# Polarimetry from the Stratosphere with Spider and BLASTPol

by

Jamil Aly Shariff

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Astronomy & Astrophysics
University of Toronto

# Abstract

Polarimetry from the Stratosphere with SPIDER and BLASTPol

Jamil Aly Shariff

Doctor of Philosophy

Graduate Department of Astronomy & Astrophysics

University of Toronto

2015

This thesis presents the hardware development and flight performance of two balloon-borne experiments. The SPIDER experiment is a millimetre-wavelength polarimeter designed to measure $B$-mode polarization in the Cosmic Microwave Background at degree scales. This pattern is the imprint of the primordial gravitational waves predicted to have been produced by inflation. The BLASTPol experiment is a submillimetre-wavelength polarimeter designed to measure the linearly-polarized emission from aligned dust grains in Galactic molecular clouds, inferring the directions of the magnetic fields there. One goal of this measurement is to understand the role of magnetic fields in the earliest stages of star formation.

SPIDER had a Long-Duration Balloon flight around Antarctica in January 2015. BLASTPol had two such flights, in December 2010 and 2012. Analysis of SPIDER data is underway. Results of BLASTPol 2012 data analysis are presented herein.

The design and performance of the SPIDER pointing control system is presented. A new pivot motor control mode was developed, in which the servo drive controlled motor velocity, not current. This mode enabled sinusoidal azimuth scans at a peak speed of 5 deg/s, with a peak acceleration of 0.5 deg$\cdot$s$^{-2}$, in flight. The pointing stability in flight was $1''$ to $2''$ RMS. A new elevation drive system was designed and built for SPIDER.

The SPIDER observing strategy is presented. It enabled observation of a 10% patch of sky, avoiding the sun and Galactic plane, with uniform coverage in declination, and

good cross-linking.

A model of the BLASTPol 2012 PSF was developed, allowing centroiding, flat-fielding, and map deconvolution. The latter was attempted in Fourier space, and using the Lucy–Richardson method.

A net linear polarization of the dust emission in the Carina Nebula was measured by BLASTPol. The mean fractional polarization $p$ is 6.75% ± 0.015%, 6.84% ± 0.016% and 7.06% ± 0.019%, at 250 µm, 350 µm, and 500 µm respectively. A falling polarization spectrum was found, in contrast with the V-shaped spectrum measured in other molecular clouds. The median ratios of the fractional polarization between bands have been measured to be 1.0155 ± 0.00035 between 250 and 350 µm, and 0.9376 ± 0.00056 between 500 and 350 µm.

*We come spinning out of nothingness, scattering stars like dust.*

— Rumi

# Acknowledgements

First, I would like to thank my thesis supervisor: Barth Netterfield. I have been fortunate and privileged to work with a supervisor who acts as a true advisor, mentor, and colleague. To work with Barth is to be treated to a regular stream of fascinating insights, many of which are on topic, and all of which are delivered with personality and good humour. If even a fraction of his keen physical intuition and experimental knowledge has been transferred to me, then I count myself lucky.

My labmates in what is (to me) the *original* Toronto balloon astrophysics crew, have been just as important to me: Steve Benton, Laura Fissel, Natalie Gandilo, and Juan Soler. I learned so much from them, and had a lot of fun in the process. I will miss the long days in the high bay, soldering, drilling, cutting, bolting, and gluing. I'll also miss the late nights testing star cameras and making things rotate. Here's to us: the Gondoliers.

It's hard to find words to express my admiration for the Spider instrument team, with whom I spent nearly three months on the ice. This includes most of the aforementioned, plus Sean Bryan, Cynthia Chiang, Jeff Filippini, Anne Gambrel, Jón Guðmundsson, John Hartley, Bill Jones, Ziggy Kermish, Lorenzo Moncelsi, Johanna Nagy, Ivan Padilla, Sasha Rahlin, John Ruhl, Don Wiebe, and Ed Young. They made the already incredible Antarctic experience that much more memorable. We persevered through trying and sometimes even dire circumstances, where lesser teams would have fragmented. Together, we achieved something amazing. I will repeat what I said during Thanksgiving dinner at McMurdo Station: you are the finest group of people I have ever had the privilege of working with.

The Spider collaboration is deeply indebted to the British Antarctic Survey team, led by Sam Burrell, who ventured to Spider's remote landing site in West Antarctica, where they successfully recovered all of the data recorded by the experiment.

I'd like to give my personal thanks to the two aerospace engineers in our lab: Steven

Li for his tireless last-minute design effort on an adapter piece for the Spider pivot motor, and Javier Romualdez for reviewing my theoretical treatment of the pointing control system.

I am grateful to Giles Novak for insights into polarimetry and dust grain alignment, and for direct feedback on a portion of this manuscript. I would also like to thank the professors on my thesis and exam committees: Roberto Abraham, J. Richard Bond, Gary Hinshaw, and Keith Vanderlinde.

Through my experience in the astronomy graduate program, I made fast, and I hope, lifelong friends. I could not have gotten through this experience without Mubdi Rahman. Thanks for being my mentor, and for pushing me to think outside the box. I will never forget all the discussions, debates, dinners, sponteaneous road trips, pranks, Nerf gun fights, and so much more! To Stephen Ro: my comrade-in-arms. I can always count on you for advice and encouragement, and to enliven nearly any situation, often with a hilarious non sequitur. I'll meet you at Jimmy's Coffee. Thanks Ilana MacDonald and Charles Zhu, for being such good friends over the years, and for being my partners in the astronomy public outreach that keeps us inspired, and reminds us why we do all of this in the first place.

Marie: you bring out the best in me. Over the past year, you have always been there for me, with love and understanding. Thanks for supporting me during the long days spent coding and writing. And thanks for putting up with my long absence in one of the most remote places on Earth! You bring me joy and happiness, every single day that we are together.

Mom and Dad: thank you for making all of this possible. Thanks for instilling in me a sense of curiosity, and an interest in science. Thank you for teaching me to work hard, and not to give up too easily. Thanks for always being there whenever I was worried, or just needed to talk. I dedicate this thesis to you.

Toronto, July 24, 2015

# Brief Contents

# Contents

# List of Tables

# List of Figures

# List of Acronyms and Abbreviations

| | |
|---|---|
| AC | Alternating Current. 126 |
| ACS | Attitude Control System. 41, 44, 46, 62, 63, 71, 72, 75, 81, 129, 219–221, 223 |
| ACT | Atacama Cosmology Telescope. 17, 19 |
| ADC | Analogue-to-Digital Converter. 129 |
| AGM | Absorbed Glass Mat. 53, 57 |
| ALMA | The Atacama Large Millimeter/submillimeter Array. 123 |
| APEX | Atacama Pathfinder Experiment. 216 |
| AR | Anti-Reflection (coating). 31, 36 |
| AWG | American Wire Gauge. 45, 46, 57, 106 |
| Az | Azimuth. 29, 49, 67–69, 71, 72, 76, 87–90, 92–97, 100, 101, 107–109, 111, 112, 132 |
| BAO | Baryon Acoustic Oscillations. 17, 18 |
| BAS | British Antarctic Survey. 112, 114 |
| BICEP | Background Imaging of Cosmic Extragalactic Polarization. 19–21, 23, 29, 31, 35, 36 |
| BIT | Balloon-borne Imaging Testbed. 215, 216 |
| *BLAST03* | BLAST test flight, Fort Sumner, New Mexico, August 2003. 120 |
| *BLAST05* | BLAST science flight, Kiruna, Sweden, June 2005. 120 |
| *BLAST06* | BLAST LDB flight, Antarctica, December 2006. 120 |
| *BLASTPol10* | BLASTPol LDB flight, Antarctica, December 2010. 120, 133, 134, 139 |

| | |
|---|---|
| *BLASTPol12* | BLASTPol LDB flight, Antarctica, December 2012. 120, 126, 128, 133–135, 139, 140, 172, 173, 216 |
| BOOMERanG | Balloon Observations Of Millimetric Extragalactic Radiation and Geophysics. 19 |
| BSC | Bore-sight Star Camera. 44, 100, 220 |
| CCD | Charge-Coupled Device. 133, 216 |
| CIB | Cosmic Infrared Background. 20, 119 |
| CMB | Cosmic Microwave Background. 1, 2, 7, 11, 13–15, 17–23, 79, 85, 94, 100, 105 |
| COBE | Cosmic Background Explorer. 19 |
| CREAM | Cosmic Ray Energetics And Mass (balloon experiment). 53 |
| CSBF | Columbia Scientific Balloon Facility. 25–27, 30, 43, 44, 53, 55, 76, 111 |
| DAC | Digital-to-Analogue Converter. 70–72, 75, 77 |
| DAS | Data Acquisition System. 41, 46, 129, 220 |
| DASI | Degree Angular Scale Interferometer. 19 |
| DC | Direct Current. 45, 55, 61, 74, 75, 81, 131, 223 |
| DC-DC | DC-to-DC Converter. 61, 81, 224 |
| Dec | Declination. 89, 94, 96, 97, 101, 104, 105, 132, 137, 173 |
| DFT | Discrete Fourier Transform. 188–192 |
| dGPS | Differential Global Positioning System. 63, 133, 220 |
| DIP | Dual In-line Package. 57, 58 |
| DSP | Digital Signal Processor. 71, 72, 75, 77, 85, 129 |
| El | Elevation. 29, 62, 79, 80, 83, 87–90, 93–97, 100, 101, 107, 108, 132, 223 |

