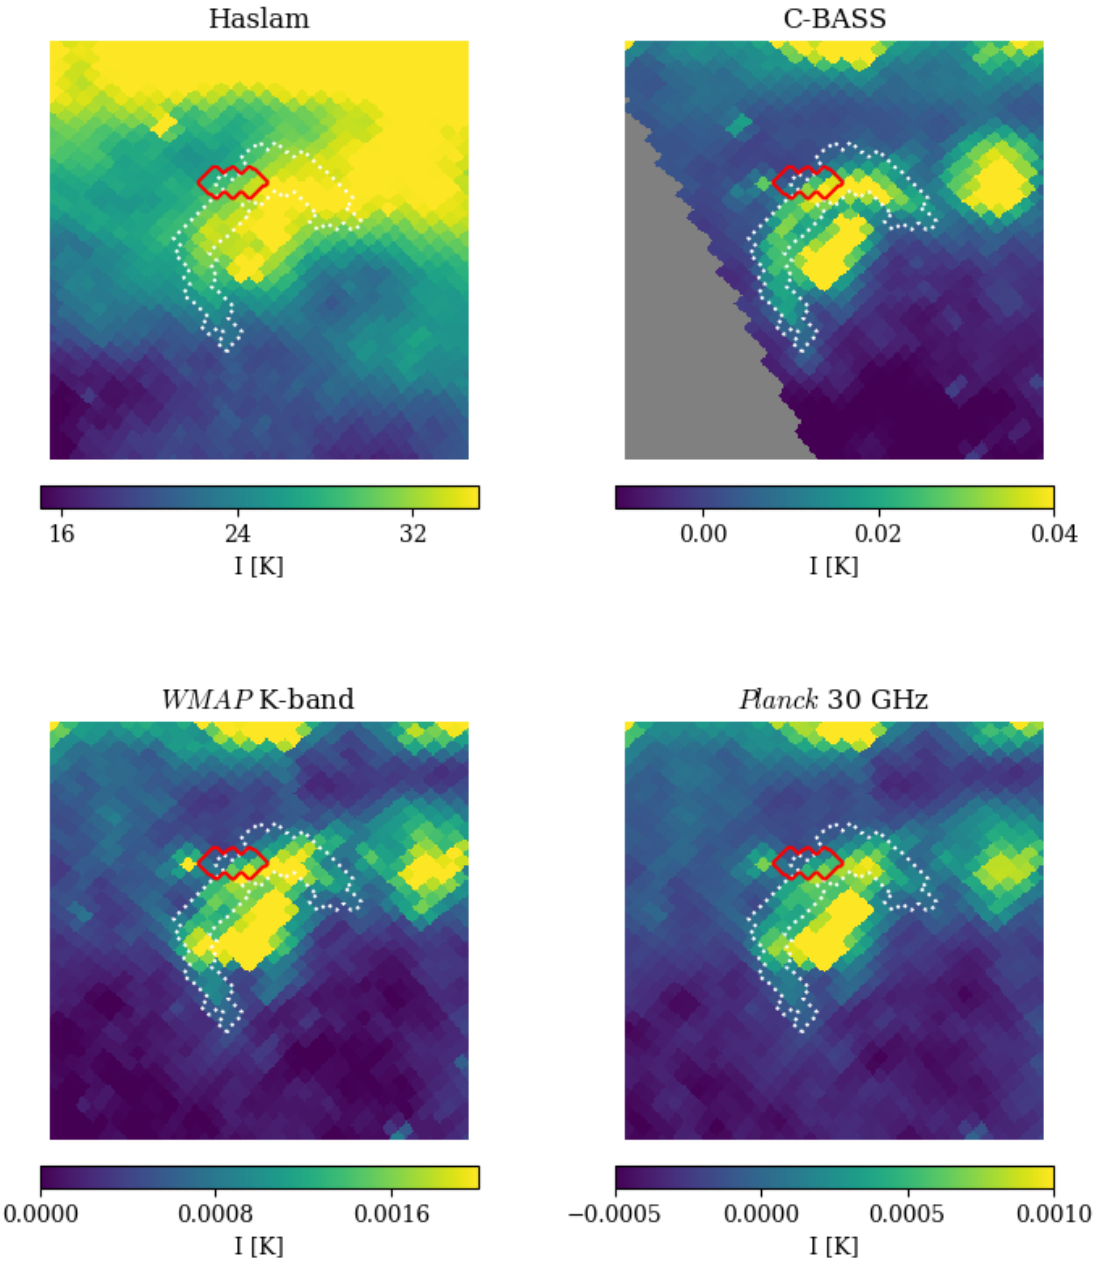

Figure 5.21: Gnomonic projections of Barnard's Loop, the images are 30° across. Each map has been smoothed to 1° and is in brightness temperature. The *solid red line* outlines region 1 and *dotted white line* outlines region 2.

prior on its slope and a population of outliers).

The T-T plots for region 1 are shown in Figure 5.22. The estimated values of β between C-BASS and the other three maps are listed in Table 5.4. The spectral index between C-BASS/Haslam is shallower than the spectral index between C-BASS/*WMAP* and C-BASS/*Planck*. The free-free spectral index is expected to flatten off at lower frequencies and the measured values are consistent with that. The spectral index between C-BASS/*WMAP* and C-BASS/*Planck* is steeper than expected. This could be caused by a gain error in the C-BASS map or a mixture of foreground components and emission from extra-galactic sources.

For small gain errors δG in the C-BASS map, the bias introduced to the estimated spectral index is $\delta\beta \simeq \delta G/\log(\nu_1/\nu_2)$, which corresponds to $\delta\beta \simeq 0.005$. A 5% gain error will bias the spectral indices by $\delta\beta \simeq 0.025$. This systematic error is not included in the uncertainty on the spectral index. Multiplying the C-BASS map by a factor slightly less than 1 would bring the spectral indices closer to the expected values.

The T-T plots for region 2 are shown in Figure 5.22. We have not included the T-T plot between C-BASS and Haslam here as in this larger region there is significant synchrotron emission at 0.408 GHz. The estimated values of β between C-BASS and *WMAP* and *Planck* in this larger region are also listed in Table 5.4.

There are pixels in the C-BASS/*WMAP* T-T plot, and to a lesser extent the C-BASS/*Planck* T-T plot, that do not obey a linear relationship and are too bright at frequencies above 5 GHz. This is likely caused by AME. There are pixels in the 3-colour map slightly more blue than typical cyan coloured free-free at both the top of region 2 and to the left of Orion A, two thirds of the way down Barnard's Loop.In this larger region, not all of the pixels appear to be dominated by a single emission mechanism with constant spectral parameters in all pixels.

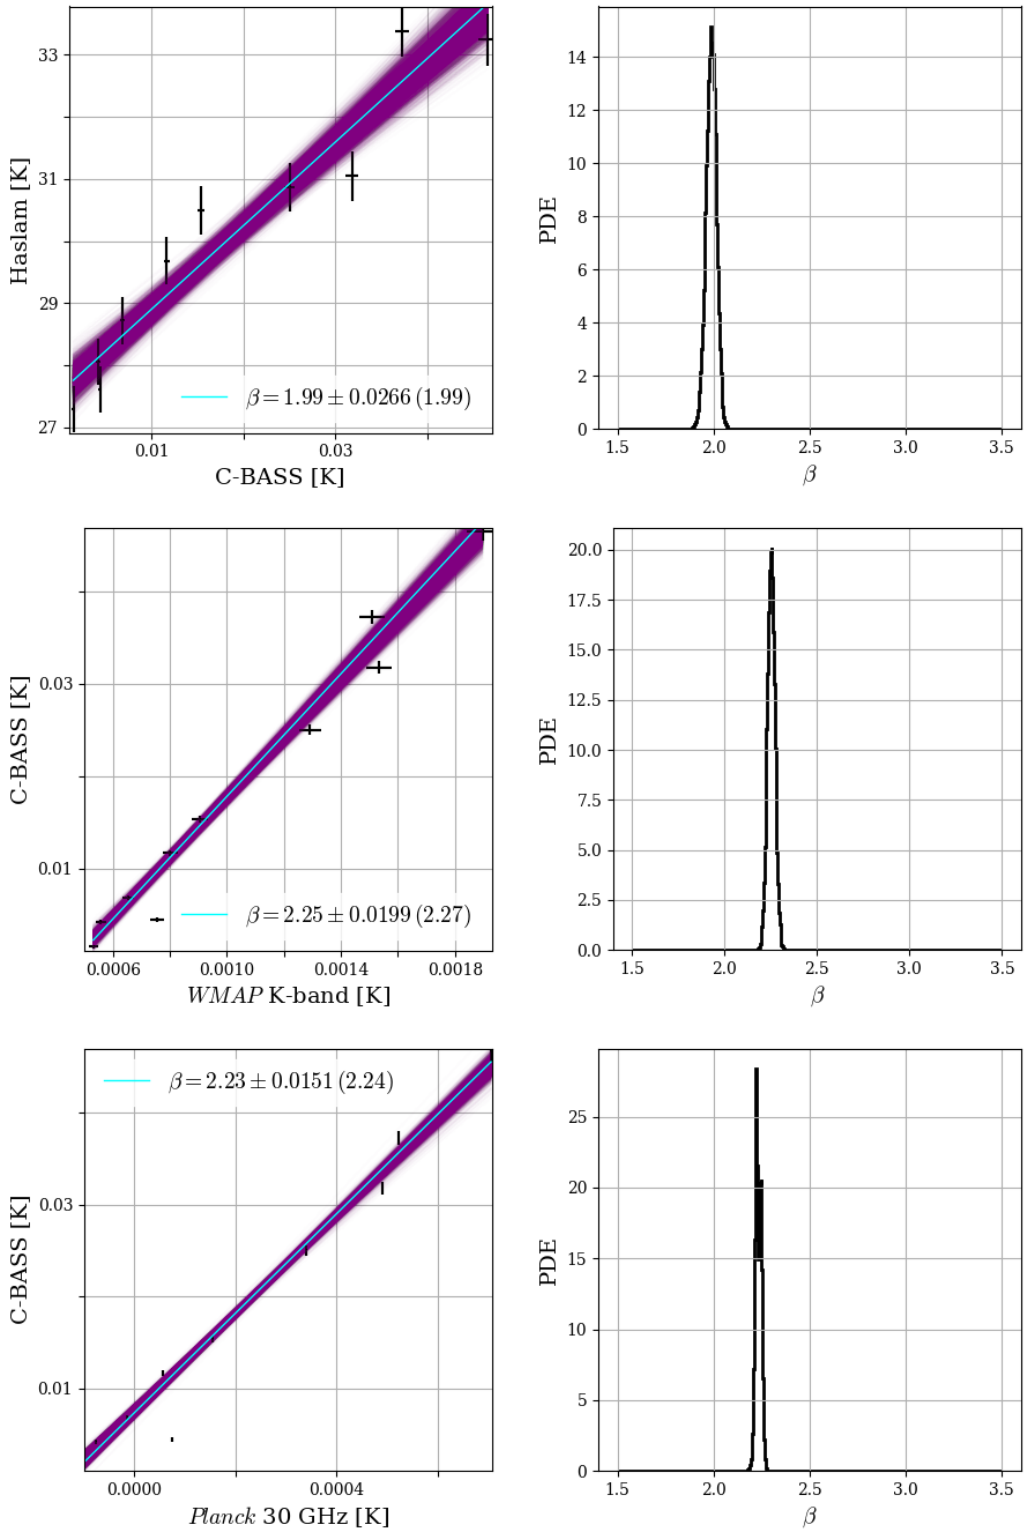

Figure 5.22: T-T plots of Barnard's loop in region 1 between C-BASS and Haslam, *WMAP* K-band and *Planck* 30 GHz. The *purple* lines are a thinned sample of the converged MCMC chains and the *cyan* line is the maximum posterior estimate. The mean and standard deviation of the marginalised PDE of the spectral index is labelled on each plot with the maximum posterior spectral index in brackets.

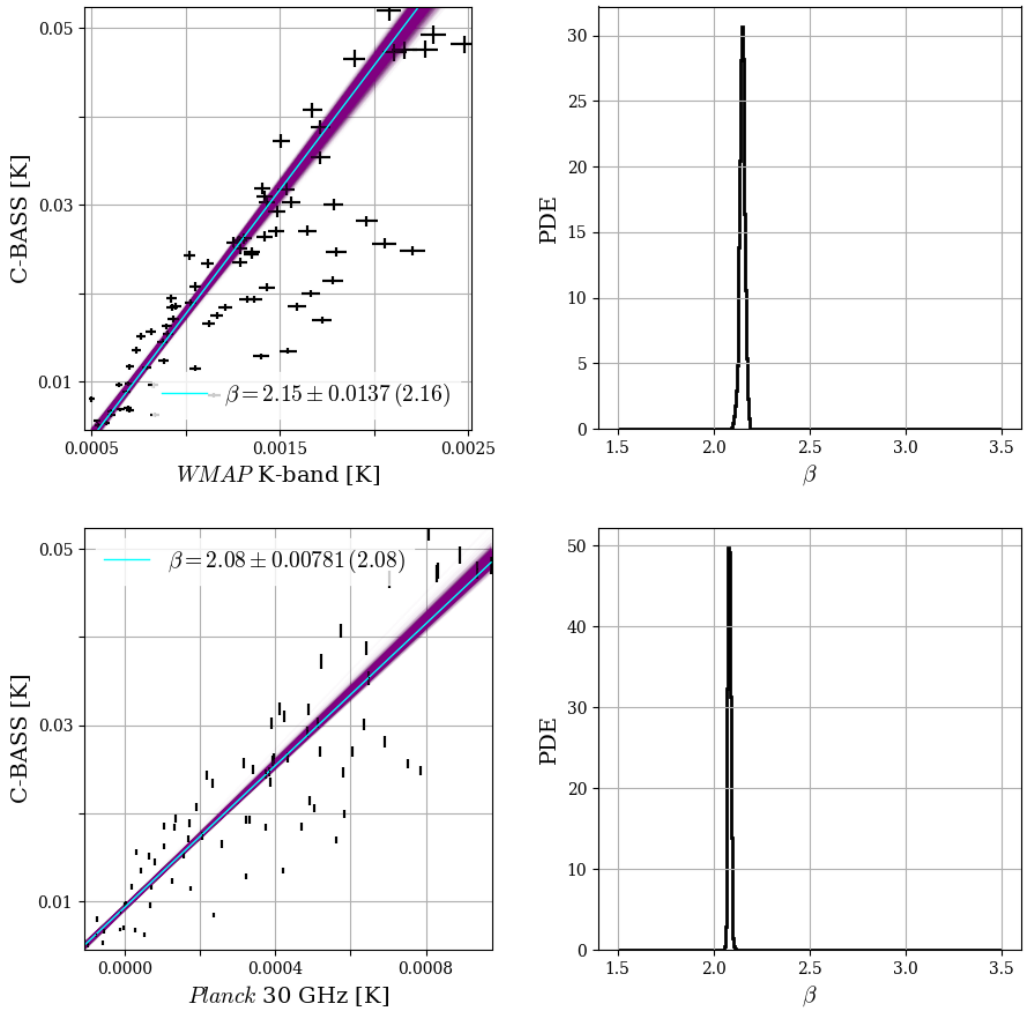

Figure 5.23: T-T plots of Barnard's loop in region 2 between C-BASS and *WMAP* K-band and *Planck* 30 GHz. The *purple* lines are a thinned sample of the converged MCMC chains and the *cyan* line is the maximum posterior estimate. The mean and standard deviation of the marginalised PDE of the spectral index is labelled on each plot with the maximum posterior spectral index in brackets.

5.3 Discussion

In this chapter we have analysed the frequency spectrum of low-frequency diffuse Galactic emission in total intensity.

We have demonstrated the complexity of the emission by creating a 3-colour map of the sky from the Haslam, C-BASS and *WMAP* maps.

We used the T-T plot technique to estimate the spectral index between the C-BASS and Haslam maps. If the emission in a region obeys a power-law frequency spectrum then a plot of the temperature at one frequency versus the temperature at

another will obey a linear relationship. T-T plots are typically used on regions that are known to be dominated by a single emission mechanism and we developed the standard method to try and measure the synchrotron spectral index blindly across the whole sky.