| | |
|---|---|
| EMF | Electromotive Force. 69 |
| | |
| FFT | Fast Fourier Transform. 188 |
| FLC | Flight Logic Computer (itsy or bitsy). 40–42, 44, 62, 63, 113 |
| FLRW | Friedmann-Lemaître-Robertson-Walker (spacetime metric). 4 |
| FOV | Field Of View. 31, 103 |
| FPGA | Field-Programmable Gate Array. 81 |
| FTS | Fourier Transform Spectrometer. 124 |
| FWHM | Full Width at Half Maximum. 123, 144, 153, 155, 157, 160, 165–167, 188, 190–195 |
| | |
| GMC | Giant Molecular Cloud. 135, 183, 185, 186 |
| GPS | Global Positioning System. 63, 112 |
| | |
| HDD | Hard Disk Drive. 42 |
| HDPE | High-Density Polyethylene. 31 |
| HEALPix | Hierarchical Equal-Area iso-Latitude Pixelization (of the sphere). 95, 105 |
| HK | Housekeeping (electronics). 41, 46, 62, 220, 221 |
| HWP | Half-Wave Plate. 33, 36, 37, 41, 97–100, 119, 124, 125, 132, 137, 161, 162, 176, 180, 186, 220, 221 |
| | |
| IC | Integrated Circuit. 224 |
| IR | Infrared. 33, 121, 122, 139, 212 |
| IRAC | Infrared Array Camera (*Spitzer Space Telescope* instrument). 184, 185 |
| IRAS | Infrared Astronomical Satellite. 140, 141 |
| ISC | Integrating Star Camera. 133, 135, 172, 173 |
| ISM | Interstellar Medium. 120 |

JFET Junction gate Field-Effect Transistor. 129

LABOCA Large APEX Bolometer Camera. 216

$\Lambda$CDM Lambda Cold Dark Matter (cosmological model). 7, 12, 17, 18

LDB Long Duration Balloon(ing). 25, 26, 29–31, 38, 42, 112–114, 120, 134

LED Light-Emitting Diode. 59

LHe Liquid Helium. 37, 124, 125

$LN_2$ Liquid Nitrogen. 124, 125

LOS Line Of Sight. 26, 63, 72, 107, 111, 204, 220

L–R Lucy–Richardson (deconvolution method). 194–197, 204, 216

LST Local Sidereal Time. 89, 96, 97

MAD Median Absolute Deviation. 205, 210, 211

MAS MCE Acquisition Software. 41

MCC MCE Control Computer. 41, 42, 44, 61, 62, 113, 220

MCE Multi-Channel Electronics. 41, 44, 62, 220

MLI Multi-Layer Insulation. 37

MPPT Maximum Power Point Tracking. 55, 59, 64, 67

Mux Multiplexer. 36

NASA National Aeronautics and Space Administration. 25, 26, 43

NEMA National Electrical Manufacturers Association. 79

NiMH Nickel–Metal Hydride. 52, 53

NTD-Ge Neutron Transmutation Doped Germanium. 126, 129

| | |
|---|---|
| OD | Outer Diameter. 69, 75 |
| OSC | Other Star Camera. 133, 135 |
| PBOB | Power Breakout Board. 46, 60–62, 133, 219–223 |
| PCB | Printed Circuit Board. 46, 60, 61, 222, 223 |
| PCI | Peripheral Component Interconnect. 41 |
| pcm | The SPIDER master control program. 40, 72, 75, 77, 79, 81–83, 85, 87–94, 97 |
| PI | Proportional-Integral (control loop). 70 |
| PSD | Power Spectral Density. 109–111 |
| PSF | Point Spread Function. 139, 187–190, 192–195, 216 |
| PTFE | Polytetrafluoroethylene. 31, 46 |
| PV | Photovoltaic (cell). 47 |
| PV | Pressure Vessel. 42, 44 |
| PVC | Polyvinyl chloride. 46 |
| PWM | Pulse Width Modulation. 81 |
| RA | Right Ascension. 89, 94, 96, 97, 101, 104, 105, 132, 137, 173 |
| RAT | Radiative Alignment Torque. 121, 212, 217 |
| RC | Resistor-Capacitor (circuit). 82, 223 |
| RMS | Root Mean Square. 133, 145, 153 |
| RSC | Rotating Star Camera. 44, 220 |
| RTD | Resistive Temperature Detector. 41 |
| RW | Reaction Wheel. 62, 63, 71, 131 |
| SFT | Superfluid Tank. 37–39, 220 |
| SIP | Support Instrumentation Package (CSBF flight hardware). 26, 44, 113, 129, 219, 221 |

| | |
|---|---|
| SPARO | Submillimeter Polarimeter for Antarctic Remote Observations. 185–187 |
| SPIRE | Spectral and Photometric Imaging Receiver (*Herschel* Space Observatory instrument). 20, 119 |
| SPT | South Pole Telescope. 17, 19–21 |
| SQUID | Superconducting Quantum Interference Device. 35, 36, 41 |
| SSA | SQUID Series Array. 36 |
| SSD | Solid-State Drive. 42, 135, 136, 172 |
| SSR | Solid-State Relay. 219–222, 224 |
| TDRSS | Tracking and Data Relay Satellite System. 26, 136, 137 |
| TES | Transition-Edge Sensor. 34, 35 |
| TOAST | Time Ordered Astrophysics Scalable Tools. 173, 195, 198, 203 |
| TOD | Time-Ordered Data. 136, 137 |
| TTL | Transistor–Transistor Logic. 60 |
| UHMWPE | Ultra-High-Molecular-Weight Polyethylene. 33 |
| ULDB | Ultra-Long-Duration Balloon(ing). 27, 216 |
| UTC | Coordinated Universal Time. 62, 95, 101, 112, 134, 135 |
| UV | Ultraviolet. 184, 216 |
| VCS | Vapour-Cooled Shield. 37, 125 |
| VDC | Volts DC. 55, 79, 81, 82, 224 |
| WMAP | Wilkinson Microwave Anisotropy Probe. 17, 19, 133 |

# Chapter 1

# Introduction

Cosmology aims to understand the universe in its entirety, from its origins, to the precise history of its evolution over time into the vast and intricate cosmos of stars and galaxies that is observed at the present day. At the outset, the questions involved perhaps seem too ambitious in scope for progress to be made in answering them through scientific inquiry. Yet, thanks to a now decades-long effort, marked by tremendous theoretical insights and experimental advancements, a great deal is understood. Modern cosmological models are able to make quantitative, testable predictions about the universe: predictions that, in large part, continue to be borne out by observations.

At the heart of our current model is the standard "hot Big Bang" scenario, in which the universe was hotter and denser at earlier times, and has subsequently expanded and cooled. This idea rests upon several strong observational foundations. One is the observed redshift of the light from distant galaxies, indicating their recessional motion [1]. Another is the measured relative abundances of the light chemical elements (deuterium, helium, and lithium), which can be explained by their formation in a period of nucleosynthesis in the first minutes after the Big Bang [2]. A third piece of evidence is the existence and near perfect blackbody spectrum of the Cosmic Microwave Background (CMB) radiation, the remaining thermal glow from the hot early stages of the universe [3, 4].

This standard Big Bang model has been supplemented with a few additional elements, such as the presence of small primordial fluctuations (which manifest as CMB temperature anisotropy) in the otherwise homogeneous early universe. These acted as the seeds for the growth of all present-day structure. The presence of some non-baryonic form of dark matter is essential to this picture, producing the level of inhomogeneity observed at the present day. The observation that the expansion of the universe appears to be accelerating has been attributed to the mysterious dark energy, about which very little is presently known [5].

The initial conditions of the hot Big Bang, otherwise seemingly finely-tuned, find a natural explanation in a theoretical paradigm known as inflation [6]. The combination of standard Big Bang cosmology and inflation is able to reproduce the observed properties of our universe remarkably well. However, inflation raises many questions of its own. The physical mechanism that drove it, and the energy scale at which this occurred, are not presently known. Therefore, many recent observational efforts have focused on detecting direct evidence of an inflationary epoch in the early universe.

The main goal of this chapter is to provide the scientific motivation for the development of the balloon-borne SPIDER experiment (Chapter 2), which has been the focus of much of the work of this thesis. To that end, a brief review of the mathematical foundations of the standard cosmological model is given, along with a brief overview of inflation (Section 1.1). This is followed by a description of the properties of the Cosmic Microwave Background (Section 1.2). Past and ongoing experimental efforts to measure the temperature and polarization anisotropy of the CMB are then discussed, along with the major results to date (Section 1.3). The discussion culminates with an explanation of the experimental challenges involved in detecting a primordial $B$-mode pattern of polarization in the CMB: a signature of inflation. The ways in which SPIDER attempts to meet these challenges are outlined (Section 1.4). Information on scientific ballooning (Section 1.5), polarimetric conventions (Section 1.6), and the work conducted for this thesis (Section 1.7) is also provided.

The other balloon-borne experiment to which the remainder of this thesis work was devoted—BLASTPol—is introduced in Chapter 3. A description of its scientific goals can be found there. BLASTPol data analysis work is presented in Chapters 4 and 5.

## 1.1 Theoretical Underpinnings

Our view of the cosmos changed dramatically with the discovery by Edwin Hubble that the "spiral nebulae" are all in fact galaxies akin to our own Milky Way, and that they are located well outside of it, at vast distances [7]. Hubble followed this discovery a short time later with the revelation that distant galaxies all appear to be moving away from us. Their observed redshifts, interpreted as recessional velocities, follow a linear relation with distance, Hubble's law [1]:

$$v = Hd \tag{1.1}$$

The Hubble parameter $H$ is independent of distance. Therefore, its present-day value $H_0$ is known as the Hubble constant. This is often written as $H_0 = 100h$ km·s$^{-1}$·Mpc$^{-1}$. At present, the measured value is approximately $h = 0.7$ [8]. This relation between recessional velocity and distance is what would be expected for an expansion of space, in which the distances between objects are increasing uniformly.

### 1.1.1 The Friedmann World Models

The development of the General Theory of Relativity [9] provided, for the first time, a theoretical framework for describing the universe as a whole. In particular, it provided a natural explanation for the uniform expansion just described. General Relativity explains gravitation as arising from the curvature of spacetime. This is expressed by the Einstein field equations, which relate the geometry of spacetime to its mass-energy content. A key element of this geometric description is the metric tensor $g_{\mu\nu}$, which relates coordinate separations to physical distances in an arbitrary four-dimensional spacetime. In particular, given the metric, the invariant spacetime interval $ds$ can be written as

$$ds^2 = \sum_\mu \sum_\nu g_{\mu\nu} dx^\mu dx^\nu \tag{1.2}$$

where $x^\mu$ are spacetime coordinates, with $\mu, \nu \in [0, 3]$. In the construction of a cosmological model, General Relativity is typically combined with a central set of assumptions, namely that the universe is isotropic and homogeneous. These assumptions are supported by observations of the distribution of matter on the largest scales, and motivated by philosophical ideas such as the Copernican principle. It has been shown [10, 11] that the metric for any isotropic, homogeneous, uniformly-expanding spacetime is the Friedmann-Lemaître-Robertson-Walker (FLRW) metric. With the FLRW metric, the invariant interval (Eq. 1.2) can be written (setting $c = 1$) as

$$ds^2 = -dt^2 + a^2(t)\left[\frac{dr^2}{1-\kappa r^2} + r^2(d\theta^2 + \sin^2\theta d\phi^2)\right] \tag{1.3}$$

In this equation, $t$ is proper time as measured by comoving observers. The value of $\kappa$ determines the spatial curvature, with $\kappa > 0$, $\kappa = 0$, and $\kappa < 0$ corresponding to the three possible cases of positive, zero, or negative curvature respectively. The dimensionless factor $a(t)$ scales the spatial part of the metric. It is the ratio of the separation of any two points at time $t$ to their separation now. Thus, the evolution of $a(t)$ captures the dynamics of the expansion of the universe. This evolution is described by General Relativity. If $g_{\mu\nu}$ is the FLRW metric, the Einstein field equations can be used to derive the following two relations, known as the Friedmann equations:

$$\left(\frac{\dot{a}}{a}\right)^2 \equiv H^2(t) = \frac{8\pi G}{3}\rho - \frac{\kappa}{a^2} \tag{1.4}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3P) \tag{1.5}$$

In these equations, $\rho$ is the energy density of a perfect (isotropic) fluid permeating the universe, and $P$ is its pressure. The density $\rho$ includes contributions from all the constituents of the universe, such as radiation, matter, and dark energy: $\rho = \rho_r + \rho_m + \rho_{\rm de}$. Thus, at any given epoch, the dynamics of the expansion will be determined by whichever component dominates the energy density of the universe. It is informative to explore this idea in more detail. A generic equation of state for the perfect fluid can be parameterized as

$$P = w\rho \tag{1.6}$$

A simple argument can be used to derive the variation of $\rho$ with scale factor for different constituents. Consider a volume $V$ in the universe with internal energy $U = \rho V$. For the case of adiabatic expansion

$$dU = -PdV \tag{1.7}$$

The volume at any time can be related to the present volume by $V = V_0(a/a_0)^3$, where $a_0$ is the scale factor at present, taken to be unity here. Based on this, Eqs. 1.6 and 1.7 can be used to derive:

$$\frac{d\rho}{da} + \frac{3}{a}\rho(1+w) = 0 \tag{1.8}$$

Eq. 1.8 can be derived more rigorously by requiring that the covariant derivative of the energy-momentum tensor for the perfect fluid must equal zero (see Eq. 2.55 of [12]). Solving this differential equation results in the relation

$$\rho(a) = \rho_0 a^{-3(1+w)} \tag{1.9}$$

Here, $\rho_0 = \rho(a_0 = 1)$: the energy density at the present epoch. The variation of energy density with scale factor can be examined for different constituents:

- For ordinary (non-relativistic) matter, $w = 0$, and $\rho_m(a) = \rho_m(a_0)\, a^{-3}$. This result is as expected for the case where the total number of particles is conserved, and the number density is simply diluted by the expansion.

- For radiation[1], $w = 1/3$, and $\rho_r(a) = \rho_r(a_0)\, a^{-4}$. This result is also intuitive,

[1] It can be derived from the Planck law for blackbody radiation that $P = \frac{1}{3}\rho$ for an isotropic radiation field.

corresponding to the case where the particle number density scales as $a^{-3}$, and the energy per particle (i.e. per photon) scales as $\lambda^{-1} \sim a^{-1}$

- For the simplest form of dark energy, we can choose $w = -1$, leading to $\rho_{\rm de}(a) =$ constant. This corresponds to the contribution from a cosmological constant term on the right-hand side of Eq. 1.4: $\frac{\Lambda}{3} = \frac{8\pi G}{3}\rho_{\rm de}$

Thus, the universe has undergone a transition from being radiation-dominated to matter-dominated in the past, and is now just starting to be dark-energy-dominated.

From Eq. 1.5, we can see that the criterion for accelerated expansion ($\ddot{a} > 0$) is $P < -\rho/3$. For the forms of dark energy typically considered, this criterion is satisfied ($w < -1/3$). Another result from the Friedmann equations is that there is a critical value for the density $\rho_{\rm cr}$ at which the geometry of the universe is flat (Euclidean). For $\rho < \rho_{\rm cr}$, the geometry is open (negative curvature) and for $\rho > \rho_{\rm cr}$, the geometry is closed (positive curvature). By imposing $\kappa = 0$ in Eq. 1.4, and evaluating it at the present epoch ($t = t_0$), we find that

$$\rho_{\rm cr} = \frac{3H_0^2}{8\pi G} \tag{1.10}$$

Therefore, $\rho_{\rm cr} = 1.88h^2 \times 10^{-26}$ kg·m$^{-3}$. If we define the *density parameter* for the $i^{\rm th}$ constituent as $\Omega_i \equiv \rho_i/\rho_{\rm cr}$, then the first Friedmann equation (Eq. 1.4) can be re-written in a different form (the derivation is left as an exercise to the reader):

$$\left(\frac{H}{H_0}\right)^2 = \Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_k a^{-2} + \Omega_\Lambda \tag{1.11}$$

In this equation, $\Omega_k \equiv -(\kappa/H_0^2) = 1 - \Omega$, where $\Omega = \rho/\rho_{\rm cr}$ is the total density parameter. Different models with different values for this set of parameters correspond to different expanding (or contracting) universes with different geometries, present ages, expansion histories, and ultimate fates. These are the Friedmann world models.

A major success of observational cosmology has been the measurement, to high precision, of the cosmological parameter values that are most likely to describe our universe.

This has been accomplished using observations of the CMB anisotropy, and other data sets (Section 1.3). The standard cosmological model has come to be known as $\Lambda$CDM (Lambda Cold Dark Matter), since a model with cold dark matter and a non-zero cosmological constant ($\Lambda$) is in best agreement with the data. So far this discussion has only focused on the smooth, expanding universe. Another essential feature of $\Lambda$CDM is the presence of primordial fluctuations around this smooth background. For instance, primordial density fluctuations were the seeds for the growth (under gravity) of all the present-day structure in the universe. The fluctuations can be introduced as perturbations to the spacetime metric. A theoretical treatment of this is beyond the scope of this thesis.

### 1.1.2 Inflation

Theories of inflation postulate that the universe underwent a period of exponential expansion in the first instants after the Big Bang. At the end of this inflationary phase, whatever field or fields drove this inflation decayed into the known elementary particles of the standard model. This process of *reheating*, along with inflation itself, produced the initial conditions of our observable universe. The motivation for this idea came from considering several outstanding problems in standard Big Bang cosmology.

The first of these problems is known as the *horizon problem.* The universe appears remarkably homogeneous on the largest spatial scales. Perhaps the most striking example of this is the isotropy of the CMB, which is uniform in temperature over its last scattering surface (see Section 1.2) down to a level of $10^{-5}$. Considering the horizon scale at the epoch in question, areas separated by more than $\sim 2^\circ$ [13] on this surface could not possibly have influenced each other by any causal mechanism. Yet, the CMB appears essentially uniform in temperature across the entire sky.

A second problem is known as the *flatness problem.* In the matter-dominated case (safely ignoring the contribution of $\Omega_\Lambda$ at early times), Eq. 1.11 can be used to derive the relation

$$1 - \frac{1}{\Omega} = a\left(1 - \frac{1}{\Omega_0}\right) \tag{1.12}$$

where $\Omega$ is the total density parameter in the past when the scale factor is $a$, and $\Omega_0$ is its present-day value. Examining this equation, if $\Omega = 1$ at any time, then $\Omega_0 = 1$. However, if $\Omega > 1$ or $\Omega < 1$, then $\Omega_0 \gg 1$ or $\Omega_0 \ll 1$. Slight deviations from flatness in either direction increase with time. In order to have $\Omega_0$ be of order unity today, the universe must have been inordinately close to flat initially. This can be considered a type of fine-tuning problem.