The signal-to-noise on both axes of a T-T plot are often similar and so estimates of the slope and intercept of the line found using the method of least-squares is biased. We use the posterior distribution first identified by Gull (1989), which provides unbiased estimates of the slope and intercept of a line in this regime. We extended the standard T-T plot model with a mixture that allows pixels to belong to one of several components. Specifically we included two straight lines, one with an informative prior on its slope (and therefore spectral index) and one with an uninformative prior on its slope, and a population of outliers.

The T-T plot method requires the sky to be divided into regions over which the spectral parameters and map monopoles are constant. We used the Voronoi binning technique to divide the sky into the smallest possible regions that also met a minimum signal level.

From simulations we found that the measured spectral index between 0.408 and 5 GHz is likely caused by a mixture of synchrotron and free-free emission. We also found that including a second line in the mixture model with an informative prior on its spectral index did not improve the estimates of the synchrotron spectral index.

We used this method on the real C-BASS and Haslam maps, only including a single straight line component and a population of outliers in our mixture model. We found that the spectral index varied across the sky between 2–3 but could not determine whether that variation was due to the synchrotron spectral index changing or the ratio of free-free to synchrotron emission changing. At 5 GHz the foregrounds across most of the sky are too complicated to easily measure the spectral parameters of a single component using T-T plots and so we have to use more sophisticated techniques, for example pixel-based parametric model fitting.

T-T plots are typically used on regions that are known to be dominated by a single emission mechanism and so we used them on Barnard's Loop to verify the calibration of the C-BASS I map. In the smaller region 1 we found that the C-BASS data point may be higher than expected. In the larger region 2 we found that there was contamination from other foreground components and even this relatively small region has a more complicated spectrum than just being dominated by free-free emission.

In the following section we analyse the frequency spectrum of polarized emission.

Chapter 6

Constraints on the Polarized Diffuse Galactic Synchrotron Spectral Index

In this chapter we present measurements of the polarized synchrotron spectral index between 5 and 30 GHz in $N_{\rm side} = 64$ maps.

Estimating the spectral index of polarized synchrotron emission is more straight forward than the total intensity case. Firstly, unlike maps of Stokes I, maps of Stokes Q and U (and therefore P) are not usually subject to large global offsets. Secondly, the polarized diffuse Galactic emission is far simpler than the total intensity emission. Synchrotron and thermal dust are the only known major components of polarized emission and thermal dust is negligible at C-BASS frequencies. AME may be polarized but there are currently only upper limits on its polarization fraction.

In Section 6.1 we describe the C-BASS and *Planck* 30 GHz polarization maps that we use in this analysis. We verify the calibration of the C-BASS P map by measuring the polarization fraction of TauA. We also calculate the spectral index very simply using Equation 5.2. This simple calculation does not account for the probability distributions of P maps and does not provide an estimate for the uncertainty of the estimate.

In Section 6.2 we fit a model spectrum, pixel-by-pixel, to the C-BASS and *Planck* 30 GHz P maps, modelling measurements of P as Rician random variables. We test the method on simulated data to show that it recovers unbiased estimates of the spectral index even in low signal-to-noise regions of the sky. We then use the method on real data.

6.1 Polarization datasets and a rudimentary estimate of the spectral index between 5 and 30 GHz

In this section we describe the C-BASS and *Planck* 30 GHz P maps, test the calibration of the C-BASS map by measuring the polarization fraction of TauA and make a rudimentary estimate of the spectral index between the two maps.

6.1.1 Datasets

We used the C-BASS P and σ_P maps that have been smoothed to 1° and downgraded to $N_{\rm side} = 64$, as described in Section 4.5.2. The maps were scaled to brightness temperature using an effective frequency of $\nu_{\rm eff} = 4.78$ GHz, which is appropriate for emission with spectral index of $\beta = 3$.

We used the *Planck* 30 GHz $N_{\rm side} = 256$ maps from the 2015 data release (Planck Collaboration et al., 2016a) and converted them from thermodynamic temperature scale to brightness temperature using the colour correction for emission with a spectral index of 3.

We smoothed the Q and U maps to a Gaussian with FWHM $\theta_{\rm FWHM}^{\rm target} = 60'$ by assuming that the telescope beam was equal to the equivalent Gaussian, with effective FWHM $\theta_{\rm FWHM}^{\rm equivalent} = 32.4'$ (Planck Collaboration et al., 2015f). The smoothing kernel was then a Gaussian with $\theta_{\rm FWHM}^{\rm smoothing} = \sqrt{\theta_{\rm FWHM}^{\rm target\ 2} - \theta_{\rm FWHM}^{\rm equivalent\ 2}} = 50.3'$. The variance maps then had to be scaled by the factor $\eta_{\rm smooth} = 0.033$, calculated using Equation 4.17.

We downgraded the smoothed Q, U, σ_Q and σ_U maps to $N_{\rm side} = 64$ using the HEALPix UD_GRADE routine. Simulations showed that an $N_{\rm side} = 256$ variance map that has been smoothed by a Gaussian with FWHM $\theta_{\rm FWHM} = 50.3'$ and scaled by a factor of $\eta_{\rm smooth} = 0.033$ requires a further scaling of $\eta_{\rm dg} = 0.63$ after downgrading to $N_{\rm side} = 64$.

From the smoothed and downgraded maps we constructed $P = \sqrt{Q^2 + U^2}$ and approximated $\sigma_P = (\sigma_Q + \sigma_U)/2$. To the downgraded *Planck* σ_P map we added 0.3 $\mu{\rm K_{CMB}}$ of extra noise to account for systematics and 0.35% of the P map, in quadrature, to account for calibration uncertainty (Planck Collaboration et al., 2016c).

The C-BASS and *Planck* 30 GHz P and σ_P maps are shown in Figure 6.1. The effective frequencies and typical sensitivities of both maps are shown in Table 6.1.

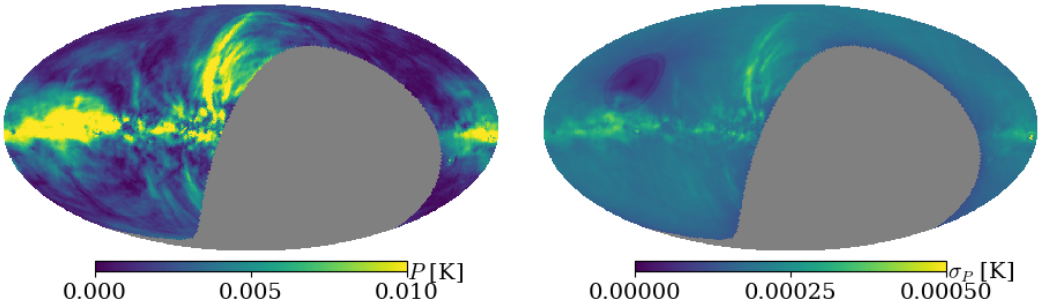

(a) C-BASS P map

(b) C-BASS σ_P map

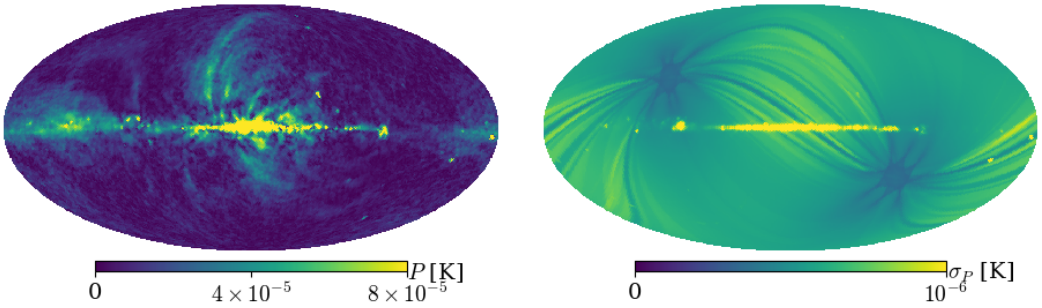

(c) *Planck* 30 GHz P map

(d) *Planck* 30 GHz σ_P map

Figure 6.1: C-BASS and *Planck* 30 GHz P and σ_P maps downgraded to $N_{\rm side} = 64$. All maps are in brightness temperature and on linear colour scales.

We have not used the *WMAP* polarization maps in this work. The *WMAP* Q and U maps have a small number of poorly constrained modes at low ℓ that introduce large scale gradients and offsets into the maps and as such P can not be directly calculated (Jarosik et al., 2011; Bennett et al., 2013).

6.1.2 Checking the calibration of the C-BASS P map

The astronomical calibration subroutine in the pipeline calculates the amplitude of the noise diode signal in I TOD from observations of the total intensity emission of

Table 6.1: A summary of the effective frequencies and typical sensitivities (in an $N_{\rm side} = 64$ pixel) of the C-BASS and *Planck* 30 GHz maps used in this analysis.

Map	Effective frequency [GHz]	Typicall $\sigma_P(1^\circ)$
C-BASS	4.78	0.22 $\mathrm{mK_{RJ}}$
Planck 30 GHz	28.4	0.59 $\mu\mathrm{K_{RJ}}$

Table 6.2: Measurements of the polarization fraction of TauA at different frequencies and angular resolutions.

Frequency [GHz]	Resolution	Polarization fraction	Reference
90.0	5.0′	$8.8 \pm 0.2\%$	Aumont et al. (2010)
90.0	10.0′	$7.7 \pm 0.2\%$	–”–
92.9	13.2′	$7.13 \pm 0.43\%$	Weiland et al. (2011)
60.2	21.0′	$7.00 \pm 0.19\%$	–”–
40.4	30.6′	$6.97 \pm 0.27\%$	–”–
32.8	39.6′	$6.94 \pm 0.18\%$	–”–
22.5	52.8′	$7.08 \pm 0.25\%$	–”–
5.0	6.1′	$5.6 \pm 0.1\%$	Sastry, Pauliny-Toth & Kellerman (1967)
5.0	60.0′	$4.7 \pm 0.1\%$	This work
2.65	7.3′	$3.8 \pm \%0.5$	Gardner & Davies (1966)
1.41	14.1′	$1.6 \pm \%0.1$	–”–
0.96	19.0′	$0.9 \pm \%0.2$	–”–

TauA, which is well known. This calibration is applied to the polarization data via the internal temperature-stabilised load. This is the first attempt at the calibration and to verify that it has worked we extract the polarization fraction of TauA from the C-BASS maps. In this section we also briefly discuss future and ongoing work to check and improve on the calibration of C-BASS North using C-BASS South data.

TauA is the primary astronomical calibrator of the C-BASS maps and it is well characterised in total intensity, for examples see Baars et al. (1977) or Hafez et al. (2008). We are not aware of any reliable measurements of the polarization fraction of TauA at 5 GHz and 1° resolution. Table 6.2 lists some measurements of the polarization fraction of TauA at other frequencies and resolutions.

The polarization fraction of TauA increases with frequency, for example Gardner & Davies (1966) find that TauA has a polarization fraction of $0.9 \pm 0.2\%$ at 0.96 GHz and $3.8 \pm 0.5\%$ at 2.65 GHz. The polarization fraction also typically increases with the angular resolution. For example, at 90 GHz Aumont et al. (2010) find a polarization fraction of $7.7 \pm 0.2\%$ at 10.0′ resolution that increases to $8.8 \pm 0.2\%$ at 5.0′ resolution. At 5 GHz Sastry, Pauliny-Toth & Kellerman (1967) report a polarization fraction of $5.6 \pm 0.1\%$, but this is with an angular resolution approximately ten times greater than that of C-BASS. We therefore expect the polarization fraction of TauA in the C-BASS map to be close to, but smaller than, the value of Sastry, Pauliny-Toth & Kellerman (1967).

We estimated the polarization fraction of TauA using the C-BASS maps smoothed to 1° and at $N_{\rm side} = 1024$. We divided the P and I values in the pixel closest to the centre of TauA by one another. This very simple process gives a polarization fraction

of $4.7 \pm 0.1\%$. This is consistent with the expectation from measurements at other frequencies and resolutions.

The reported error in the polarization fraction was calculated using

$$\sigma_f = f\sqrt{\left(\frac{\sigma_P}{P}\right)^2 + \left(\frac{\sigma_I}{I}\right)^2}, \tag{6.1}$$

where f is the polarization fraction. We have approximated the noise on P as Gaussian with standard deviation σ_P. We have not taken account of the Rician distribution of measurements of P (see Equation 6.3 in Section 6.2.1) and we have not attempted to subtract a background signal.

Work to accurately measure the properties of radio point-sources in the C-BASS map is currently ongoing.

In future we will have an independent cross-check of the polarized sky in the overlap region from C-BASS South. During the commissioning of C-BASS South we made measurements that will allow us to make a direct measurement of the polarization angle of the telescope and transfer this to a measurement of TauA. We transmitted a signal at 5.104 GHz from a nearby hilltop using a horn and a signal generator. The hilltop was far enough away to ensure that the telescope was in the far-field region of the transmitted radiation. We modulated the signal at 10 Hz, which allows us to remove emission from background sources by demodulating the TOD. The transmitting horn was mounted such that its polarization angle could be rotated by multiples of 45° and we made raster observations of it transmitting $+Q$, $-Q$, $+U$ and $-U$ signals. We intend to fit the C-BASS South beam to the raster images to determine the polarization angle of the telescope and the level of cross-polarization. Immediately after the observations of the radio horn we made a series of raster images of TauA. With these observations we will be able to accurately measure the polarization angle of TauA at 1° resolution at 5 GHz.

A further future cross-check will be to estimate the polarization fraction from 5 GHz VLA data. However, it is difficult to measure the total power of TauA at $\sim 1°$ resolution with interferometers due to missing baselines. For example, the largest angular scale of detectable emission of the VLA at C-band is $4'$ (in D configuration) [1] whereas TauA is $\sim 5'$ across (Green & Green, 2014).

[1] VLA Observational Status Summary: https://science.nrao.edu/facilities/vla/docs/manuals/oss

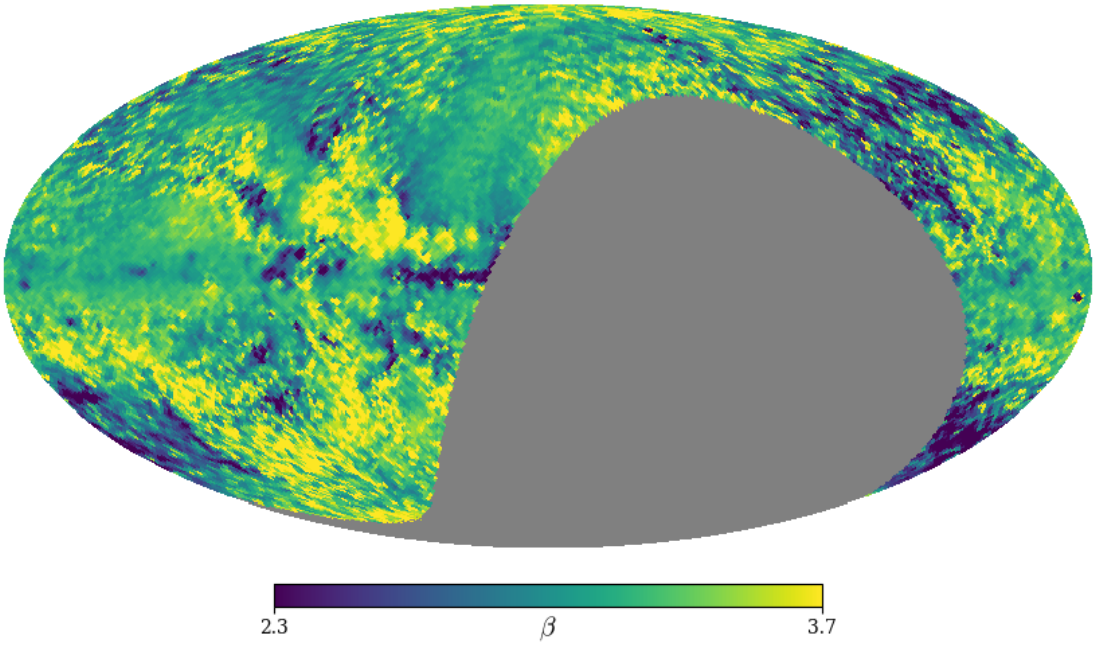

Figure 6.2: Map of the polarized spectral index values between the C-BASS and *Planck* 30 GHz maps. The spectral index was calculated directly from the P maps with no corrections to account for the noise bias.