A third problem is how the initial perturbations originated. Standard Big Bang cosmology offers no explanation for their existence and the properties of their spectrum, considering them simply to be "primordial".

Alan Guth came up with inflation as a solution to the problem that theories in high-energy particle physics predict the existence of magnetic monopoles, which have never been observed. However, he soon realized that it could be a solution to the flatness and horizon problems [6]. The qualitative explanations are straightforward. Inflation expands regions that are initially causally-connected to a much larger physical scale, solving the horizon problem. The exponential expansion by many orders of magnitude will tend to decrease spatial curvature, producing a universe very close to flat[2]. The perturbations are explained as random quantum-mechanical fluctuations that are expanded to a macroscopic scale, with a nearly scale-invariant spectrum.

Guth's idea was subsequently expanded upon or modified by others [14, 15]. Although myriad inflationary models exist, this discussion will focus only on one of the simplest, known as single-field, slow-roll inflation. In this model, inflation is driven by a scalar field $\phi$, sometimes called the inflaton field. This field has potential energy density $V(\phi)$. From

[2]Deriving Eq. 1.12 for the dark-energy-dominated case, which would produce exponential expansion just like inflation, one finds the factor of $a$ replaced with $a^{-2}$: an inflating universe tends towards flatness.

the energy-momentum tensor for a scalar field, the following relations can be derived

$$\rho = \frac{1}{2}\dot{\phi}^2 + V(\phi) \tag{1.13}$$

$$P = \frac{1}{2}\dot{\phi}^2 - V(\phi) \tag{1.14}$$

The energy density has "kinetic" and "potential" energy terms. To satisify the criterion (from Eq. 1.5) for accelerated expansion that $P < -\rho/3$, the potential energy term must be larger than the kinetic term. One way in which this is achieved is to have the field slowly roll down a shallow potential, so that $\phi$, and the energy density, are nearly constant. To see how a homogeneous scalar field evolves with time in an expanding universe, the expressions for $\rho$ and $P$ can be substituted into the Friedmann equations to obtain

$$\ddot{\phi} + 3H\dot{\phi} + V'(\phi) = 0 \tag{1.15}$$

During slow-roll inflation, the energy density, and therefore $H$, is nearly but not exactly constant. This is typically quantified by the slow-roll parameters. In terms of the Planck mass $m_{\mathrm{P}}^2 \equiv (8\pi G)^{-1}$, these are:

$$\epsilon \equiv \frac{m_{\mathrm{P}}^2}{2}\left(\frac{V'}{V}\right)^2 \tag{1.16}$$

$$\eta \equiv m_{\mathrm{P}}^2\left(\frac{V''}{V}\right) \tag{1.17}$$

The equations above can be used to show that $\epsilon \ll 1$ and $|\eta| \ll 1$ are necessary for inflation [16].

The relevant primordial fluctuations take the form of scalar and tensor perturbations to the metric. These are characterized by their power spectra $\mathcal{P}_{\mathrm{s}}(k)$ and $\mathcal{P}_{\mathrm{t}}(k)$ describing the power in the perturbations as a function of spatial frequency.

$$\mathcal{P}_{\mathrm{s}}(k) = A_{\mathrm{s}}(k/k_0)^{n_{\mathrm{s}}-1} \tag{1.18}$$

$$\mathcal{P}_{\mathrm{t}}(k) = A_{\mathrm{t}}(k/k_0)^{n_{\mathrm{t}}} \tag{1.19}$$

A parameter known as the tensor-to-scalar ratio $r$ is defined as

$$r \equiv \frac{A_{\mathrm{t}}}{A_{\mathrm{s}}} \tag{1.20}$$

Within the context of slow-roll inflation, the following relations hold [17]:

$$n_{\mathrm{s}} - 1 = -6\epsilon + 2\eta \tag{1.21}$$

$$r = 16\epsilon \tag{1.22}$$

$$n_{\mathrm{t}} = -2\epsilon = -\frac{r}{8} \tag{1.23}$$

Furthermore, the value of $r$ is related to the characteristic energy scale during inflation:

$$V^{\frac{1}{4}} \approx r^{\frac{1}{4}} \left(3 \times 10^{16} \text{ GeV}\right) \tag{1.24}$$

where $V$ is in units of $\text{GeV}^4$ when $\hbar = c = 1$. If there were some experimental means of detecting tensor modes and determining the value of $r$ (Section 1.2.3), the inflationary paradigm would be on very sure footing. A measurement of the energy scale would help constrain or rule out various inflationary models. Furthermore, such a measurement would be the first direct probe of such extremely high-energy processes in the early universe, where quantum-gravitational effects are relevant.

## 1.2 The Cosmic Microwave Background

The early universe was sufficiently hot that the (baryonic) matter in it was entirely in the form of elementary charged particles. This *primordial plasma* eventually came to be composed of electrons and protons, with some light nuclei. High-energy photons in the hot thermal bath kept the matter ionized until the universe had expanded and cooled sufficiently, at $T \approx 3000$ K. At this point, electrons could combine with protons and other nuclei to form stable, neutral atoms for the first time. This event is known (in somewhat of a misnomer) as *recombination*. The estimated redshift of recombination is $z_* = 1090$.

In the primordial plasma, radiation was well-coupled to the matter by the mechanism of Thomson scattering of photons off charged particles. The photon mean free path was therefore very low. During recombination, as the matter in the universe underwent a transition from ionized to neutral, Thomson scattering ceased, and photons began freely-streaming through space. It is these photons from the hot, glowing primordial plasma, arriving from all directions, and redshifted by a factor of $z_*$ that are observed as the CMB. A given CMB photon was last-scattered in a narrow range of redshifts around $z_*$. Therefore, the CMB emission originates from a thin spherical shell known as the last scattering surface. Precise measurements of the CMB spectrum reveal it is that of a nearly perfect blackbody with temperature $T_0 = 2.72548 \pm 0.00057$ K [18].

### 1.2.1 CMB Temperature Anisotropy

Although the CMB is remarkably uniform in temperature over the last scattering surface, small fluctuations are present. This anisotropy is due to density perturbations in the primordial plasma. For the baryonic matter and radiation, these perturbations took the form of acoustic waves, driven to oscillation by the inward force of gravity, and the outward pressure of the photon-baryon fluid. Upon recombination, the radiation and matter became decoupled. While the matter perturbations continued to grow over time, the variations in the radiation intensity field became frozen in, reaching observers at the present day. Thus the CMB gives us a snapshot of the fluctuations as they were during recombination, at an estimated time of only 380 000 years after the Big Bang.

Since the pattern of temperature fluctuations is distributed over a sphere, the usual formalism for describing the pattern is to decompose it into spherical harmonic functions $Y_{\ell m}$, in a manner analogous to Fourier decomposition:

$$\frac{\Delta T}{T}(\hat{\mathbf{n}}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) \tag{1.25}$$

where $\hat{\mathbf{n}}$ is the unit direction vector on the sky. A spherical harmonic function of degree

$\ell$ has $2\ell + 1$ possible values for its order $m$. The degree $\ell$ is usually referred to as the *multipole moment*, with larger values of $\ell$ corresponding to smaller angular scales on the sky. From the orthonormality condition for spherical harmonics, the coefficients $a_{\ell m}$ can be computed by integrating over the sphere:

$$a_{\ell m} = \int_{4\pi} Y^*_{\ell m}(\hat{\mathbf{n}}) \frac{\Delta T}{T}(\hat{\mathbf{n}})\, d\Omega \tag{1.26}$$

A common assumption (consistent with inflation) is that the coefficients are described by an underlying Gaussian probability distribution with mean $\langle a_{\ell m} \rangle = 0$ and with a variance given by:

$$\langle a_{\ell m} a^*_{\ell' m'} \rangle = \delta_{\ell \ell'} \delta_{m m'} C_\ell \tag{1.27}$$

The quantity $C_\ell$ describes the variance in power in a given $\ell$-mode, and is known as the *angular power spectrum* of the fluctuations. It is also written as $C_\ell^{TT}$ to distinguish it from the angular power spectra of different modes of the polarization anisotropy (Section 1.2.3).

Figure 1.3 (Section 1.3.1) shows the temperature angular power spectrum predicted by a best-fit $\Lambda$CDM model, along with measured data. There are various interesting features at different angular scales, reflecting how the initial scale-invariant density perturbation spectrum is processed by early-universe plasma physics. Different physical effects are relevant to the power at different spatial scales. These effects will not be reviewed here, except for a brief mention of the oscillatory features in the spectrum, known as "acoustic peaks". The peak positions (angular scales) are related to the physical scale of the sound horizon at the time of recombination, and its harmonics. Heuristically, this is because the power in the anisotropy is enhanced for perturbations (acoustic waves) with wavelengths such that, by the time of recombination, they would be in state of maximum over- or under-density.

### 1.2.2 CMB Polarization Mechanism

The CMB radiation is expected to be partially linearly-polarized, which has been confirmed observationally (Section 1.3.2). Figure 1.1 illustrates the mechanism for this polarization. A radiation field with a quadrupole anisotropy is incident on an electron ($e^-$). The quadrupole anisotropy refers to radiation from hotter and colder areas, separated in their direction of incidence by 90°. Due to the dependence of the Thomson scattering cross-section on polarization direction [19], the two orthogonal components of the polarization of the scattered radiation are contributed by radiation from each of the two different directions of incidence. One of these components is more intense, leading to a net linear polarization.

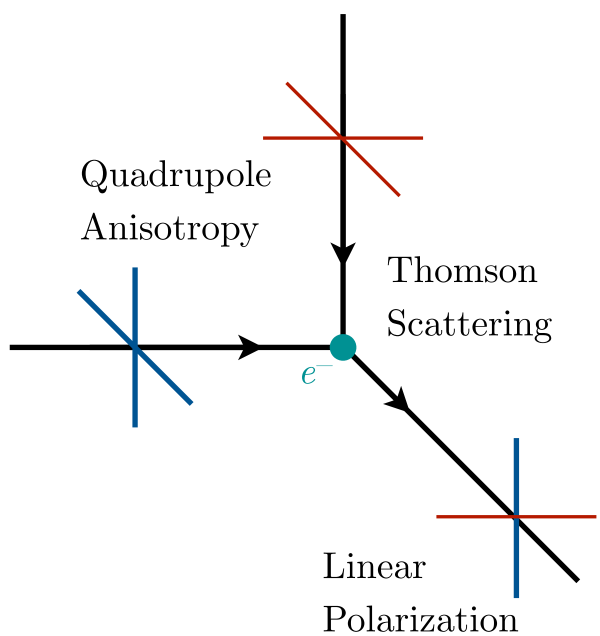


Figure 1.1: Thomson scattering in a radiation field with a quadrupole anisotropy produces a net linear polarization of the scattered light. Blue (thick) lines represent radiation from a hot spot, and red (thin) lines represent radiation from a cold spot. This figure is adapted from Figure 1 of [19].

Rapid Thomson scattering leads to a randomization of the photon directions that destroys any quadrupole anisotropy and linear polarization. Thus, the CMB polarization arises only at the time of "last" scattering: at recombination. Some additional scattering occurs later at the epoch of reionization, when the neutral matter was reionized by radiation from the first luminous objects to form in the universe.

### 1.2.3 *E*-Modes and *B*-Modes

The quadrupole anisotropy that creates linear polarization results from different types of perturbations to the metric. The two relevant geometric types are scalar perturbations, which correspond to density inhomogeneity, and tensor perturbations, which are gravitational waves. These two types create polarization patterns with different symmetry properties. A detailed explanation is in the excellent pedagogical introduction by Hu and White [19], which is only summarized here.

One can imagine an observer in the primordial plasma at the crest of a plane-wave density perturbation of wave vector $\mathbf{k}$. In this observer's coordinate system, there is hotter radiation towards the poles due to infalling material, and colder radiation around the equator. Thus the quadrupole anisotropy resulting from this perturbation is described by the $Y_{20}$ spherical harmonic, which is azimuthally-symmetric (symmetric under a rotation around $\mathbf{k}$). From Figure 1.1, the polarization direction of Thomson-scattered radiation must therefore be either parallel or perpendicular to $\mathbf{k}$, depending on the sign of the quadrupole.

In contrast, an observer at the crest or trough of a plane gravitational-wave perturbation sees a quadrupole pattern of radiation described by the $Y_{2\pm2}$ spherical harmonic, which is not symmetric under rotation around $\mathbf{k}$. This pattern occurs because the wavefronts can be considered to be areas in which a circle of test particles is stretched into ellipses whose semi-major and semi-minor axes oscillate back and forth. The same stretching occurs to photon wavelengths.

Heuristically, one can imagine the superposition of several plane-wave perturbations, all converging radially towards a central point. In the case of scalar perturbations, in which the polarization directions must be parallel or perpendicular to $\mathbf{k}$ for each wave, the superposition will create a tangential or radial pattern of polarization directions around CMB temperature hot spots or cold spots. These patterns are shown in the left-hand column of Figure 1.2. However, the superposition of the gravitational-wave perturbations, which have no requirement for symmetry around $\mathbf{k}$, could equally-well produce the 45° rotations of these patterns, with no particular correlation with temperature. These are shown in the right-hand column of Figure 1.2. The "curl-free" patterns on the left and the "curl-like" patterns on the right are named $E$-modes and $B$-modes respectively, by analogy with the $\mathbf{E}$ and $\mathbf{B}$ vector fields of electromagnetism. An arbitrary pattern of polarization directions on the sky can be decomposed into these two modes.

The crucial result of the above discussion is that scalar perturbations can only produce $E$-modes of polarization, while tensor perturbations can produce both $E$-modes and $B$-modes. Therefore, the ratio of $B$-mode power in the polarization anisotropy is a measure of the ratio of tensor modes to scalar modes: $r$.

A more formal description of $E$- and $B$-modes is as follows. The linear polarization on the sky can be described by the Stokes parameters $Q$ and $U$ (Section 1.6). However, these are not invariant under a rotation in the plane perpendicular to $\hat{\mathbf{n}}$. The quantity $Q \pm iU$ transforms under a rotation by angle $\phi$ as follows:

$$(Q \pm iU)'(\hat{\mathbf{n}}) = e^{\mp 2i\phi}(Q \pm iU)(\hat{\mathbf{n}}) \tag{1.28}$$

This is an example of a spin-2 function (a definition is given in [21]). There is a set of spin-2 spherical harmonics appropriate for expanding this type of function on the sphere, analogous to the regular spherical harmonics $Y_{\ell m}$ used to expand the scalar quantity $T(\hat{\mathbf{n}})$:

$$(Q \pm iU)(\hat{\mathbf{n}}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} {}_{\pm 2}a_{\ell m}\, {}_{\pm 2}Y_{\ell m}(\hat{\mathbf{n}}) \tag{1.29}$$

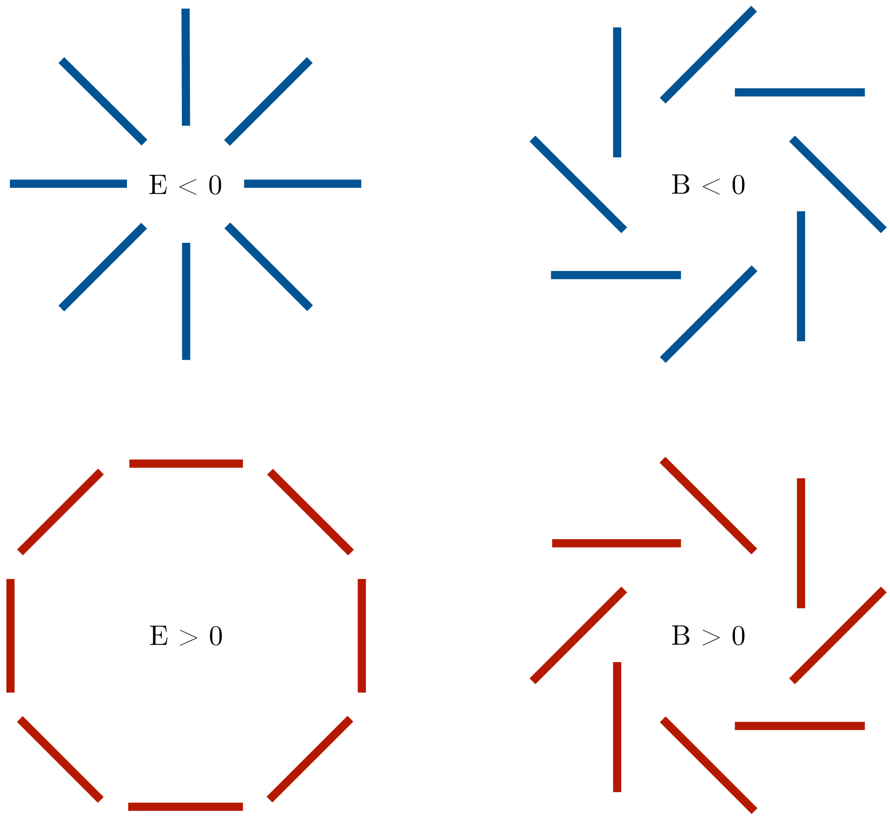


Figure 1.2: A schematic representation of $E$-modes and $B$-modes around temperature hot and cold spots on the sky. Note that each $B$-mode pattern is the rotation of the polarization directions in the $E$-mode pattern of the same sign by 45°. The two modes behave differently under the parity transformation $\hat{\mathbf{n}} \to -\hat{\mathbf{n}}$. $E$-modes are unchanged, while $B$-modes flip signs. This is adapted from Figure 1.4 of [20].

The coefficients of the expansion are determined in a similar way:

$$ {}_{\pm 2}a_{\ell m} = \int_{4\pi} {}_{\pm 2}Y^*_{\ell m}(\hat{\mathbf{n}})(Q \pm iU)(\hat{\mathbf{n}})\, d\Omega \tag{1.30} $$

The coefficients of the $E$- and $B$-modes are defined by linear combinations of the spin-2 coefficients:

$$ a^E_{\ell m} = -({}_2a_{\ell m} + {}_{-2}a_{\ell m})/2 \tag{1.31} $$

$$ a^B_{\ell m} = i({}_2a_{\ell m} - {}_{-2}a_{\ell m})/2 \tag{1.32} $$

Thus, the angular power spectra of the polarization anisotropy can be defined as the variances of the underlying probability distributions of these coefficients:

$$ \langle a^E_{\ell m} a^{*E}_{\ell' m'} \rangle = \delta_{\ell\ell'}\delta_{mm'}C^{EE}_\ell \tag{1.33} $$

$$ \langle a^B_{\ell m} a^{*B}_{\ell' m'} \rangle = \delta_{\ell\ell'}\delta_{mm'}C^{BB}_\ell \tag{1.34} $$

## 1.3 CMB Observations

### 1.3.1 CMB Temperature Measurements

Figure 1.3 shows the CMB temperature angular power spectrum released by *Planck* [22] in 2013. The main cosmological result of the *Planck* observations is that the data are well-fitted by a basic six-parameter ΛCDM model, in keeping with the results from previous experiments. Table 1.1 lists these best fit parameter values, obtained from the combination of *Planck* data, WMAP polarization data, high-$\ell$ data from ACT and SPT, and BAO data [8]. BAO refers to Baryon Acoustic Oscillations: the characteristic length scale of the acoustic oscillations in the primordial plasma is imprinted into the distribution of baryonic matter, such that it can still be observed at late times in the distribution of galaxies.