6.1.3 Rudimentary estimate of spectral index

We made a rudimentary estimate of the spectral index between the C-BASS and *Planck* 30 GHz P maps using Equation 5.2 on each pixel. A map and a histogram of the spectral indices recovered are shown in Figures 6.2 and 6.3 respectively. The structures with the shallowest spectral indices ($2 < \beta < 2.5$) are caused by depolarization from compact HII regions and Faraday rotation at 5 GHz. Many of the high-signal regions corresponding to loops and spurs from the Galactic plane have spectral indices $2.9 < \beta < 3.3$. Outside of these regions are large areas with low signal-to-noise and the spectral indices can range from less than 2 to more than 4.

These are biased estimates of β and have non-trivial statistical properties. In the following section we solve for the spectral index between the two maps using a rigorous Bayesian framework that provides unbiased estimates of β with representative error bars.

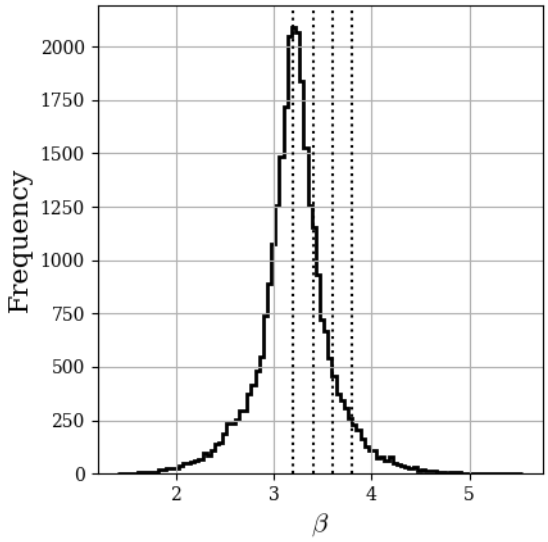

Figure 6.3: Histogram of the polarized spectral index values between pixels in the C-BASS and *Planck* 30 GHz maps. The spectral index was calculated directly from the P maps with no corrections to correct for the bias. The vertical dotted lines are for reference at $\beta = 3.2$, 3.4, 3.6 and 3.8.

6.2 Bayesian analysis of the spectral index of polarized emission between 5 and 30 GHz

We used Bayesian methods to estimate the spectral index in $N_{\rm side} = 64$ pixels between the C-BASS and *Planck* 30 GHz maps. Section 6.2.1 contains a description of the model that we fitted to the maps, pixel-by-pixel, and two choices of prior distribution for the free parameters. Section 6.2.2 shows the recovered estimates of the synchrotron spectral index from simulated data using both sets of priors. Only one of the choices of prior recovers unbiased estimates of the spectral index in low signal-to-noise regions. In Section 6.2.3 we show the results from real data.

6.2.1 Model and choice of priors

In this section we describe the parametric model that we fit to the data and two sets of priors that we later test on simulated data.

We fit a model spectrum with two free parameters, an amplitude A_0 and a spectral index β, pixel-by-pixel to a pair of P maps at different frequencies. The model takes the form

$$P_i(n) = A_0(n) \left(\frac{\nu_i}{\nu_0} \right)^{-\beta(n)} , \tag{6.2}$$

where $P_i(n)$ is the linear polarized intensity in the n^{th} pixel of the map at frequency ν_i, $A_0(n)$ is the linear polarized intensity at a reference frequency ν_0 and $\beta(n)$ is the spectral index. We choose the geometric mean of the map frequencies as the reference frequency, $\nu_0 = 12\,\text{GHz}$.

If a map pixel has a true polarized linear intensity P_0 and $\sigma_P \equiv \sigma_Q = \sigma_U$ then the measured linear polarized intensity, P, is a random variable drawn from a Rician distribution (Equation 4.22). P is a biased estimator of P_0 and it is common to construct the debiased estimator P' of Wardle & Kronberg (1974) where

$$P' \equiv \begin{cases} 0 & P/\sigma_P < \sqrt{2} \\ \sqrt{P^2 - \sigma_P^2} & P/\sigma_P > 3 \end{cases}. \tag{6.3}$$

P' is assumed to be normally distributed about P_0. This is an approximation to the maximum likelihood estimator (Carobbi & Cati, 2008).

To estimate the spectral index in pixels with low signal-to-noise, instead of fitting models to debiased estimates of P_0, we use the full Rician distribution in the likelihood function and use uninformative priors on the free parameters. The likelihood function for a pixel n in maps at frequencies ν_1 and ν_2 is then

$$p(P_1, P_2 | A_0, \beta, \sigma_{P1}, \sigma_{P2}) = \frac{P_1}{\sigma_{P1}^2} e^{-\frac{P_1^2 + (A_0(\nu_1/\nu_0)^{-\beta})^2}{2\sigma_{P1}^2}} I_0\left(\frac{P_1 A_0(\nu_1/\nu_0)^{-\beta}}{\sigma_{P1}^2}\right) \times \tag{6.4}$$

$$\frac{P_2}{\sigma_{P2}^2} e^{-\frac{P_2^2 + (A_0(\nu_2/\nu_0)^{-\beta})^2}{2\sigma_{P2}^2}} I_0\left(\frac{P_2 A_0(\nu_2/\nu_0)^{-\beta}}{\sigma_{P2}^2}\right). \tag{6.5}$$

The maximum likelihood estimate of P_0 from a single measured value of P and σ_P has two solutions, one of which is spurious (Carobbi & Cati, 2008). Using an appropriate choice of prior will remove the spurious maximum from the posterior distribution. The Jeffreys prior for a Rician distribution is $p(P_0, \sigma_P) = \frac{1}{P_0^{1-\epsilon}\sigma_P^{1+\epsilon}}$ where ϵ is a real number (Lauwers et al., 2009). When $\epsilon > 1$ the maximum posterior estimate of P_0 has only one solution. In our case σ_P is known and so we drop the dependence on it from the prior. A uniform distribution in the $Q_0 - U_0$ plane corresponds to a prior $p(P_0) \propto P_0$. A_0 is the estimate of P_0 at ν_0 and so we adopt the prior

$$p(A_0) \propto A_0. \tag{6.6}$$

We test two different prior distributions for β. The spectral index is the slope of a straight line in logarithmic space and for our first choice we set the prior on the angle of that line to be uniform. This gives the prior for the spectral index,

$$p(\beta) \propto \frac{1}{(1+\beta^2)^{3/2}}. \tag{6.7}$$

Our second choice of prior on the spectral index comes from assigning $p\left((\nu_i/\nu_0)^{-\beta}\right) \propto$ constant, which gives

$$p(\beta) \propto |\beta|. \tag{6.8}$$

The two choices of prior are then Prior A,

$$p(A_0)p(\beta) \propto \frac{A_0}{(1+\beta^2)^{3/2}}, \tag{6.9}$$

and Prior B

$$p(A_0)p(\beta) \propto A_0|\beta|. \tag{6.10}$$

We explore the posterior distribution for each pixel using the Metropolis-Hastings algorithm, implemented using PyMC. We discard the first 150,000 steps as a burn in period, during which time we tune the width of the step proposal distributions. After the burn in period we sample for a further 200,000 steps and thin the chain by only taking every fourth sample. These parameters were chosen by inspecting MCMC chains from simulated data during development for convergence. On the simulated data, the chains safely converged early in the burn-in period and the remaining thinned chains well sampled the posterior probability distribution.

For each pixel we calculate the mean and standard deviation of the chain of β values and use these as our estimate of the spectral index and our uncertainty on that measurement. In low signal-to-noise pixels the posterior distributions of β have non-zero skew and kurtosis. This is not unexpected as the likelihood becomes increasingly non-Gaussian at low signal-to-noise and the priors are also non-Gaussian.

6.2.2 Tests on simulated datasets

In this section we show that the method described above recovers unbiased estimates of the spectral index from simulated C-BASS and *Planck* data when using Prior B.

We used PySM to simulate the Stokes Q and U diffuse Galactic synchrotron emission at 5 GHz and 1° resolution in an $N_{\rm side} = 64$ map. We extrapolated both maps to 30 GHz using a spectral index of $\beta = 3$. We added noise realisations to all the maps and created P maps from the noisy maps of Q and U.

The noise realisation for each pixel was generated by multiplying a number drawn from a standard normal distribution by the average Q and U sensitivity for that pixel from the *Planck* 30 GHz white noise maps downgraded to $N_{\rm side} = 64$ and a C-BASS map produced from the E37S60 dataset, smoothed to 1°, downgraded to $N_{\rm side} = 64$ and left in units of antenna temperature.

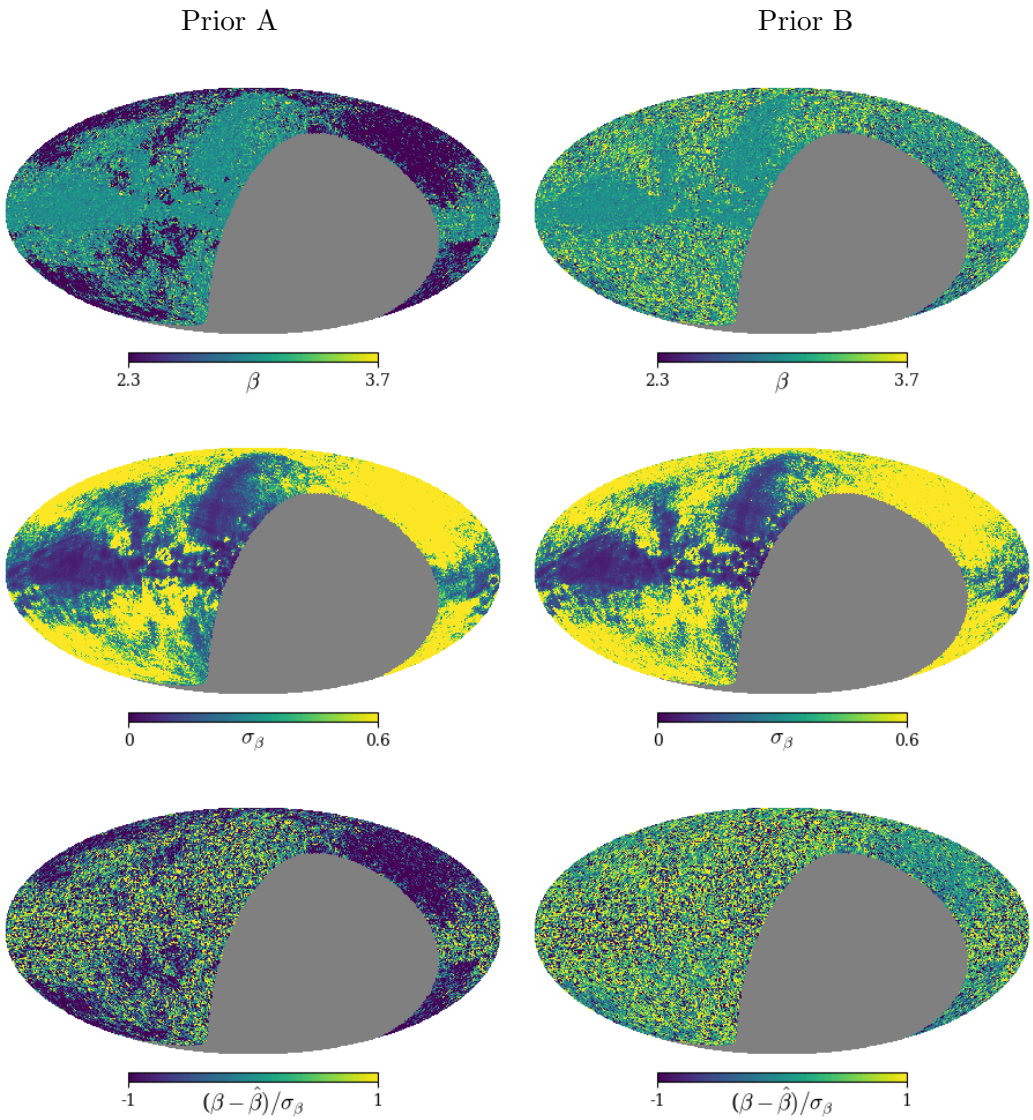

Figure 6.4: From *top* to *bottom*, maps of the estimated mean, standard deviation and normalised deviation from the true spectral index ($\hat{\beta} = 3.0$) using Prior A (*left*) and Prior B (*right*) on the simulated C-BASS and *Planck* 30 GHz maps.

Maps of the recovered estimates of β and their uncertainties, using both priors, are shown in Figure 6.4. The figure also shows maps of the normalised deviation from the true spectral index, $\hat{\beta} = 3$. In low signal-to-noise regions Prior A gives estimates of β that are biased low.

The nature of this bias is most clearly demonstrated by 2D histograms of the estimated β values against the signal-to-noise ratios. These are shown in Figure 6.5, where $R_{\rm CBASS} = P_{\rm CBASS}/\sigma_{P\rm CBASS}$ is the simulated C-BASS signal-to-noise ratio and $R_{\rm Planck} = P_{\rm Planck}/\sigma_{P\rm Planck}$ is the simulated *Planck* 30 GHz signal-to-noise ratio. We also compare the recovered estimates of the spectral index to the geometric mean of

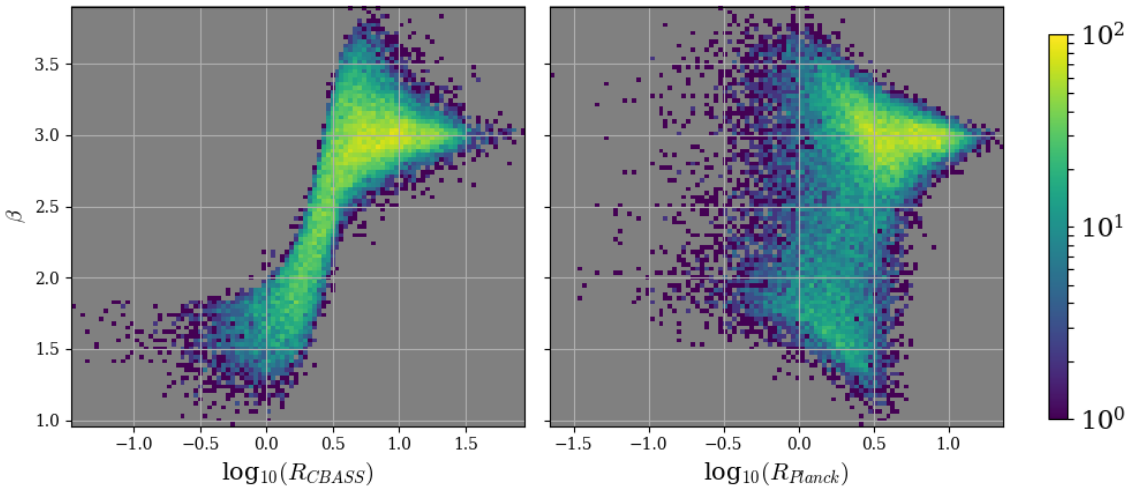

(a) Prior A

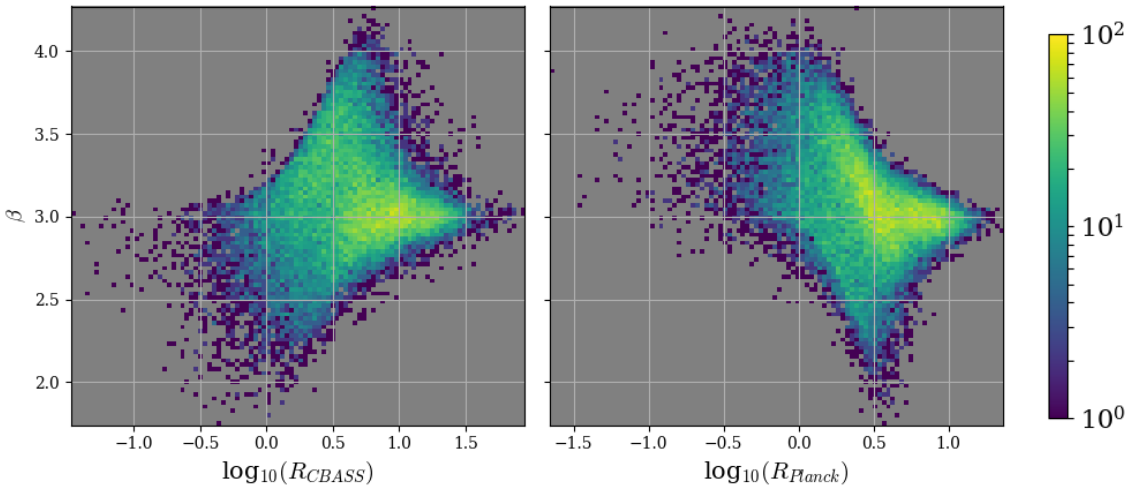

(b) Prior B

Figure 6.5: 2D histograms of the estimate of β in each pixel against the signal-to-noise ratio in the simulated C-BASS (*left*) and *Planck* 30 GHz maps (*right*). The *top* subplots are when using Prior A and the *bottom* subplots are when using Prior B. The true spectral index used to simulate the data was $\hat{\beta} = 3.0$. The colour of each histogram bin shows the number count.

the signal-to-noise ratios in each map in Figure 6.6, $R_{\text{geo}} = \sqrt{R_{\text{CBASS}} R_{\text{Planck}}}$.

1D histograms of the estimated β values are shown in Figure 6.7. The histograms exclude pixels with R_{geo} below four thresholds; firstly all of the pixels were included and then only those with $\log_{10}(R_{\text{geo}}) > 0.5$, 0.75 or 1 were included. These thresholds roughly correspond to $R_{\text{geo}} > 3$, 5 and 10 respectively. Histograms of the normalised deviations, with the same thresholding applied, are shown in Figure 6.8.

The (2D) histograms show that the estimates of β found using Prior A are biased

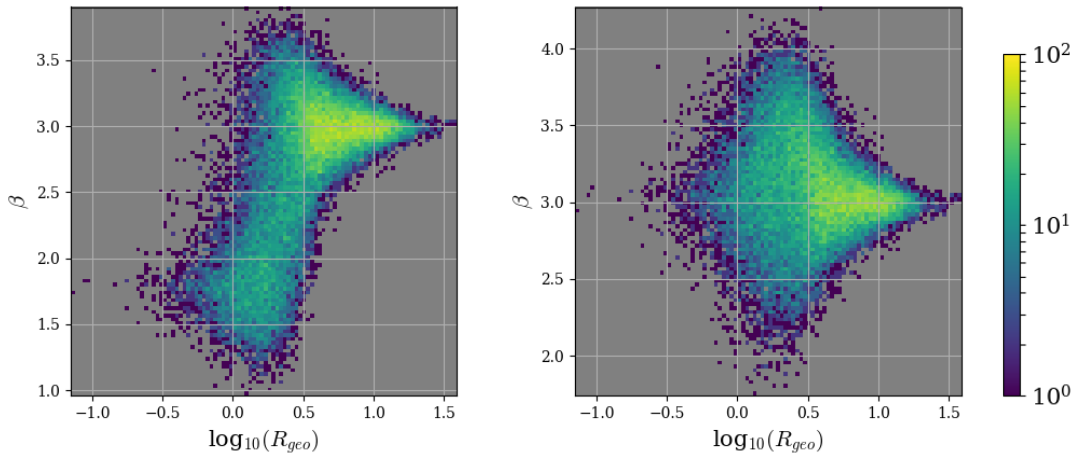

(a) Prior A (b) Prior B

Figure 6.6: 2D histograms of the estimate of β in each pixel against the geometric mean of the signal-to-noise ratio in the simulated C-BASS and *Planck* 30 GHz maps. The *left* subplot is when using Prior A and the *right* subplot is when using Prior B. The true spectral index used to simulate the data was $\hat{\beta} = 3.0$. The colour of each histogram bin shows the number count.

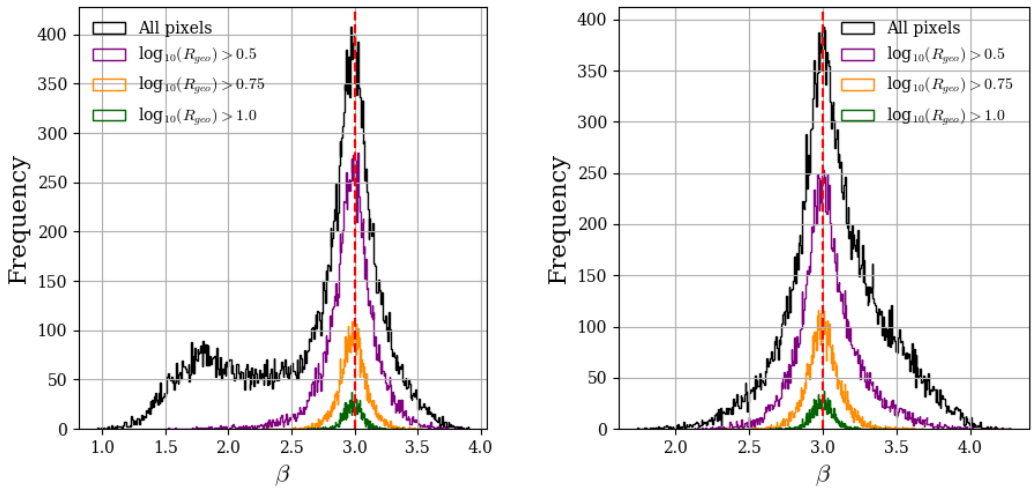

(a) Prior A (b) Prior B

Figure 6.7: Histograms of the estimated β values, excluding pixels below several $R_{\rm geo}$ thresholds. The *dashed red* line is the true spectral index that was used to simulate the data, $\hat{\beta} = 3.0$. The *left* subplot is when using Prior A and the *right* subplot is when using Prior B.

low when $\log_{10}(R_{\rm geo}) < 0.5$ ($R_{\rm geo} \sim 3$), instead preferring $\beta \sim 1.75$. In the high S/N regime the normalised deviation histogram shows that this method recovers

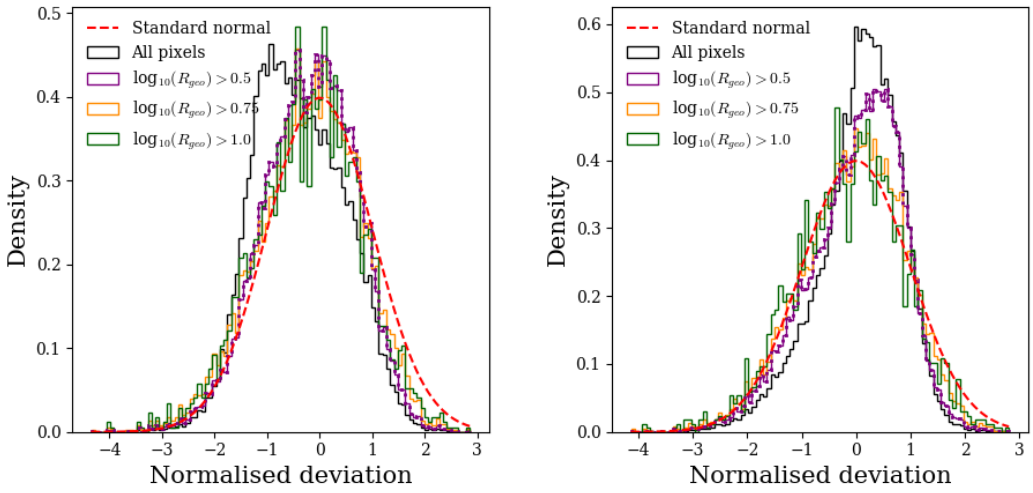

(a) Prior A

(b) Prior B

Figure 6.8: Histograms of the normalised deviation from the true spectral index ($\hat{\beta} = 3.0$), excluding pixels below several $R_{\rm geo}$ thresholds. The *dashed red* line is the standard normal distribution for comparison. The *left* subplot is when using Prior A and the *right* subplot is when using Prior B.

unbiased estimates of β with approximately Gaussian uncertainties. The debiasing approximation in Equation 6.3 is that $P = 0$ when $R_{\rm geo} \sim 3$, and so using Prior A is no better than using this approximation.

When using Prior B the estimates of β are not significantly biased even at small $R_{\rm geo}$. The normalised deviation from the true value is only consistent with the standard normal distribution when $\log_{10}(R_{\rm geo}) > 0.75$. It is not surprising that the posterior distributions for β are not Gaussian because neither the likelihood nor the priors are Gaussian. This method does however allow us to find largely unbiased estimates of β in the low signal-to-noise regime and has similar performance at high signal-to-noise. We will therefore use Prior B when estimating the spectral index from real data.

6.2.3 Results from real data

Maps of the estimated spectral index between the real C-BASS and *Planck* 30 GHz P maps and the uncertainty of those estimates, found using the method described above, are shown in Figure 6.9.

The maps show shallow spectral indices along the Faraday Ghosts and other depolarized features in the C-BASS map. Across the rest of the Galactic plane and the

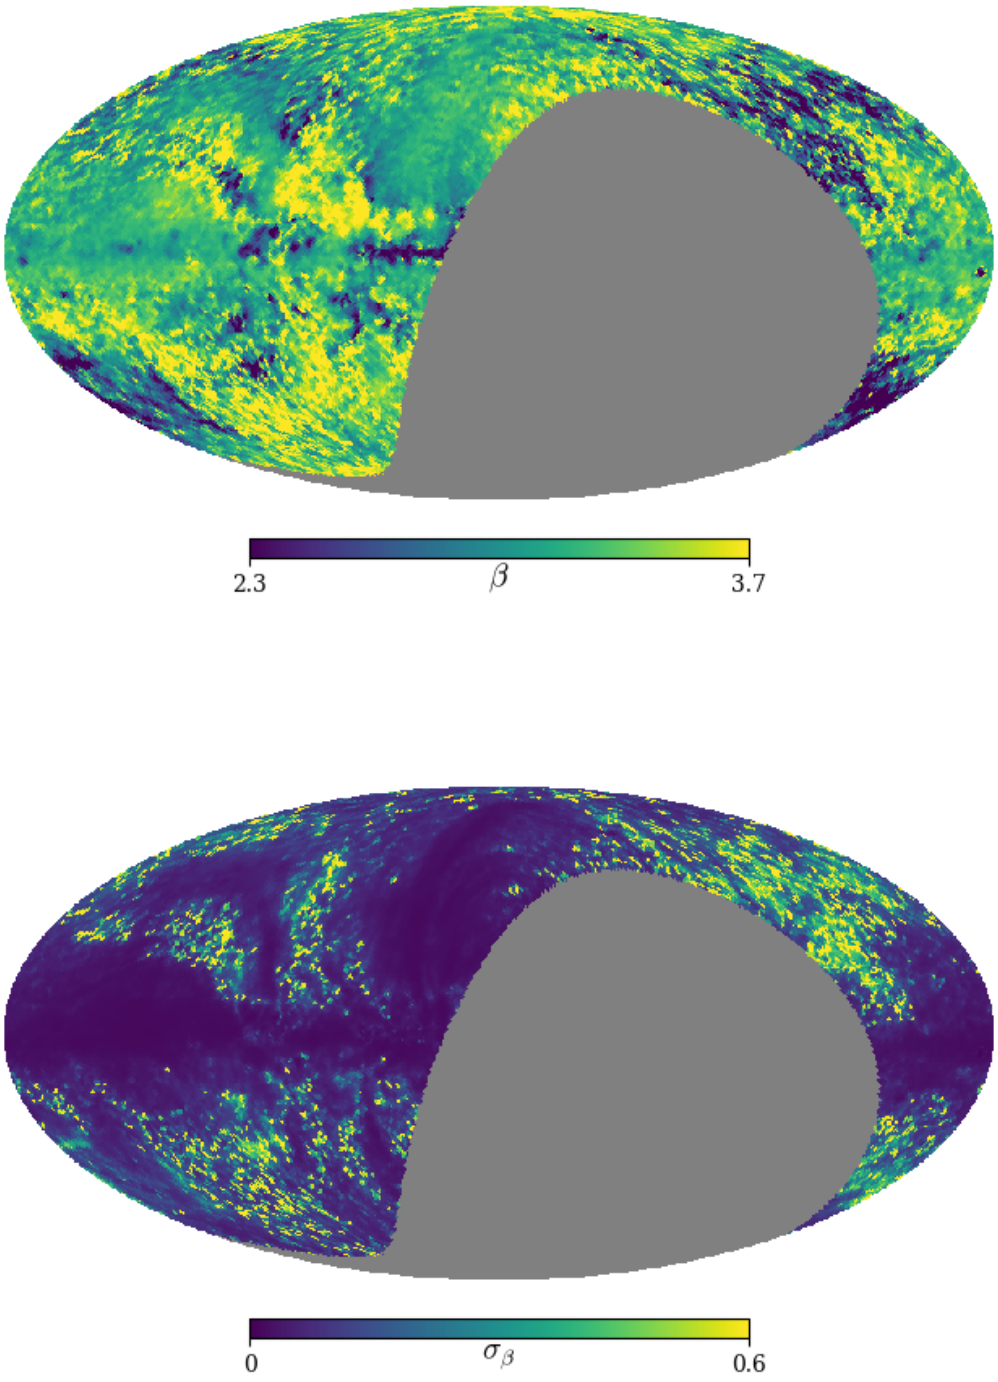

Figure 6.9: Maps of the estimates of β (*top*) and σ_β (*bottom*) between the real C-BASS and *Planck* 30 GHz P maps.

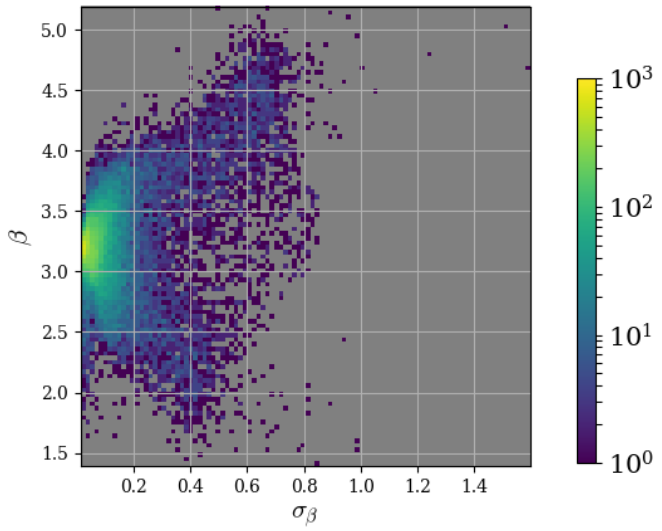

Figure 6.10: 2D histogram of the estimates of β against σ_β between the real C-BASS and *Planck* 30 GHz P maps. The colour of each histogram bin shows the number count.

synchrotron loops and spurs the spectral index is typically 3.0–3.4. Further off the plane are larger fluctuations.

A histogram of β against σ_β is shown in Figure 6.10. The highest density of points are those with small σ_β close to $\beta = 3.2$. The pixels with shallow spectral indices and small error bars are those that correspond to the depolarized features in the C-BASS maps. The weighted average spectral index across the whole sky is $\beta = 3.1658 \pm 0.0002$. There are lots of pixels with β many σ_β away from this weighted average value. The spectral index genuinely changes across the sky.

Histograms of the recovered estimates of β, only including pixels above several $R_{\rm geo}$ thresholds, and a 2D histogram of β against $R_{\rm geo}$ are shown in Figure 6.11. The small number of pixels with large $R_{\rm geo}$ and small β correspond to depolarization in the C-BASS map.

The histograms do not show the spatial structure of variations in the spectral index. In the maps we see structure on both large and small scales. Towards the Galactic centre and in small patches along the Galactic plane are shallower spectral indices caused by depolarization at 5 GHz. Along the rest of the Galactic plane and in the spurs and loops the spectral index is typically close to 3.2. Further off the plane and away from loops and spurs the signal to noise drops and the uncertainties on the estimated spectral indices increases.