In Table 1.1, $\Omega_b$ and $\Omega_c$ are the individual density parameters for baryonic matter and cold dark matter respectively ($\Omega_m = \Omega_b + \Omega_c$). The parameter $\tau$ is the Thomson

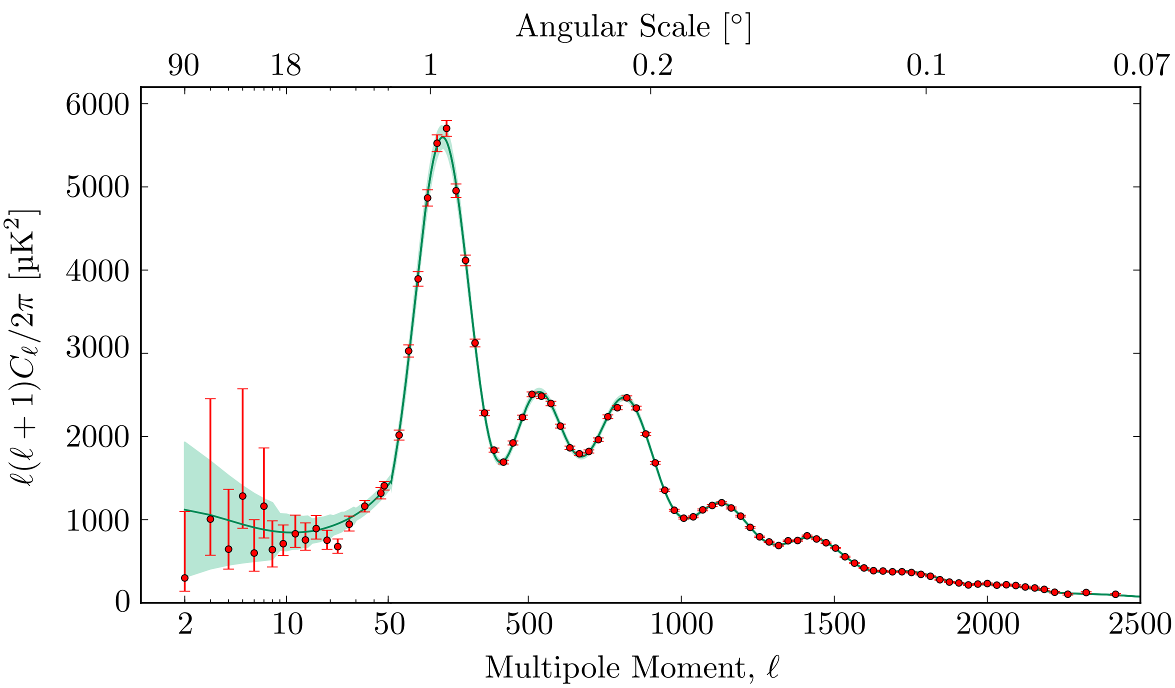


Figure 1.3: The CMB temperature angular power spectrum measured by *Planck*. The green shaded area around the best fit model indicates the uncertainty due to cosmic variance. This uncertainty is also included in the error bars, along with measurement error. This is Figure 37 of [23].

Table 1.1: The best fit six-parameter ΛCDM values [8]

| Parameter | 68% limits (*Planck* + WP + high$\ell$ + BAO) |
|---|---|
| $\Omega_b h^2$ | $0.02214 \pm 0.00024$ |
| $\Omega_c h^2$ | $0.1187 \pm 0.0017$ |
| $100\theta_*$ | $1.04162 \pm 0.00056$ |
| $\tau$ | $0.092 \pm 0.013$ |
| $n_{\mathrm{s}}$ | $0.9608 \pm 0.0054$ |
| $\ln(10^{10} A_{\mathrm{s}})$ | $3.091 \pm 0.025$ |

scattering optical depth to the epoch of reionization. The angular size, on the surface of last scattering, of the comoving sound horizon at recombination is $\theta_*$. The remaining parameters have to do with the power spectrum of primordial scalar perturbations: $\mathcal{P}_\mathrm{s}(k) = A_\mathrm{s}(k/k_0)^{n_\mathrm{s}-1}$. The pivot scale $k_0$ is 0.05 $\mathrm{Mpc}^{-1}$. The base model assumes $\Omega = 1$ and $w = -1$. Therefore, given these values of the six base parameters, other parameter values can be derived, such as $\Omega_\Lambda = 0.692 \pm 0.010$, $H_0 = 67.80 \pm 0.77$ km·s$^{-1}$·Mpc$^{-1}$, and the age of the universe $t_0 = 13.798 \pm 0.037$ Gyr. Contraints on the tensor-to-scalar ratio can be obtained from CMB temperature data alone, although there are strong parameter degeneracies. The *Planck* 2013 upper limit is $r_{0.002} < 0.11$ at 95% confidence, which is quoted for a value of the pivot scale of $k_0 = 0.002$ $\mathrm{Mpc}^{-1}$ rather than 0.05 $\mathrm{Mpc}^{-1}$.

For the sake of brevity, only one recent set of CMB temperature anisotropy measurements is presented here. There is, of course, a long history of such measurements. The first detection of the anisotropy was by the Cosmic Background Explorer (COBE) [24]. This detection was followed by early groundbreaking measurements of the acoustic peaks in the spectrum. Perhaps one of the most notable of these measurements was by the BOOMERanG balloon experiment, which was among the first to determine the standard cosmological parameter values [25]. In terms of more recent observations, one can compare the *Planck* results above to the WMAP nine-year values [26]. The temperature spectrum has been measured out to even higher $\ell$ by large-dish, ground-based telescopes such as ACT [27] and SPT [28].

### 1.3.2 CMB Polarization Measurements

The first detection of CMB ($E$-mode) polarization was by the Degree Angular Scale Interferometer (DASI) [29], and this was followed by a succession of measurements from other CMB polarization experiments (Figure 1.4). Measurements of $E$-mode polarization allow contraints on $\tau$, providing information about reionization. A notable CMB polarization result is from the BICEP experiment at the South Pole, which detected for the first time the peak in $C_\ell^{EE}$ expected at $\ell \sim 140$. BICEP also set the first meaningful constraint on tensor modes to come from CMB polarization alone: $r < 0.72$ at 95% confidence [30].

The analysis of data from BICEP2 led to the detection of a $B$-mode signal on the sky at degree angular scales, at 150 GHz. This signal, if cosmological in origin, was consistent with a level of tensor modes at $r = 0.2^{+0.07}_{-0.05}$ [31]. This result was announced with much fanfare in March 2014, and published in *Physical Review Letters* some time thereafter. However, concerns were raised over the assumptions underlying the Galactic foreground models used in the analysis. At the time, these models were theoretical, or data-driven using preliminary *Planck* results, since the final *Planck* polarization data were not yet ready for release. Even under these foreground-level assumptions, the hypothesis of a Galactic dust origin for the signal could only be rejected at the $2\sigma$ level. It was soon shown that under different but still arguably reasonable assumptions, the dust foreground levels in the BICEP2 field could account for nearly all of the measured signal [32]. Following the public release of the *Planck* polarization data, a joint analysis by the BICEP2 and *Planck* teams using both data sets found that this was indeed the case. The analysis showed a strong correlation between the signal in the BICEP2 and *Keck Array* 150 GHz maps, and the signal in *Planck* maps of the same field at 353 GHz, where the dust emission dominates. While there was a residual $B$-mode amplitude in the BICEP2 data after accounting for the *Planck*-measured dust foreground level, it was not nearly strong enough to constitute a detection. The revised result was an upper limit of $r < 0.12$ at 95% confidence [33].

Recently, experiments have detected the expected $B$-mode signal at arcminute scales resulting from $E$-modes that have been gravitationally-lensed by intervening structure in the universe. This signal was first detected by SPTpol in cross-correlation with a template of the lensing potential, which was constructed from *Herschel*-SPIRE maps of the Cosmic Infrared Background (CIB) [34]. Lensing $B$-modes were subsequently detected by POLARBEAR through direct measurement of the $C_\ell^{BB}$ band powers in four multipole bins [35]. A detection of lensing also resulted from the joint BICEP2/*Keck* and *Planck* analysis mentioned above.