To verify whether the large scale structure in the spectral index map is real or

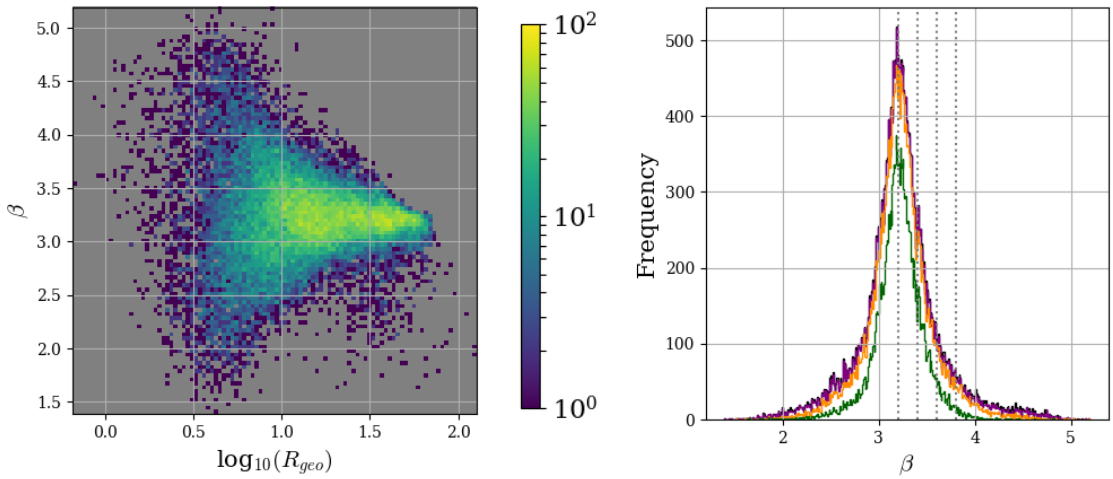

Figure 6.11: Histograms of the recovered estimates of β between C-BASS and *Planck* 30 GHz P maps. The *left hand* plot is a 2D histogram of the spectral index against $\log(R_{\rm geo})$. The colour of each histogram bin shows the number count. The *right hand* plot shows cuts through it: the *black* line includes all map pixels; the *purple*, *orange* and *green* lines include map pixels with signal-to-noise above $\log(R_{\rm geo}) = 0.5$, 0.75 and 1.0 respectively. The *dotted lines* are for reference at $\beta = 3.2$, 3.4, 3.6, and 3.8.

just noise we downgraded the spectral index map from $N_{\rm side} = 64$ to $N_{\rm side} = 16$ and to $N_{\rm side} = 4$. Instead of using the HEALPix UD_GRADE subroutine we set each super-pixel to the weighted average of its children pixels. The downgraded maps are shown in Figure 6.12 and large scale features have not been averaged away.

Off the plane in low-signal pixels the results are more sensitive to systematic errors in the maps such as monopoles and dipoles. Once the polarization angle calibration is accurate in the C-BASS maps we can use T-T plots of the Stokes Q and U values in these regions to estimate the spectral index using a method that is insensitive to offsets in the maps. Furthermore, we only tested this method of estimating the spectral index on simulated data with a perfectly accurate noise model. We did not test how sensitive the results are to inaccuracies in the noise model. To verify these results we need to use the method on simulated data with an inaccurate noise model.

There is also structure in the spectral index map on small angular scales. For example, Figure 6.13 shows a gnomonic projection of the spectral index map around the NPS. Two regions in the map have been identified. The spur outlined in *solid black* has a steep spectral index and the uncertainty on those pixels is low. The spur outlined in *dashed red* has a shallower spectral index, again with high certainty. Beyond these regions are degree-scale fluctuations at the 10% level along the NPS. Higher signal-to-noise polarization maps are needed to confirm whether this structure

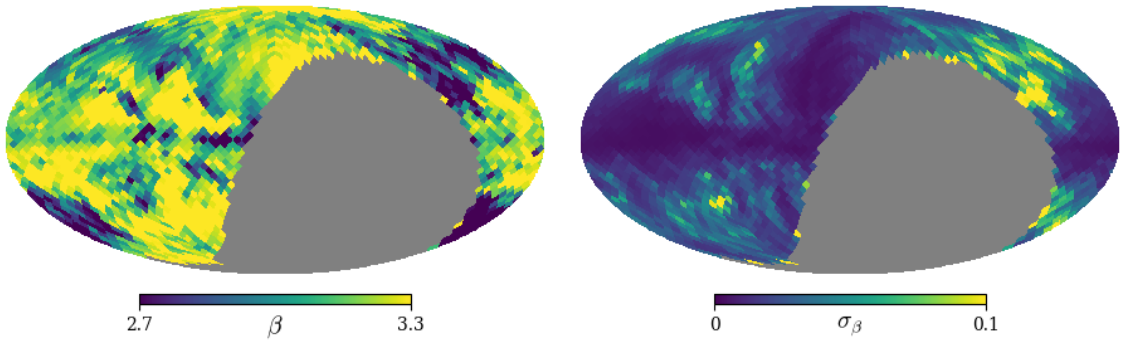

(a) $N_{\rm side} = 16$

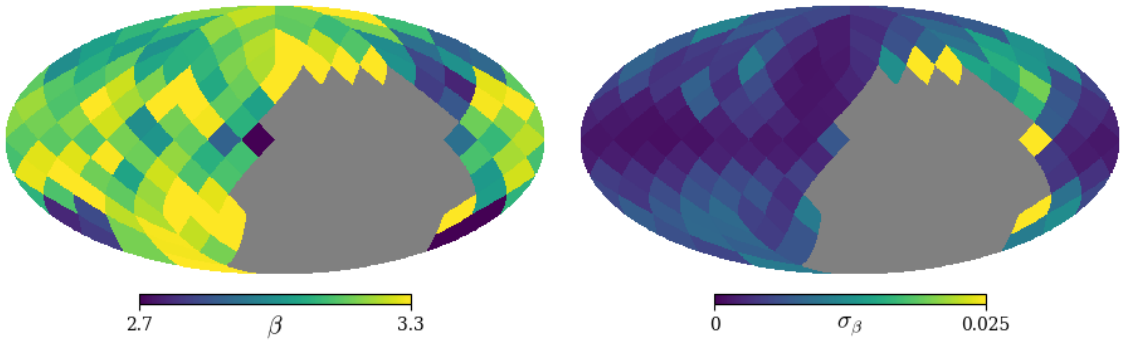

(a) $N_{\rm side} = 4$

Figure 6.12: Downgraded β (*left*) and σ_β (*right*) maps between the real C-BASS and *Planck* 30 GHz P maps. The *top* maps are at $N_{\rm side} = 16$ and the *bottom* maps are at $N_{\rm side} = 4$. The pixels in these maps are the weighted average of the original $N_{\rm side} = 64$ maps.

persists further from the Galactic plane and away from loops and spurs.

Previous analysis by Fuskeland et al. (2014) found that the polarized spectral index between the *WMAP* K and Ka-bands steepened from 2.88 ± 0.01 on the plane to 3.12 ± 0.04 off the plane. We calculated the weighted average spectral index in different Galactic latitude ranges and they are listed in Table 6.3. We see no evidence of the spectral index steepening with latitude.

We found steeper spectral indices than Fuskeland et al. (2014), this either means that the synchrotron spectral index is shallower at higher frequencies or that the C-BASS map is incorrectly calibrated. The measured polarization fraction of TauA showed that the C-BASS P and I maps are consistently calibrated with one another but not whether they are both on the same temperature scale as external datasets.

In both this analysis and the analysis of the previous chapter we found steeper spectral indices between 5 and 30 GHz than expected. In Section 5.2.8 we found

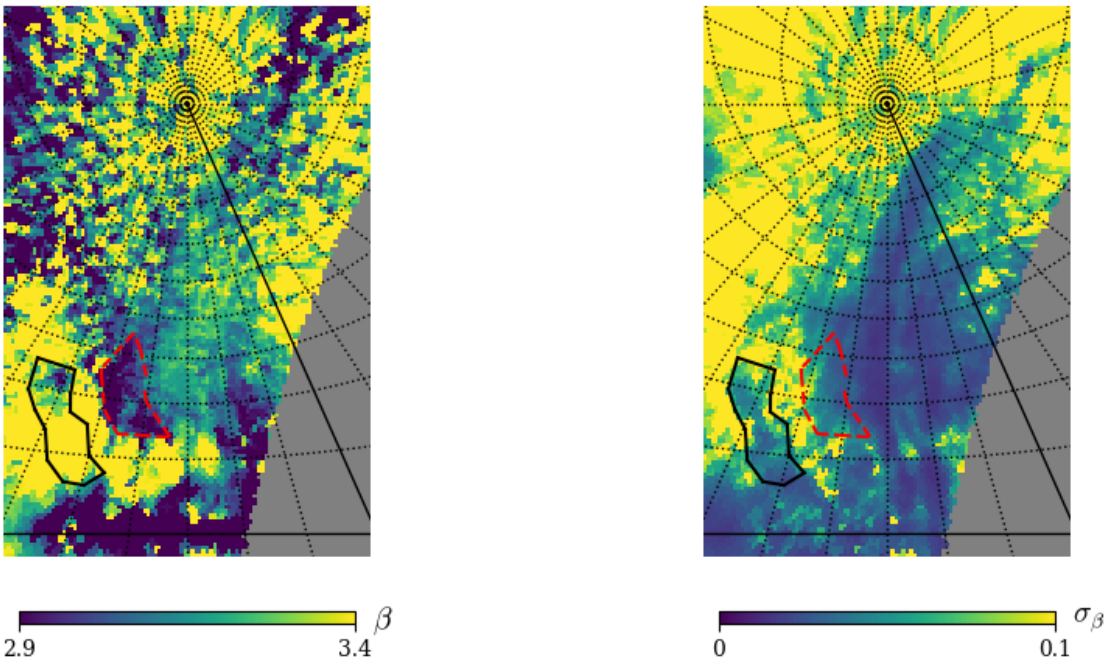

Figure 6.13: The measured polarized spectral index and its uncertainty around the NPS region between the real C-BASS and *Planck* 30 GHz maps. The graticules show lines of constant Galactic latitude and longitude and are separated by 10° intervals.

Table 6.3: Weighted average spectral indices in different Galactic latitude ranges between the real C-BASS and *Planck* 30 GHz maps.

	β	
Galactic latitude range	Weighted average	Error
Whole sky	3.1658	0.0002
$\lvert b\rvert < 10°$	3.1462	0.0004
$10° < \lvert b\rvert < 20°$	3.2078	0.0006
$20° < \lvert b\rvert < 30°$	3.1643	0.0007
$30° < \lvert b\rvert < 50°$	3.1856	0.0006
$50° < \lvert b\rvert < 75°$	3.1596	0.0009
$75° < \lvert b\rvert < 90°$	3.2391	0.0027

that the spectral index of total intensity emission in region 1 of Barnard's Loop was steeper between C-BASS and *Planck* 30 GHz maps than typical free-free emission.

Any error in the calibration of the C-BASS maps would introduce global shifts into the estimated polarized spectral indices but would not remove the spatial structure that we have observed. To confirm these results we need to more rigorously test the calibration of the C-BASS maps.

Chapter 7

Summary and Future Work

> If it be true that good wine needs
> no bush, 'tis true that a good play needs no
> epilogue; yet to good wine they do use good bushes,
> and good plays prove the better by the help of good
> epilogues. What a case am I in then, that am
> neither a good epilogue nor cannot insinuate with
> you in the behalf of a good play!
> –*William Shakespeare, As You Like It*–

In this thesis I have described some of my work on the C-BASS project. In this chapter I give a summary of this thesis, highlighting future work, before briefly discussing C-BASS South.

7.1 The CMB, its foregrounds and C-BASS

I began this thesis with a brief discussion of the current status of CMB cosmology and introduced the C-BASS project.

In Chapter 1 I discussed the CMB, its power spectrum, the difficulties in accurately measuring the CMB B-modes that are predicted to have been produced by primordial gravitational waves and the current status of CMB observations. The intrinsic CMB B-mode power spectrum is a direct probe of the energy scale of inflation. It is weaker than the diffuse Galactic foregrounds at all frequencies and therefore an accurate characterisation of the frequency spectrum of the foreground emission is vital.

In Chapter 2 I gave an overview of the C-BASS project, which aims to produce an all-sky map in intensity and polarization at a central frequency of 5 GHz with resolution $\lesssim 1^\circ$. The C-BASS map is low enough in frequency to be dominated by

synchrotron and free-free emission but high enough for Faraday rotation and depolarization to be small across most of the sky. I gave a brief description of the C-BASS systems and simulated a pixel-based parametric fitting process. Results from the simulations showed that a C-BASS data point reduced the biases and variances of estimates of the CMB and foreground components.

7.2 C-BASS data analysis

In Chapters 3 and 4 I described the C-BASS North data analysis pipeline that converts raw TOD into fully calibrated maps of the sky, I explained how I measured the noise on TOD, tested the C-BASS data for systematics using jackknives and presented preliminary maps of the northern sky.

I described the C-BASS TOD pipeline in Section 3.1. The pipeline makes pointing corrections, calibrates the gain of the system using a noise diode, removes a contaminating signal at 1.2 Hz, flags RFI events, calibrates the TOD into astronomical units and subtracts a scan synchronous ground signal. The reduced data are then saved as FITS files. The TOD pipeline, as it currently exists, was largely developed by others. Currently the polarization angle calibration has not been applied and so there is an arbitrary offset in the C-BASS polarization angles. In future work we will fully test and implement the polarization calibration.

In Section 3.3 I investigated the noise properties of C-BASS TOD. An accurate description of the noise on TOD is needed to improve the mapmaking procedure. In this section I described the statistical properties of periodograms and the TOD noise model. I tested two methods of fitting the noise model to the data; firstly a least-squares approach and secondly a Bayesian approach using the Whittle likelihood. Of the two methods, only the Bayesian method with flat priors on the noise parameters could recover reliable estimates of the noise properties from the short chunks of continuous data in C-BASS survey scans.

I also measured the level of residual 1.2 Hz contamination in the C-BASS TOD. After removal, the variance introduced into the TOD by this contaminating signal is significantly lower the typical white noise levels. Further work is needed to calculate the probability distributions from which we expect $D^{1.2}$ and $D^{1.7}$ values to be drawn so that we can accurately quantify the level of residual 1.2 Hz contamination in the data and propagate the effect of this contamination through to the maps. It may be necessary to tune or improve the 1.2 Hz removal subroutine.

The TOD FITS files are then passed to the DESCART mapmaker to produce the sky maps. In Section 4.3 I tested the various sources of error introduced by the mapmaker and modified the code to allow pixels to be masked when estimating the baseline amplitudes. This masking significantly reduced the size of signal errors in the final maps. DESCART only implements traditional destriping and to reduce the residual correlated noise in the maps will require covariant destriping to be implemented.

I performed a suit of Jackknife tests on the C-BASS maps to test for systematics. The C-BASS data did not formally pass any of the jackknife tests. The main causes of the failures were residual correlated noise in the maps, an incorrect noise model used by DESCART that underestimated the white noise power in the maps by $\sim 12\%$, inconsistencies in the calibration of the TOD at the $\sim 2\%$ level, residual solar emission and residual ground emission.

To account for these results we increased the covariance matrices by 12%, added 2% of the signal maps to the white noise maps and only used night-time survey scans. The residual correlated noise has at least an order of magnitude less power than the sky signal in I, Q and U on all angular scales and the residual ground was only significant in the polarization maps. Further work testing and developing the ground removal subroutine is required, particularly in polarization.

Finally I presented maps of I, σ_I, P and σ_P. We deconvolved the C-BASS beam from the maps and smoothed them to 1° resolution, downgraded them to $N_{\rm side} = 64$ and scaled them to brightness temperature. I produced maps of P because the polarization angle is not currently calibrated in the C-BASS data.

7.3 Measurements of diffuse Galactic emission

I used the C-BASS I and P maps along with surveys at other frequencies to measure the frequency spectrum of low-frequency Diffuse Galactic emission, focussing first on total intensity emission in Chapter 5 and then on polarized emission in Chapter 6.