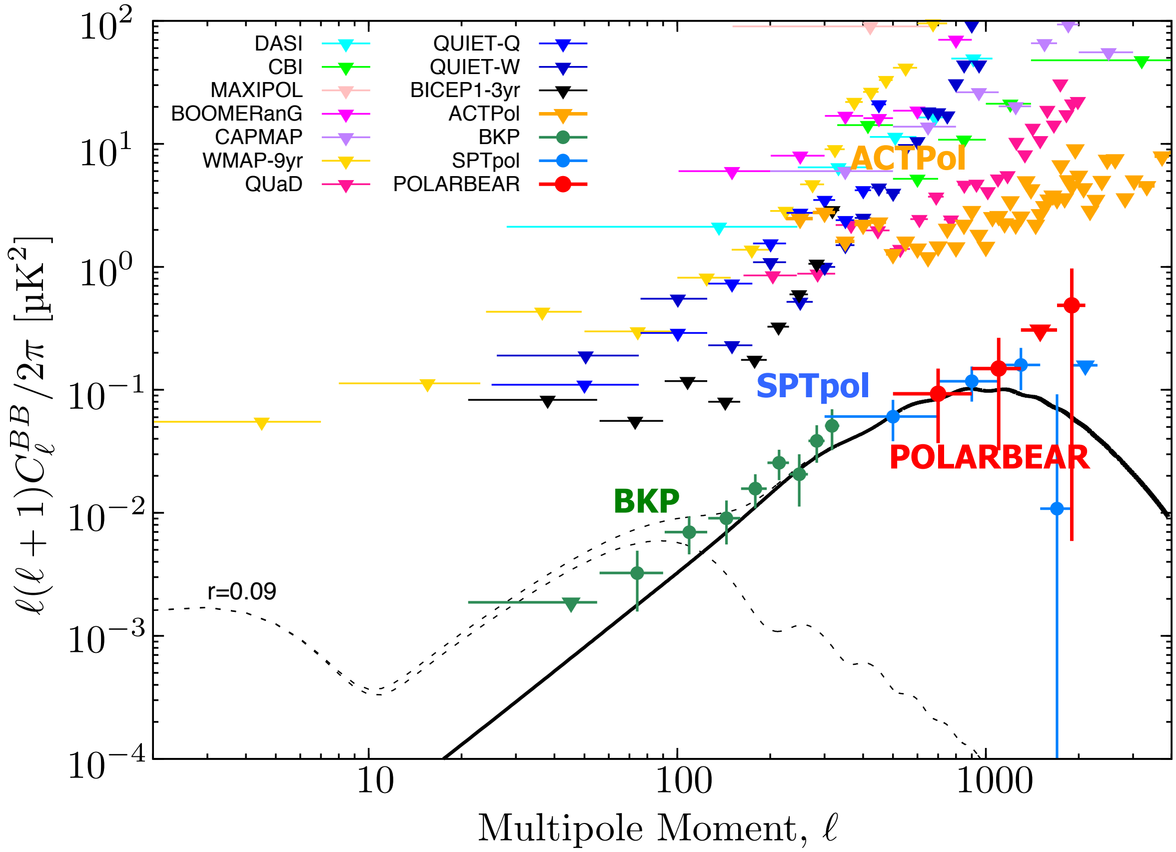


Figure 1.4: Upper limits on $C_\ell^{BB}$ band powers from previous and ongoing CMB polarization experiments. The measured band powers from BICEP2/*Keck*+*Planck* (BKP), POLARBEAR, and SPTpol are also shown. The dashed lines show the primordial $B$-mode power spectrum for $r = 0.09$, along with the combined primordial + lensing spectrum. The solid line shows the lensing spectrum alone. This figure is from [36].

### 1.3.3 Polarized Foregrounds

A major hindrance to measurements of CMB polarization is the presence of in-band foreground emission from within our own Galaxy. These foregrounds dominate the sky signal. The two expected sources of polarized foreground emission are synchrotron radiation, and thermal emission from dust grains in the interstellar medium that are aligned with the local magnetic field. Synchrotron radiation is radio emission arising from relativistic charged particles spiralling around magnetic field lines, and it is expected to be brighter at lower frequencies, relative to 100 GHz. Dust emission is thermal radiation from 10–20 K dust, and is expected to increase towards higher frequencies. The polarization arises because the long axes of the dust grains can become aligned, and the grains have a higher emission cross-section along that axis. More details are given in Section 3.1.1.

Synchrotron emission is modelled as a power law with frequency: $I_\nu \sim \nu^{-\alpha}$. Dust emission is modelled using a greybody spectrum of the form $I_\nu = (\nu/\nu_0)^\beta B_\nu(T)$ where $\beta$ is the power law spectral index for the dust emissivity, and $B_\nu$ is the Planck function. Multiple dust components at different temperatures and with different values of $\beta$ are sometimes considered.

Until recently, the amplitude of *polarized* foregrounds was not well-characterized, especially for dust emission. The best data set to date is from *Planck*, which has mapped the whole sky in polarization at 353 GHz and higher frequencies. This has produced a measurement of $\ell(\ell+1)C_\ell^{BB}/2\pi$ for the diffuse dust emission at intermediate Galactic latitudes (Figure 2 of [37]). Another *Planck* analysis has found a mean dust temperature of $T_d = 19.6$ K, and mean spectral index (in polarization) of $\beta = 1.59 \pm 0.02$, also at intermediate Galactic latitudes [38].

## 1.4 Experimental Challenges

Based on the previous sections, a number of experimental challenges to the detection of primordial $B$-modes can be identified. The CMB temperature fluctuations are at a level of one part in $10^5$. The amplitude of the $B$-mode fluctuations is 2 to 3 orders of magnitude lower than this, at the peak scale of $\ell \sim 80$, for $r \sim 0.1$ (Figures 1.3 and

1.4). Therefore, significant improvements in instrumental sensitivity are required in order to measure primordial $B$-modes. A second challenge is the contamination of the primordial spectrum by lensed $B$-modes at smaller angular scales. Third, as explained in the previous section, polarized emission from Galactic foregrounds can obscure the CMB signal. Multi-frequency coverage in the field of observation is required in order to adequately model and subtract the foreground contribution. Lastly, in order to achieve the polarimetric precision required, instrumental systematics that could produce a spurious polarization signal ($T \to E/B$ or $E \to B$) must be eliminated or lessened severely.

The Spider experiment is a balloon-borne microwave polarimeter whose experimental design attempts to address these challenges. Observing from a stratospheric balloon at 36 km altitude, above 99% of the air mass, allows for much greater sensitivity than ground-based experiments, which experience in-band loading from atmospheric emission. Photon-noise-limited bolometric detectors are used, meaning that overall instrumental sensitivity can be increased with the addition of more detectors. A lithographic fabrication technique has allowed the creation of large-format arrays of polarization-sensitive bolometers. These arrays are located in the focal planes of compact, monochromatic, refractive microwave telescopes, of the design first developed for BICEP. The design is inherently scalable: *Keck Array* [39] uses five telescopes of this type, while Spider flew six, for a total of 2400 bolometers.

The telescopes have small (∼0.3 m) apertures, so that they are sensitive to degree angular scales and larger ($10 \lesssim \ell \lesssim 300$) where the primordial component of the signal dominates over lensing. The small aperture and compact design allows for mid-to-far-field characterization of the instrument in the laboratory. The optical system is on-axis with axially-symmetric elements. A stepped half-wave plate at each telescope aperture rotates the sky polarization, allowing instrumental systematic effects to be identified and removed. The optics and detectors are described in more detail in Section 2.1. More details of systematics that could affect Spider, and the results of simulations to determine their minimum acceptable levels, are in [40], [41], and [42].

A suborbital experiment also allows for larger sky coverage than ground-based experiments. Lower $1/f$ noise from atmospheric drift means that less aggressive filtering of bolometer timestreams is required, so that larger spatial modes survive in the resulting maps. SPIDER was also designed to scan as widely and as quickly as possible (Section 2.8.4). Larger sky coverage means more uncorrelated $\ell$-bins, and less sensitivity to the selection of a particularly "clean" (foreground-free) patch of sky. SPIDER was designed to cover approximately 10% of the sky. A map-based maximum-likelihood simulation indicates that for $r \sim 0.1$, this sky fraction produces the tightest constraint on $r$, assuming a fixed integration time [43]. Finally, a suborbital experiment has access to higher frequency bands that cannot be observed from the ground due to atmosphere. The proposed SPIDER configuration always included a $> 200$ GHz channel for measuring dust emission, but detector arrays at these frequencies were not available for the first SPIDER flight. A dust channel will be included in a second flight. The combination of all of the above experimental design features was intended to allow SPIDER to set an upper limit of $r < 0.03$ at $3\sigma$.

## 1.5 Scientific Ballooning

Antarctica (or "The Ice"[3]) is a particularly good location for launching high-altitude balloons carrying experiments in astrophysics, particle physics, and other fields. The advantages of Antarctic ballooning include:

1. Flights can occur over completely uninhabited areas.

2. In the austral summer, the stratospheric winds are circumpolar, swirling counter-clockwise around the South Pole [44]. This wind pattern allows balloons to be

[3] As it is colloquially referred to by those who have been deployed there as part of the United States Antarctic Program

launched, circle the continent in two to three weeks, and return to a location that is convenient for recovery.

3. In the austral summer, the sun remains above the horizon continuously. Therefore:
   (a) There is a continuous source of power for experiments.
   (b) The lack of large diurnal temperature variations means that the balloon altitude is relatively stable (Figure 2.43), allowing it to stay afloat longer.

Regarding 3b: the high-altitude helium balloons that are used are equal-pressure balloons. Their final altitude (referred to herein as "float altitude") occurs when their internal pressure is equal to the ambient pressure. At that point, any excess helium (used to provide additional lift during ascent) is vented out of openings. A day-night cycle would cause the balloon to contract and sink, losing gas in the process. This effect can only be mitigated by carrying ballast: extra mass to be dropped intermittently throughout a flight.