In Section 5.1 I presented a 3-colour map made from the C-BASS, Haslam and *WMAP* maps. The inputs maps were scaled so that emission with a spectral index of $\beta = 2.7$ appears *colourless*. Emission with a spectral index of $\beta = 2.1$ appears *cyan*, emission with spectral index steeper than $\beta = 2.7$ appears *orange* or *red* and pixels with significant AME appear *blue*. The structure in the 3-colour map demonstrates the complexity of diffuse Galactic foregrounds in total intensity.

The monopoles of the maps must be accurately known to directly measure the spectral index pixel-by-pixel. If they are not accurately known then the method of T-T plots can be used to estimate the spectral index from a region on the sky. In Section 5.2 I extended the standard T-T plot method to try and measure the synchrotron spectral index across the whole sky. I used a Voronoi binning technique to divide the sky into the smallest regions that also meet a given signal threshold. I fitted lines to the T-T plots using a method that recovers unbiased estimates of the slope and intercept when the signal-to-noise on the x abscissas and y ordinates are similar and the errors are Gaussian. I extended the standard T-T plot model using a mixture model to allow pixels to belong to one of several populations. Tests on simulated maps that included a synchrotron component, a free-free component and noise showed that the spectral index of synchrotron and free-free emission could not be unambiguously measured. Fits to the real C-BASS and Haslam maps showed a broad scatter of spectral indices between $\beta = 2$–3. This tells us little about the true spatial variation of the synchrotron spectral index.

We are currently using COMMANDER to fit a parametric model, pixel-by-pixel, to the C-BASS, Haslam, *WMAP* and *Planck* maps to try and measure the synchrotron spectral index across the sky.

In Section 6.1.2 I checked the calibration of the P map by measuring the polarization fraction of TauA. The polarization fraction of TauA in the C-BASS map is consistent with the literature values, although there are no direct observations of TauA at 5 GHz and 1° resolution against which to directly compare our results. In future we could estimate the polarization fraction using VLA data. We have made observations with C-BASS South that should enable us to calibrate the C-BASS polarization maps to high precision.

In Section 6.1.3 I made a rudimentary map of the polarized spectral index between C-BASS and the *Planck* 30 GHz map. To produce this map I made no account of the approximately Rician distribution of measurements of the linear polarized intensity and so this map is likely biased. I did not quantify the error of the estimated spectral indices.

In Section 6.2 I again estimated the polarized spectral index between the C-BASS and *Planck* 30 GHz maps but this time modelled the polarized intensity measurements as Rician random variables. Tests on simulated data show that the method returned unbiased estimates of the spectral index. From the real data I found that the spectral index varied across the sky. There are shallow spectral indices in small regions due to depolarization at 5 GHz. Along much of the Galactic plane and the synchrotron loops

and spurs the spectral index was typically 3.1–3.2. Further off the plane the spectral index varied by larger amounts. We found structure in the spectral index map on large scales and on degree scales in high signal regions. We saw no evidence of the spectral index steepening with Galactic latitude. We only tested the fitting routine on simulations with an accurate noise model and we need to test the robustness of the method to inaccuracies in the noise model to see if the large scale structure persists. We need higher signal-to-noise measurements to determine whether there is further small scale structure in the low signal regions.

7.4 C-BASS South

C-BASS South is currently in the final stages of commissioning in the Karoo desert, South Africa. The main difference between the Southern and Northern systems is that C-BASS South has a digital radiometer/polarimeter with spectral capability. A large part of future work on the C-BASS project will involve developing the pipeline to maximally utilise this spectral resolution.

In this thesis I have not discussed my work on the commissioning of the Southern system. My commissioning work includes making measurements of the system temperature, using absorbing collars to remove resonances most likely caused by leakage of the signal from the thermal gap between the OMT and horn section, measuring the balance temperature across the band and using photogrammetry to align the primary mirror.

There is significant overlap in sky coverage between C-BASS North and C-BASS South. This overlap will provide many opportunities to cross-check the ground removal and calibration.

For example, in Section 6.1.2 I briefly mentioned observations that we made to measure the polarization angle of C-BASS South. We made raster observations of a radio horn transmitting a signal with a known polarization angle. We are yet to analyse these observations but by fitting the co- and cross-polar beams to them we should be able to accurately measure the polarization angle of C-BASS South and the level of cross-polarization. Directly after observing the radio horn we made raster images of TauA. The measurements of TauA made by C-BASS South could then be used to improve the calibration of the C-BASS North map.

Appendix A

Temperature definitions

Four definitions of temperature are commonly used in CMB cosmology, they are: the Planck or thermodynamic temperature ($T_{\rm thermo}$); the Rayleigh-Jeans or brightness temperature ($T_{\rm RJ}$); the CMB temperature, sometimes also called thermodynamic temperature ($T_{\rm CMB}$); and the antenna temperature ($T_{\rm A}$). $T_{\rm CMB}$ should not be confused with the temperature of the CMB (T_0)! $T_{\rm thermo}$, $T_{\rm RJ}$ and $T_{\rm CMB}$ are properties of the sources in the sky where as $T_{\rm A}$ is antenna specific. This appendix contains a summary of the four commonly used temperature definitions and the relationships between them, including colour corrections.

Flux and brightness The flux, S, is defined as the rate at which energy, E, crosses an area, A, perpendicular to the direction of propagation such that

$$S = \frac{dE}{dt}\frac{1}{A}. \tag{A.1}$$

S is a function of frequency and the specific flux or flux density, S_ν is defined as

$$S_\nu = \frac{dS}{d\nu}. \tag{A.2}$$

S_ν is often measured in Janskys, $1\mathrm{Jy} = 10^{-26}\mathrm{Wm}^{-2}\mathrm{Hz}^{-1}$. Brightness or intensity, I, is the flux per unit solid angle and similarly the specific brightness or specific intensity, I_ν, is the flux per unit solid angle per unit frequency. Brightness is often measured in either $\mathrm{Jy\,Sr}^{-1}$ or $\mathrm{Jy\ beam}^{-1}$. Importantly, in free space, brightness is not a function of the separation between the observer and the emitter.

Planck/thermodynamic temperature The brightness of a source can be compared to the temperature of the equivalent blackbody that would reproduce the

observed intensity (blackbodies are both ideal and diffuse emitters that absorb all incident radiation). Blackbody emission obeys Planck's Law

$$I_\nu = \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/kT_{\rm thermo}} - 1} \tag{A.3}$$

where h is Planck's constant and k is the Boltzmann constant. The temperature required to reproduce the observed emission is $T_{\rm thermo}$. This convention is used even when the emitting source is not a blackbody emitter.

Brightness temperature Most observations in radio astronomy are in the Rayleigh-Jeans limit, where $h\nu \ll kT_{\rm thermo}$. In this limit

$$I_\nu = \frac{2k\nu^2 T_{\rm RJ}}{c^2}. \tag{A.4}$$

The temperature of an equivalent blackbody in the Rayleigh-Jeans limit is $T_{\rm RJ}$. This convention is used even when the emitting source is not a blackbody emitter and when not in the Rayleigh-Jeans limit.

CMB temperature $T_{\rm CMB}$ is defined as the required change in $T_{\rm thermo}$ from T_0 to reproduce the observed I_ν. It is found by differentiating Equation A.4 with respect to $T_{\rm thermo}$ and setting $\delta T_{\rm thermo} = T_{\rm CMB}$ so that

$$T_{\rm CMB} = T_{\rm RJ} \frac{(e^x - 1)^2}{x^2 e^x} \tag{A.5}$$

where $x = h\nu/kT_0$.

This conversion from $T_{\rm CMB}$ to $T_{\rm RJ}$ is only true for monochromatic emission or for observations with delta-function bandwidths. If either of these conditions are not met then the effective frequency of the observation must be changed or equivalently the temperatures must be scaled using colour corrections.

The effective frequency of a source is

$$\nu_{\rm eff} = \frac{\int \nu g(\nu) \nu^{-\beta} d\nu}{\int g(\nu) \nu^{-\beta} d\nu}, \tag{A.6}$$

where $g(\nu)$ is the passband and β is the spectral index. For example, Jarosik et al. (2003) estimate the effective frequency of the *WMAP* K-band map for free-free emission ($\beta = 2.1$) as 22.5 GHz and for synchrotron emission ($\beta = 3.0$) as 22.4 GHz. The C-BASS North effective frequency as a function of spectral index is plotted in Figure A.1.

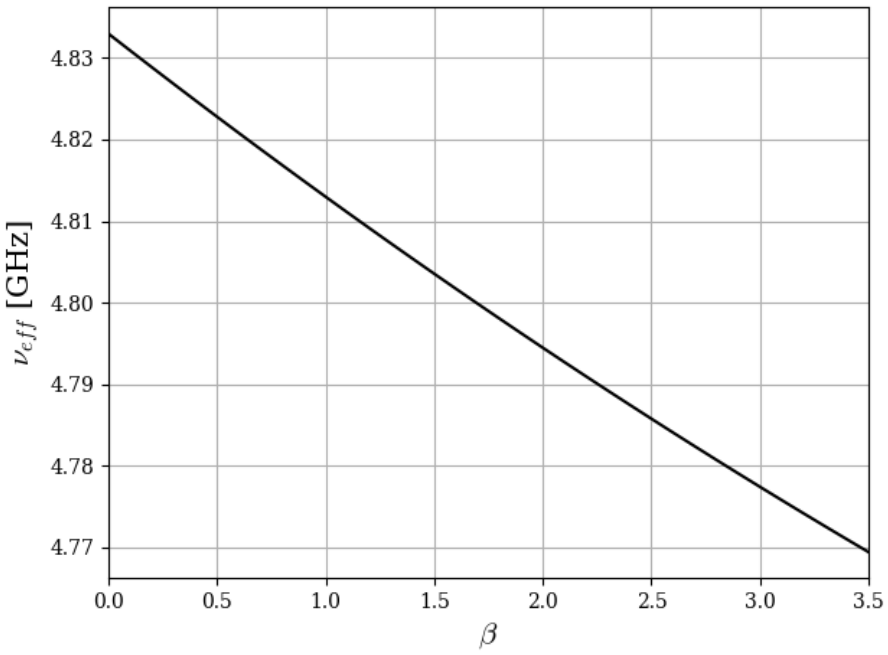

Figure A.1: C-BASS North effective frequency as a function of temperature spectral index.

Instead of changing the effective frequency of an observation, $T_{\rm RJ}$ can be scaled using colour corrections. Letting $\eta(\nu) = x^2 e^x/(e^x-1)^2$ then

$$T_{\rm RJ} = T_{\rm CMB}\eta(\nu_0)C(\alpha), \tag{A.7}$$

where $C(\alpha)$ is the colour correction factor, ν_0 is the gain averaged centre frequency for monochromatic radiation and $\alpha = 2-\beta$.

Planck Collaboration et al. (2015e) approximate the colour corrections for the *Planck* LFI maps with a Taylor expansion

$$C(\alpha) = C_0 + C_1\alpha + C_2\alpha^2. \tag{A.8}$$

The *Planck* 30 GHz map has $\nu_0 = 28.4\,$GHz and $(C_0, C_1, C_2) = (1.005, 0.003, -0.003)$. This means that for free-free radiation $C = 1.005$ and for synchrotron radiation $C = 0.999$.

Antenna temperature The antenna temperature is the temperature of the equivalent resistor (outside the atmosphere) that would result in the same power spectral density from an antenna, P_ν, as the astronomical source. For an unpolarized source it is defined as

$$T_{\rm A} = \frac{P_\nu}{k} = \frac{S_\nu A_{\rm eff}}{2k}. \tag{A.9}$$

where A_{eff} is the effective area of an antenna and the factor of 2 arises as the power is shared equally between the two orthogonal polarizations.

A_{eff} is proportional to the collecting area of the telescope with a proportionality constant called the aperture efficiency. A_{eff} is found from the directivity of the telescope, D, using

$$D = \frac{4\pi A_{\text{eff}}}{\lambda^2}, \tag{A.10}$$

where the directivity is the gain at 0° compared to an isotropic radiator.

T_{A} is antenna-specific, where as the other temperature definitions in this section are properties of the sources on the sky. T_{A} is easy to measure and well defined, so data products are often provided as antenna temperatures. For subsequent astrophysics these temperatures need to be converted to, for example, brightness temperature.

The brightness temperature is related to the antenna temperature via

$$T_{\text{A}} = \frac{A_{\text{eff}}}{\lambda^2} \int \int T_{\text{RJ}}(\theta, \phi) P(\theta, \phi) d\Omega, \tag{A.11}$$

where θ and ϕ are the position angles on the sky, $P(\theta, \phi)$ is the antenna power pattern (normalised to unity at its maximum) and the integral is over the solid angle of the entire sky.

Without knowing the detailed structure of the astronomical emission we can only estimate the brightness temperature from the antenna temperature for a range of typical circumstances, the two most useful being the cases of point source emission and of diffuse emission.

If the emission is from a point source then $P \sim 1$ over the extent of that emission and the integral in Equation A.11 simplifies to

$$T_{\text{A}} = \frac{A_{\text{eff}}}{\lambda^2} \Omega_{\text{source}} \hat{T}_{\text{source}}, \tag{A.12}$$

where Ω_{source} is the solid angle of the source and $\hat{T}_{\text{source}}$ is its average brightness temperature.

If the source size is comparable to the beam and has a constant brightness temperature across its extent, T_{constant}, then the integral in Equation A.11 simplifies to

$$T_{\text{A}} = \frac{A_{\text{eff}}}{\lambda^2} \Omega_{\text{b}} T_{\text{constant}}, \tag{A.13}$$

where Ω_{b} is the solid angle of the beam and side lobes that fall on the source.

Finally, if the source fills the sky such that it is much larger than the beam then Equation A.11 simplifies to

$$T_{\text{A}} = \frac{A_{\text{eff}}}{\lambda^2} \Omega_{\text{A}} T_{\text{constant}}, \tag{A.14}$$

where Ω_{A} is the solid angle of the beam. This conversion is sometimes referred to as the full-beam scale and is useful when comparing the diffuse emission across the sky. For an ideal, lossless system $A_{\mathrm{eff}}/\lambda^2 = 1/\Omega_{\mathrm{A}}$ and $T_{\mathrm{A}} = T_{\mathrm{constant}}$. If there are ohmic losses in the system then this is not the case.

Appendix B

Bayesian statistics and Markov Chain Monte Carlo methods

Bayes theorem states that the probability of event A is related to the probability of event B by

$$p(A|B)p(B) = p(B|A)p(A), \tag{B.1}$$

where $p(A)$ and $p(B)$ are the probabilities of events A and B occurring, $p(A|B)$ is the probability of event A occurring given that event B has occurred and vice versa (Bayes & Price, 1763; LaPlace, 1814).

We often *have* a generative model of measured data and can calculate the probability of getting those data given a set of model parameter values, θ. We often *want* the probability of θ given the data, this distribution is called the posterior probability and can be found using Bayes theorem,

$$p(\theta|\mathbf{d}) = \frac{p(\mathbf{d}|\theta)p(\theta)}{p(\mathbf{d})}, \tag{B.2}$$

where $\mathbf{d}$ is the data. $p(\theta)$ is called the prior and represents our state of knowledge of θ before we analyse the new measured data. The prior is multiplied by $p(\mathbf{d}|\theta)$, which is called the likelihood. The likelihood is the probability of obtaining the data that we have for a given θ. $p(\mathbf{d})$ is the evidence, in parameter estimation it can be treated as a normalising constant but is important in model selection. The posterior distribution then is our new knowledge of θ given the data and the prior.

The choice of prior distribution is often non-trivial and sometimes controversial. If we have testable information about θ then $p(\theta)$ can be chosen using the principle of maximum entropy (Jaynes, 1968).[1] If θ is a continuous parameter then the entropy,

[1] Testable information is information about a probability distribution that is easily verified or falsified, for example its expectation value.

S, is

$$S = -\int p(\theta) \log\left(\frac{p(\theta)}{m(\theta)}\right) d\theta, \tag{B.3}$$

where $m(\theta)$ is the Lebesgue measure, which ensures that the entropy is invariant under a change of variables. S can be maximised using Lagrange multipliers to enforce any constraints from our testable information. Any distribution with lower entropy contains implicit assumptions about extra information.

An example of testable information is the expectation value of a parameter. If this is the only testable information known then the corresponding maximum entropy prior is the exponential function $p(\theta) \propto e^{-A\theta}$, where A is an unknown constant. If the variance is also known, the maximum entropy prior is the normal distribution.

If we have little or no prior information about a parameter then we should choose an uninformative prior. Jeffreys priors are an attractive choice because they are invariant under a reparametrization of the problem (Jeffrey, 1939). Jeffreys priors are found from the Fisher information matrix, $\mathcal{I}$, using

$$p(\theta) \propto \sqrt{\det(\mathcal{I})}. \tag{B.4}$$

The i, j^{th} element of $\mathcal{I}$ is

$$\mathcal{I}_{i,j} = -\mathrm{E}\left[\frac{\partial^2 p(\mathbf{d}|\theta)}{\partial\theta_i \partial\theta_j}\right]. \tag{B.5}$$

Consider a Gaussian distribution with mean, μ, and standard deviation, σ. If σ is fixed then the Jeffreys prior for μ is $p(\mu) =$ constant. The mean is a *location parameter* and a uniform prior is invariant to translations. If μ is fixed then the Jeffreys prior for σ is $p(\sigma) \propto 1/\sigma$, which is uniform in logarithmic space. σ is a *scale parameter* and a $1/X$ prior is invariant to multiplicative scaling factors.

Jeffreys priors are often improper, i.e. they do not integrate to unity. In many cases this does not matter as the resulting posterior distribution will still integrate to unity. If this is not the case it can often be remedied by placing upper and lower limits on the parameter values.

Once we have priors for all of the parameters, we can marginalise over those that we are uninterested in. If we have the probability distribution of a and b, but are uninterested in b. $p(a)$ can be found using

$$p(a) = \int_{-\infty}^{\infty} p(a, b) db = \int_{-\infty}^{\infty} p(a|b) p(b) db. \tag{B.6}$$

Depending on the complexity of the model, the posterior distribution can be highly multidimensional and very complicated. In these cases it can be very difficult,

or even impossible, to find analytic solutions for the parameter values that maximise the posterior distribution. In these cases the posterior distribution can only feasibly be explored using Markov Chain Monte Carlo (MCMC) methods.

MCMC methods sample the posterior distribution by creating Markov chains. Monte Carlo techniques estimate a probability distribution by drawing random samples from it and a chain of samples is Markovian if the $n+1^{\text{th}}$ step in the chain is determined only by the state of the chain at the n^{th} step. After an initial burn in period, the steady state distribution of an MCMC chain is the posterior distribution.

There are many MCMC algorithms and below we summarise the Metropolis-Hastings Algorithm and Gibbs sampling, both of which work by taking a random walk through the parameter space. Gibbs sampling is a special case of the Metropolis-Hastings algorithm that is useful for high-dimensional problems.

Metropolis-Hastings algorithm The Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970) is one of the most commonly used MCMC step methods and there are many variants on it. At the n^{th} step in the chain, the $n+1^{\text{th}}$ step is found as follows:

1. A proposal distribution $g(\theta_{n+1}|\theta_n)$ is used to decide the next step in the random walk.
2. The acceptance ratio $\alpha = \frac{p(\theta_{n+1}|\mathbf{d})}{p(\theta_n|\mathbf{d})}$ is calculated.
3. If $\alpha \geq 1$ then the proposed step is accepted. If $\alpha \leq 1$ the proposed step is accepted with a probability α.
4. Repeat.

The algorithm is started by choosing an arbitrary initial position θ_0. The proposal distribution must be symmetric, $g(\theta_{n+1}|\theta_n) = g(\theta_n|\theta_{n+1})$, and a common choice is a Gaussian distribution centred on θ_n with a tuneable width to give the desired average acceptance rate.

If θ_0 is in a low probability region then the early steps in the chain may not be representative of the posterior distribution. The samples from an initial burn in period should be discarded to avoid these biasing the estimated posterior distribution. The width of the proposal distribution can be tuned during this burn in period to ensure a reasonable acceptance rate, but this tuning should not continue beyond the end of

the burn in period as the resulting chain would not be Markovian. The Metropolis-Hastings algorithm produces correlated chains and this correlation can be reduced by thinning the samples in the converged chains.

Gibbs sampling Gibbs sampling is a special case of the Metropolis-Hastings algorithm that is useful for models with many free parameters (Geman & Geman, 1984). If there are m free parameters then θ is a vector $(\theta^1, \theta^2, \theta^3, ..., \theta^m)$ where θ^j is the j^{th} element of the vector. Each iteration of the sampler steps every element of θ, one at a time, such that

- θ^1_{n+1} is drawn from $g(\theta^1_{n+1}|\theta^1_n, \theta^2_n, \theta^3_n, ...)$
- θ^2_{n+1} is drawn from $g(\theta^2_{n+1}|\theta^1_{n+1}, \theta^2_n, \theta^3_n, ...)$
- θ^3_{n+1} is drawn from $g(\theta^3_{n+1}|\theta^1_{n+1}, \theta^2_{n+1}, \theta^3_n, ...)$
- and so on...

The proposal step is always accepted. It is far easier to propose steps for one θ^j at a time than for the full vector. This is sometimes called Metropolis-within-Gibbs or variable-at-a-time Metropolis-Hastings sampling.

Bibliography

Ade P. A. R. et al., 2015, Phys. Rev. Lett., 114, 101301

Ade P. A. R. et al., 2016, Astrophys. J., 833, 228

Ade P. A. R. et al., 2014a, Astrophys. J., 792, 62

Ade P. A. R. et al., 2014b, Phys. Rev. Lett., 112, 241101

Ahmed Z. et al., 2014, in Proc. SPIE, Holland W. S., Zmuidzinas J., eds., Vol. 9153, p. 91531N

Alam S. et al., 2016, eprint arXiv:1607.03155

Ali-Haïmoud Y., 2013, Adv. Astron., 2013, 1

Ali-Haïmoud Y., Hirata C. M., Dickinson C., 2009, Mon. Not. R. Astron. Soc., 395, 1055

Aller H. D., Reynolds S. P., 1985, Astrophys. J., 293, L73

Anderson L. et al., 2014, Mon. Not. R. Astron. Soc., 441, 24

Aumont J. et al., 2010, Astron. Astrophys., 514, A70

Austermann J. E. et al., 2012, in Millimeter, Submillimeter, Far-Infrared Detect. Instrum. Astron. VI. Proc. SPIE, Vol. 8452, Artic. id. 84521E, 18 pp. (2012)., Holland W. S., ed., Vol. 8452, p. 84521E

Baars J., Genzel R., Pauliny-Toth I., Witzel A., 1977, Astron. Astrophys., 61, 99

Banday A. J., Górski K. M., Bennett C. L., Hinshaw G., Kogut A., Smoot G. F., 1996, Astrophys. J., 468, L85

Barkats D. et al., 2014, Astrophys. J., 783, 67

Barnard E., 1894, Pop. Astron., 2, 151

Battye R. A., Browne I. W. A., Peel M. W., Jackson N. J., Dickinson C., 2011, Mon. Not. R. Astron. Soc., 413, 132

Baumann D. et al., 2009, in AIP Conf. Proc., AIP, pp. 10–120

Bayes M., Price M., 1763, Philos. Trans. R. Soc. London, 53, 370

Bennett C. L. et al., 2003, Astrophys. J. Suppl. Ser., 148, 97

Bennett C. L. et al., 2013, Astrophys. J. Suppl. Ser., 208, 20

Benoît A. et al., 2004, Astron. Astrophys., 424, 571

Benot A. et al., 2002, Astropart. Phys., 17, 101

Berkhuijsen E., Haslam C., Salter C., 1971, Astron. Astrophys., 14, 252

Betoule M. et al., 2014, Astron. Astrophys., 568, A22

Burke B. F., Graham-Smith F., 1997, An Introduction to Radio Astronomy. The Press Syndicate of The University of Cambridge

Burn B. J., 1966, Mon. Not. R. Astron. Soc., 133, 67

Calabretta M. R., Staveley-Smith L., Barnes D. G., 2014, Publ. Astron. Soc. Aust., 31, e007

Cappellari M., Copin Y., 2003, Mon. Not. R. Astron. Soc., 342, 345

Cardoso J.-F., Martin M., Delabrouille J., Betoule M., Patanchon G., 2008, eprint arXiv:0803.1814

Carobbi C., Cati M., 2008, IEEE Trans. Instrum. Meas., 57, 682

Carretti E., 2010, in Dyn. Interstellar Mediu. A Celebr. Can. Galact. Pl. Surv., p. 276

Carretti E. et al., 2013, Nature, 493, 66

Casadei D., Bindi V., 2004, Astrophys. J., 612, 262

Chiang H. C. et al., 2010, Astrophys. J., 711, 1123

Collaboration Planck et al., 2011, Astron. Astrophys., 536, A20

Condon J. J., 2002, in Single-Dish Radio Astron. Tech. Appl., Vol. 278, pp. 155–171

Condon J. J., Cotton W. D., Greisen E. W., Yin Q. F., Perley R. A., Taylor G. B., Broderick J. J., 1998, Astron. J., 115, 1693

Condon J. J. J. J., Ransom S. M., 2016, Essential radio astronomy. p. 361

Copley C. J., 2013, PhD thesis, University of Oxford

Davies R. D., Dickinson C., Banday A. J., Jaffe T. R., Górski K. M., Davis R. J., 2006, Mon. Not. R. Astron. Soc., 370, 1125

de Bernardis P. et al., 1999, New Astron. Rev., 43, 289

de Bernardis P. et al., 2000, Nature, 404, 955

de Gasperis G., Balbi A., Cabella P., Natoli P., Vittorio N., 2005, Astron. Astrophys., 436, 1159

de Oliveira-Costa A. et al., 2000, Astrophys. J., 542, L5

Delabrouille J., Cardoso J.-F., Le Jeune M., Betoule M., Fay G., Guilloux F., 2009, Astron. Astrophys., 493, 835

Delabrouille J., Cardoso J.-F., Patanchon G., 2003, Mon. Not. R. Astron. Soc., 346, 1089

Dicke R. H., 1946, Rev. Sci. Instrum., 17, 268

Dicke R. H., Peebles P. J. E., Roll P. G., Wilkinson D. T., 1965, Astrophys. J., 142, 414

Dickinson C., Davies R. D., Davis R. J., 2003, Mon. Not. R. Astron. Soc., 341, 369

Dodelson S., 2003, Modern cosmology. Academic Press, p. 440

Donzelli S. et al., 2006, Mon. Not. R. Astron. Soc., 369, 441

Draine B. T., 2011, Physics of the interstellar and intergalactic medium. Princeton University Press, p. 540

Draine B. T., Hensley B., 2013, Astrophys. J., 765, 159

Draine B. T., Lazarian A., 1998, Astrophys. J., 494, L19

Draine B. T., Lazarian A., 1999, Astrophys. J., 512, 740

Dunkley J. et al., 2009, in AIP Conf. Proc., Vol. 1141, AIP, pp. 222–264

Dünner R. et al., 2012, Astrophys. J., 762, 10

Einstein A., 1915, Sitzungsberichte der Königlich Preußischen Akad. der Wissenschaften, 844

Einstein A., 1916, Ann. Phys., 354, 769

Eriksen H. K. et al., 2006, Astrophys. J., 641, 665

Eriksen H. K., Jewell J. B., Dickinson C., Banday A. J., Górski K. M., Lawrence C. R., 2008, Astrophys. J., 676, 10

Eriksen H. K. et al., 2004, Astrophys. J. Suppl. Ser., 155, 227

Errard J., Feeney S. M., Peiris H. V., Jaffe A. H., 2016, J. Cosmol. Astropart. Phys., 2016, 052

Fernández-Cobos R., Vielva P., Barreiro R. B., Martínez-González E., 2012, Mon. Not. R. Astron. Soc., 420, 2162

Figueras K. M., 2016, PhD thesis, University of Manchester

Filippini J. P. et al., 2010, in Millimeter, Submillimeter, Far-Infrared Detect. Instrum. Astron. V, Holland W. S., Zmuidzinas J., eds., p. 77411N

Fink D. G., Beaty H. W., 2000, Standard handbook for electrical engineers. McGraw-Hill

Finkbeiner D. P., 2003, Astrophys. J. Suppl. Ser., 146, 407

Finkbeiner D. P., Davis M., Schlegel D. J., 1999, Astrophys. J., 524, 867

Fixsen D. J., 2009, Astrophys. J., 707, 916

Flauger R., Hill J. C., Spergel D. N., 2014, J. Cosmol. Astropart. Phys., 8, 039

Fowler J. W. et al., 2010, Astrophys. J., 722, 1148

Fraisse A. et al., 2013, J. Cosmol. Astropart. Phys., 2013, 047

Friedmann A., 1922, Zeitschrift fur Phys., 10, 377

Fuskeland U., Wehus I. K., Eriksen H. K., Næss S. K., 2014, Astrophys. J., 790, 104

Gardner F. F., Davies R. D., 1966, Aust. J. Phys., 19, 441

Geman S., Geman D., 1984, IEEE Trans. Pattern Anal. Mach. Intell., PAMI-6, 721

Génova-Santos R. et al., 2015a, Mon. Not. R. Astron. Soc., 452, 4169

Génova-Santos R. et al., 2015b, in Highlights Spanish Astrophys. VIII, pp. 207–212

Gold B. et al., 2009, Astrophys. J. Suppl. Ser., 180, 265

Gold B. et al., 2011, Astrophys. J. Suppl. Ser., 192, 15

Gorski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, Astrophys. J., 622, 759

Green D. A., Green D. A., 2014, Bull. Astron. Soc. India, 42, 47

Gregory P. C., Scott W. K., Douglas K., Condon J. J., 1996, Astrophys. J. Suppl. Ser., 103, 427

Grimes P., King O., Yassin G., Jones M., 2007, Electron. Lett., 43, 1146

Gull S. F., 1989, in Maximum Entropy and Bayesian Methods, Springer Netherlands, Dordrecht, pp. 511–518

Guth A. H., 1981, Phys. Rev. D, 23, 347

Guzmán A. E., May J., Alvarez H., Maeda K., 2011, Astron. Astrophys., 525, A138

Hafez Y. A. et al., 2008, Mon. Not. R. Astron. Soc., 388, 1775

Halverson N., Carlstrom J. E., Dragovan M., Holzapfel W. L., Kovac J., 1998, in Adv. Technol. MMW, Radio, Terahertz Telesc., Phillips T. G., ed., Vol. 3357, pp. 416–423

Hamilton J.-C., Ganga K. M., 2001, Astron. Astrophys., 368, 760

Hanany S. et al., 2000, Astrophys. J., 545, L5

Haslam C. G. T., Salter C. J., 1971, Mon. Not. R. Astron. Soc., 151, 385

Haslam C. G. T., Salter C. J., Stoffel H., Wilson W. E., 1982, Astron. Astrophys. Suppl., 47, 1

Hastings W. K., 1970, Biometrika, 57, 97

Heiles C., Haffner L. M., Reynolds R. J., Tufte S. L., 2000, Astrophys. J., 536, 335

Heilgendorff H. et al., 2017, The C-Band All-Sky Survey (C-BASS): Radio frequency interference flagging routine

Hinshaw G. et al., 2013, Astrophys. J. Suppl. Ser., 208, 19

Hogg D. W., Bovy J., Lang D., 2010, eprint arXiv:1008.4686

Holler C. M. et al., 2013, IEEE Trans. Antennas Propag., 61, 117

Hu W., White M., 1997, New Astron., 2, 323

Hubble E., 1929, Proc. Natl. Acad. Sci., 15, 168

Hyvarinen A., 1999, IEEE Signal Process. Lett., 6, 145

Ichiki K., 2014, Prog. Theor. Exp. Phys., 2014, 6B109

Ichiki K., Kaji R., Yamamoto H., Takeuchi T. T., Fukui Y., 2013, Astrophys. J., 780, 13

Imbriale W. A., 2003, Large antennas of the Deep Space Network. Wiley-Interscience, p. 302

Irfan M. O., 2014, PhD thesis, University of Manchester

Irfan M. O. et al., 2015, Mon. Not. R. Astron. Soc., 448, 3572

Jaffe T. R., Banday A. J., Eriksen H. K., Gorski K. M., Hansen F. K., 2006, Astrophys. J., 643, 616

Jarosik N. et al., 2011, Astrophys. J. Suppl. Ser., 192, 14

Jarosik N. et al., 2003, Astrophys. J. Suppl. Ser., 145, 413

Jaynes E., 1968, IEEE Trans. Syst. Sci. Cybern., 4, 227

Jeffrey H., 1939, Theory of probability. Clarendon Press, Oxford, p. 380

Jeffreys H., 1946, Proc. R. Soc. A Math. Phys. Eng. Sci., 186, 453

Jenkins A. et al., 1998, Astrophys. J., 499, 20

Jonas J. L., Baart E. E., Nicolson G. D., 1998, Mon. Not. R. Astron. Soc., 297, 977

Jones M. et al., 2017, The C-Band All-Sky Survey (C-BASS): Design and capabilities, in prep.

Keating B. G., Ade P. A. R., Bock J. J., Hivon E., Holzapfel W. L., Lange A. E., Nguyen H., Yoon K. W., 2003, in Polarim. Astron., Fineschi S., ed., Vol. 4843, p. 284

Keihänen E., Keskitalo R., Kurki-Suonio H., Poutanen T., Sirviö A.-S., 2010, Astron. Astrophys., 510, A57

Keihänen E., Kurki-Suonio H., Poutanen T., 2005, Mon. Not. R. Astron. Soc., 360, 390

Keihänen E., Kurki-Suonio H., Poutanen T., Maino D., Burigana C., 2004, Astron. Astrophys., 428, 287

Kelsall T. et al., 1998, Astrophys. J., 508, 44

King O. G., 2009, PhD thesis, University of Oxford

King O. G. et al., 2010, in Millimeter, Submillimeter, Far-Infrared Detect. Instrum. Astron. V, Holland W. S., Zmuidzinas J., eds., p. 77411I

King O. G. et al., 2014a, Mon. Not. R. Astron. Soc., 438, 2426

King O. G. et al., 2014b, Mon. Not. R. Astron. Soc., 446, 1252

King S., Lubin P., 2016, Phys. Rev. D, 94, 023501

Kitching T. D. et al., 2014, Mon. Not. R. Astron. Soc., 442, 1326

Knox L., 1995, Phys. Rev. D, 52, 4307

Kogut A., Banday A. J., Bennett C. L., Górski K. M., Hinshaw G., Smoot G. F., Wright E. L., 1996, Astrophys. J., 464, L5

Kogut A. et al., 2011, J. Cosmol. Astropart. Phys., 2011, 025

Komatsu E. et al., 2009, Astrophys. J. Suppl. Ser., 180, 330

Kovac J. M., Leitch E. M., Pryke C., Carlstrom J. E., Halverson N. W., Holzapfel W. L., 2002, Nature, 420, 772

Kuehr H., Pauliny-Toth I. I. K., Witzel A., Schmidt J., 1981, Astron. J., 86, 854

Kurki-Suonio H., Keihänen E., Keskitalo R., Poutanen T., Sirviö A.-S., Maino D., Burigana C., 2009, Astron. Astrophys., 506, 1511

LaPlace P. S., 1814, Essai philosophique sur les probabilités. Courcier (Paris)

Lauwers L., Barbe K., Van Moer W., Pintelon R., 2009, in 2009 IEEE Intrumentation Meas. Technol. Conf., IEEE, pp. 114–117

Lawson K. D., Mayer C. J., Osborne J. L., Parkinson M. L., 1987, Mon. Not. R. Astron. Soc., 225, 307

Leach S. M. et al., 2008, Astron. Astrophys., 491, 597

Lee A. T. et al., 1999, in Conf. 3K Cosmol., Vol. 476, ASCE, pp. 224–236

Leitch E. M., Readhead A. C. S., Pearson T. J., Myers S. T., 1997, Astrophys. J., 486, L23

Lemaître G., 1927, Ann. la Société Sci. Bruxelles, 47, 49

Lemaître G., 1931, Nature, 127, 706

Liu H., Mertsch P., Sarkar S., 2014, Astrophys. J., 789, L29

Macellari N., Pierpaoli E., Dickinson C., Vaillancourt J. E., 2011, Mon. Not. R. Astron. Soc., 418, 888

Maino D., Banday A. J., Baccigalupi C., Perrotta F., Gorski K. M., 2003, Mon. Not. R. Astron. Soc., 344, 544

Maino D., Burigana C., Górski K. M., Mandolesi N., Bersanelli M., 2002a, Astron. Astrophys., 387, 356

Maino D., Donzelli S., Banday A. J., Stivoli F., Baccigalupi C., 2007, Mon. Not. R. Astron. Soc., 374, 1207

Maino D. et al., 2002b, Mon. Not. R. Astron. Soc., 334, 53

Martinez-Gonzalez E., Diego J. M., Vielva P., Silk J., 2003, Mon. Not. R. Astron. Soc., 345, 1101

Mather J. C. et al., 1990, Astrophys. J., 354, L37

Matsumura T. et al., 2014, J. Low Temp. Phys., 176, 733

Meisner A. M., Finkbeiner D. P., 2014, Astrophys. J., 798, 88

Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, J. Chem. Phys., 21, 1087

Mitchell K. J., Condon J. J., 1985, Astron. J., 90, 1957

Mortlock D. J., Challinor A. D., Hobson M. P., 2002, Mon. Not. R. Astron. Soc., 330, 405

Mortonson M. J., Seljak U., 2014, J. Cosmol. Astropart. Phys., 2014, 035

Moskalenko I. V., Strong A. W., 1998, Astrophys. J., 493, 694

Mukherjee P., Coble K., Dragovan M., Ganga K., Kovac J., Ratra B., Souradeep T., 2003, Astrophys. J., 592, 692

Mukherjee P., Jones A. W., Kneissl R., Lasenby A. N., 2001, Mon. Not. R. Astron. Soc., 320, 224

NAESB, 2015, WEQ-006 Manual Time Error Correction. Tech. rep.

Narcowich F. J., Petrushev P., Ward J. D., 2006, SIAM J. Math. Anal., 38, 574

Nguyen H. T. et al., 2008, in Millim. Submillim. Detect. Instrum. Astron. IV. Ed. by Duncan, Duncan W. D., Holland W. S., Withington S., Zmuidzinas J., eds., Vol. 7020, p. 70201F

Oppenheim A. V., Schafer R. W., 1975, Digital signal processing. Prentice-Hall, p. 585

Orlando E., Strong A., 2013, Mon. Not. R. Astron. Soc., 436, 2127

Oster L., 1961, Astrophys. J., 134, 1010

Oxley P. et al., 2004, in Infrared Spaceborne Remote Sens. XII, Strojnik M., ed., p. 320

Padin S. et al., 2001, Astrophys. J., 549, L1

Padin S. et al., 2002, Publ. Astron. Soc. Pacific, 114, 83

Page L. et al., 2003, Astrophys. J. Suppl. Ser., 148, 39

Papadakis I. E., Lawrence A., 1993, Mon. Not. R. Astron. Soc., 261, 612

Partridge B., Matarrese S., Hu W., Bartelmann M., Verde L., Martínez-González E., Davies R. D., Longair M., 2010, The cosmic microwave background : from quantum fluctuations to the present universe : XIX Canary Islands Winter School of Astrophysics, Rubino-Martin J. A., Rebolo R., Mediavilla E., eds. Cambridge University Press, p. 304

Patanchon G., Cardoso J.-F., Delabrouille J., Vielva P., 2005, Mon. Not. R. Astron. Soc., 364, 1185

Patil A., Huard D., Fonnesbeck C., 2010, J. Stat. Softw., 35, 1

Penzias A. A., Wilson R. W., 1965, Astrophys. J., 142, 419

Perlmutter S. et al., 1999, Astrophys. J., 517, 565

Peucker-Ehrenbrink B., Schmitz B., 2001, Accretion of Extraterrestrial Matter Throughout Earth's History, Peucker-Ehrenbrink B., Schmitz B., eds. Springer US, Boston, MA

Plagge T. et al., 2010, Astrophys. J., 716, 1118

Planck Collaboration, 2015, Astron. Astrophys., 594, 63

Planck Collaboration et al., 2016a, Astron. Astrophys., 594, A1

Planck Collaboration et al., 2016b, Astron. Astrophys., 594, A10

Planck Collaboration et al., 2015a, Astron. Astrophys., 594, A9

Planck Collaboration et al., 2015b, Astron. Astrophys., 594, A8

Planck Collaboration et al., 2015c, Astron. Astrophys., 576, A104

Planck Collaboration et al., 2013a, Astron. Astrophys., 571, A13

Planck Collaboration et al., 2015d, Astron. Astrophys., 594, A25

Planck Collaboration et al., 2013b, Astron. Astrophys., 565, A103

Planck Collaboration et al., 2011, Astron. Astrophys., 536, A6

Planck Collaboration et al., 2014, Astron. Astrophys., 571, A12

Planck Collaboration et al., 2013c

Planck Collaboration et al., 2015e, Astron. Astrophys., 594, A2

Planck Collaboration et al., 2015f, Astron. Astrophys., 594, A4

Planck Collaboration et al., 2016c, Astron. Astrophys., 594, A3

Platania P., Burigana C., Maino D., Caserini E., Bersanelli M., Cappellini B., Mennella A., 2003, Astron. Astrophys., 410, 847

Press W. H., 2007, Numerical recipes : the art of scientific computing. Cambridge University Press, p. 1235

Quigley M. J. S., Haslam C. G. T., 1965, Nature, 208, 741

Reich P., Reich W., 1986, Astron. Astrophys. Suppl., 63, 205

Reich P., Reich W., 1988, Astron. Astrophys. Suppl. Ser., 74, 7

Reich P., Reich W., Testori J. C., 2004, in Magn. Interstellar Mediu., Uyaniker B., Reich W., Wielebinski R., eds., Antalya, pp. 63–68

Reich W., 1978, Astron. Astrophys., 64, 407

Reich W., 1982, Astron. Astrophys. Suppl., 48, 219

Reichardt C. L., 2016, in Underst. Epoch Cosm. Reionization, Astrophys. Sp. Sci. Libr., Vol. 423, Springer International Publishing, Switzerland, pp. 227–245

Reichborn-Kjennerud B. et al., 2010, Proc. SPIE, 7741, 77411C

Remazeilles M., Dickinson C., Banday A. J., Bigot-Sazy M.-A., Ghosh T., 2015, Mon. Not. R. Astron. Soc., 451, 4311

Remazeilles M., Dickinson C., Eriksen H. K. K., Wehus I. K., 2016, Mon. Not. R. Astron. Soc., 458, 2032

Riess A. G. et al., 1998, Astron. J., 116, 1009

Rohlfs K. K., Wilson T. L. T. L., 2004, Tools of radio astronomy. Springer, p. 461

Rowan-Robinson M., May B., 2013, Mon. Not. R. Astron. Soc., 429, 2894

Rubiño-Martín J. A., López-Caraballo C. H., Génova-Santos R., Rebolo R., 2012, Adv. Astron., 2012, 1

Sachs R. K., Wolfe A. M., 1967, Astrophys. J., 147, 73

Sastry C. V., Pauliny-Toth I. I. K., Kellerman K. I., 1967, Astron. J., 72, 230

Scott P. F. et al., 2003, Mon. Not. R. Astron. Soc., 341, 1076

Seiffert M. et al., 2011, Astrophys. J., 734, 6

Serkowski K., Mathewson D. L., Ford V. L., 1975, Astrophys. J., 196, 261

Sheehy C. D. et al., 2010, in Proc. SPIE, Holland W. S., Zmuidzinas J., eds., Vol. 7741, p. 77411R

Shirokoff E. et al., 2009, IEEE Trans. Appl. Supercond., 19, 517

Shukurov A., Berkhuijsen E. M., 2003, Mon. Not. R. Astron. Soc., 342, 496

Silk J., 1968, Astrophys. J., 151, 459

Silsbee K., Ali-Haïmoud Y., Hirata C. M., 2011, Mon. Not. R. Astron. Soc., 411, 2750

Singal J. et al., 2011, Astrophys. J., 730, 138

Smoot G. F. et al., 1992, Astrophys. J., 396, L1

Sokoloff D. D., Bykov A. A., Shukurov A., Berkhuijsen E. M., Beck R., Poezd A. D., 1998, Mon. Not. R. Astron. Soc., 299, 189

Stevenson M. A., 2013, PhD thesis, California Institute of Technology

Stompor R., Leach S., Stivoli F., Baccigalupi C., 2009, Mon. Not. R. Astron. Soc., 392, 216

Story K. T. et al., 2013, Astrophys. J., 779, 86

Sunyaev R. A., Zeldovich Y. B., 1972, Comments Astrophys. Sp. Phys., 4, 173

Sutton D., Johnson B. R., Brown M. L., Cabella P., Ferreira P. G., Smith K. M., 2009, Mon. Not. R. Astron. Soc., 393, 894

Sutton D. et al., 2010, Mon. Not. R. Astron. Soc., 407, 1387

Suzuki A. et al., 2016, J. Low Temp. Phys., 184, 805

Swetz D. S. et al., 2011, Astrophys. J. Suppl. Ser., 194, 41

Szapudi I., Prunet S., Pogosyan D., Szalay A. S., Bond J. R., 2001, Astrophys. J., 548, L115

Taylor A. C. et al., 2003, Mon. Not. R. Astron. Soc., 341, 1066

Taylor A. R., Stil J. M., Sunstrum C., 2009, Astrophys. J., 702, 1230

Tegmark M., 1997, Astrophys. J., 480, L87

Tegmark M., Efstathiou G., 1996, Mon. Not. R. Astron. Soc., 281, 1297

Testori J. C., Reich P., Reich W., 2008, Astron. Astrophys., 484, 733

The COrE Collaboration et al., 2011, eprint arXiv:1102.2181

The LSPE collaboration et al., 2012, ArXiv e-prints

The POLARBEAR Collaboration et al., 2014, Astrophys. J., 794, 171

Thorne B., Dunkley J., Alonso D., Naess S., 2016, eprint arXiv:1608.02841

Thornton R. J. et al., 2016, Astrophys. J. Suppl. Ser., 227, 21

Tristram M., Filliard C., Perdereau O., Plaszczynski S., Stompor R., Touze F., 2011, Astron. Astrophys., 534, A88

Tristram M. et al., 2005a, Astron. Astrophys., 436, 785

Tristram M. et al., 2005b, Astron. Astrophys., 436, 785

Troxel M. A. et al., 2017, eprint arXiv:1708.01538

Turtle A. J., Pugh J. F., Kenderdine S., Pauliny-Toth I. I. K., 1962, Mon. Not. R. Astron. Soc., 124, 297

VallsGabaud D., 1998, Publ. Astron. Soc. Aust., 15, 111

Vaughan S., 2010, Mon. Not. R. Astron. Soc., 402, 307

Vidal M., Dickinson C., Davies R. D., Leahy J. P., 2015, Mon. Not. R. Astron. Soc., 452, 656

Vidal Navarro M. A., 2014, PhD thesis, The University of Manchester

Vio R., Andreani P., 2008, Astron. Astrophys., 487, 775

Vrtilek J. M., Hauser M. G., 1995, Astrophys. J., 455, 677

Wardle J. F. C., Kronberg P. P., 1974, Astrophys. J., 194, 249

Watson R. A. et al., 2003, Mon. Not. R. Astron. Soc., 341, 1057

Weiland J. L. et al., 2011, Astrophys. J. Suppl. Ser., 192, 19

Whittle P., 1953, Ark. för Mat., 2, 423

Wolleben M., Landecker T. L., Hovey G. J., Messing R., Davison O. S., House N. L., Somaratne K. H. M. S., Tashev I., 2010, Astron. J., 139, 1681

Wolleben M., Landecker T. L., Reich W., Wielebinski R., 2006, Astron. Astrophys., 448, 411

Woody D. P., Richards P. L., 1981, Astrophys. J., 248, 18

Wright E. L., 1998, Astrophys. J., 496, 1

Wright E. L., 2004, in Meas. Model. Universe, from Carnegie Obs. Centen. Symp., Freedman W. L., ed., Cambridge University Press, p. 291

Zacchei A. et al., 2011, Astron. Astrophys., 536, A5