Items 2 and 3 above combine to enable what is known as Long-Duration Ballooning (LDB). The record duration for an LDB flight is 55 days, while the average duration is 20.1 days, as determined from an analysis of 39 LDB flights since 1990, in [45]. An empirically-determined probability density function of Antarctic flight paths is also shown there (Figure 2.5 of [45]). For experiments such as Spider, in which the duration of observing is limited by the liquid cryogens that cool the receivers, a flight duration of 15–25 days is ideal, corresponding to roughly one circuit around the continent.

The LDB program is overseen by the NASA Balloon Program Office, associated with Wallops Flight Facility. However, the actual day-to-day operation of the program is conducted at the NASA Columbia Scientific Balloon Facility (CSBF) located in Palestine, Texas. Teams of scientists provide the experimental payloads to be carried by the balloons, while CSBF provides the balloons themselves, the launch vehicles and facilities, and support personnel.

The balloons employed by CSBF are made from polyethylene sheets that are 20 µm thick. They can be quite large, with a typical[4] fully-expanded size being 140 m by 120 m—the shape is somewhat oblate [46]. The flight train, the steel cabling and parachute connecting the payload to the balloon, is approximately 60 m long. SPIDER was flown on a “34-heavy” balloon, a designation referring to the $34 \times 10^6$ ft$^3$ ($0.97 \times 10^6$ m$^3$) volume displaced by this balloon at float altitude, and to the use of a thicker material compared to other CSBF balloons. The maximum suspended weight for a 34-heavy is 8000 lbs (a mass of 3600 kg). At that suspended weight, the maximum altitude is 120 000 ft (36 km): see NASA Standard Design Balloon Load/Altitude Curves in [46].

The Antarctic LDB launch site is located in the general vicinity of Williams Field, an airfield on the McMurdo Ice Shelf, approximately 10 km from the US-operated McMurdo Station. A balloon launch requires that the wind speed be low (at most 10 knots) and that the wind directions match between surface-level winds and low-level winds, a few hundred feet up. During most of ascent, which takes 3 to 4 hours, the pointing motors cannot apply enough torque to control payload attitude, due to the atmosphere. The SPIDER pointing motors were not powered on and tested until the balloon had reached 24 km altitude. The air temperature can drop as low as $-60$°C in the troposphere, before rising again [47]. Therefore, careful thermal modelling must be done to ensure that components do not freeze or overheat, depending on their thermal environment [47]. At float altitude, only radiative cooling is possible.

For the first 24 to 36 hours of the flight, the payload is within line of sight (LOS), and therefore 1 Mbit/s communication is possible by biphase-modulated radio transmission. Once the payload goes over the horizon, communication is only possible over the Iridium or TDRSS satellite networks, at rates of a few kbit/s. Commanding is not always reliable. These conditions require that the experiment can operate autonomously. CSBF provides a Support Instrumentation Package (SIP) that enables tracking and telemetry. At the

[4]These dimensions are for a 39-light balloon.

end of the flight, a signal is sent to a set of electronics in the flight train called the Universal Terminate Package, which triggers a mechanism that physically severs the flight train, detaching the payload from the balloon. The payload is caught by the parachute, and descends to the ground. CSBF requires that payloads be designed to withstand accelerations of up to $10g$ during "chute-shock", although in practice a rip-stitch in the flight train makes the acceleration not nearly this high. Structural modelling ensured that Spider was able to withstand both launch and chute-shock [47]. Details of the Spider 2015 and BLASTPol 2012 flights are in Section 2.11 and Section 3.6 respectively.

An exciting recent technological advance is the development of enclosed super-pressure balloons, which are capable of remaining at stable altitudes through day-night cycles [46]. These balloons are still being tested, but it is hoped that they will enable so-called Ultra-Long-Duration Balloon (ULDB) flights. These would be flights of up to 100 days, and could occur at mid-latitudes.

## 1.6 Aside: Polarization Formalism

Throughout this text, the Stokes parameters $I$, $Q$, $U$, and $V$ are referenced. These conveniently describe the polarization state of light in terms of observable quantities. Specifically, the parameters can be defined operationally in terms of the irradiances (in $\mathrm{W{\cdot}m^{-2}}$) that would be measured if light were incident on a set of ideal filters [48]. This definition is in Table 1.2. When presented with completely unpolarized light, each of these filters transmits half of the total incident irradiance.

Table 1.2: An operational definition of the Stokes parameters

| Filter Type | Transmitted Irradiance | Stokes Parameter |
|---|---|---|
| Passes all polarization states equally | $I_0$ | $I \equiv 2I_0$ |
| Linear polarizer[a] at 0° | $I_1$ | $Q \equiv 2I_1 - 2I_0$ |
| Linear polarizer[a] at 45° | $I_2$ | $U \equiv 2I_2 - 2I_0$ |
| Right-handed circular polarizer | $I_3$ | $V \equiv 2I_3 - 2I_0$ |

[a]The angle of the transmission axis is measured counterclockwise from the horiziontal.

Thus, $I$ is the total irradiance, and completely unpolarized light would have $Q = U = V = 0$. A pure $Q$ state corresponds to linearly-polarized light oriented horizontally ($Q > 0$) or vertically ($Q < 0$). A pure $U$ state corresponds to linearly-polarized light oriented at $+45^\circ$ ($U > 0$) or $-45^\circ$ ($U < 0$). A pure $V$ state corresponds to circularly-polarized light with right-handedness ($V > 0$) or left-handedness ($V < 0$). Incoherent light can be partially-polarized, with $Q^2 + U^2 + V^2 \leq I^2$. The parameters are often expressed as components of a Stokes vector $\mathbf{S} = (I, Q, U, V)$.

In terms of the complex electric field vector $\mathbf{E}(t) = \hat{\mathbf{x}}E_x(t) + \hat{\mathbf{y}}E_y(t)$, it can be shown that [48]:

$$I = \langle |E_x|^2 \rangle + \langle |E_y|^2 \rangle \tag{1.35}$$

$$Q = \langle |E_x|^2 \rangle - \langle |E_y|^2 \rangle \tag{1.36}$$

$$U = 2\langle \mathrm{Re}(E_x E_y^*) \rangle \tag{1.37}$$

$$V = -2\langle \mathrm{Im}(E_x E_y^*) \rangle \tag{1.38}$$

where the angle brackets denote the time average, and $^*$ denotes the complex conjugate.

The experiments presented in this thesis measure the partial linear polarization (if present) of light from astrophysical sources. The amplitude and direction of this linear polarization can be expressed as a complex value $Q + iU = Pe^{i2\phi}$ (see also Section 5.6).

From the preceding discussion, it is clear that the Stokes parameters are coordinate-system-dependent. In a slight difference from the description above, the astronomical convention is used henceforth, in which the linear polarization directions on the sky are expressed in equatorial coordinates. Under this convention, ($Q > 0,\ U = 0$) is oriented purely north–south, ($Q < 0,\ U = 0$) is oriented purely east–west, ($Q = 0,\ U > 0$) is oriented 45° counterclockwise from north–south, and ($Q = 0,\ U < 0$) is oriented 45° clockwise from north–south. In other words, the polarization angle $\phi$ increases counterclockwise from north–south.

## 1.7 Thesis Work

Many of my direct hardware contributions to the SPIDER and BLASTPol experiments are described herein. These include the development of the flight power systems for both SPIDER and BLASTPol (Sections 2.5 and 3.5). As part of this task, I designed a power switching and current sensing electronics board (Sections 2.5.4 and A.1) which was flown on BLASTPol in both 2010 and 2012. Two instances of it were operational on SPIDER during its 2015 LDB flight. I also developed and tested the actuators and control software for an Azimuth-Elevation (Az-El) scanning mount[5]. The mount was designed to support a test cryostat containing one SPIDER/BICEP2/*Keck*-style refractive microwave telescope (described in Section 2.1.1). This apparatus has been used successfully for mid-field beam mapping at Caltech by both the SPIDER and *Keck* collaborations.

One of my most significant areas of responsibility was the development, integration, and testing of all aspects of the SPIDER pointing control system (Sections 2.6 through 2.10, and [49]). In parallel with the control software implementing the flight scan mode, I developed software that simulates the behaviour of the payload during a scan. I put this simulator to a variety of uses, including debugging the scan control algorithm, investigating available sky coverage and cross-linking, and developing an overall observing

[5]For brevity, the Az-El mount is not described here. The conceptually similar, but much more difficult development of another automated pointable platform—the SPIDER gondola—is already discussed at length.

strategy for SPIDER (Section 2.9).

I participated in two field integration campaigns at CSBF in Palestine, Texas. The first, for BLASTPol, took place in June and July 2010, during which I field-tested the solar arrays and flight battery charging system. The second was for SPIDER from June through August 2013, during which I did a wide variety of work integrating the experiment, including tests of the pointing system. Prior to this integration campaign, I spent much of the period between October 2012 and March 2013 at Princeton University, helping to integrate the warm readout electronics (BLASTbus) system. The deployment of the SPIDER instrument team to McMurdo Station, Antarctica for the 2013 LDB flight of SPIDER was postponed by a year due to the 2013 US government shutdown. I was finally deployed to Antarctica between October 27, 2014, and January 17, 2015 to participate in the field campaign for SPIDER's balloon flight.

Another area in which I've contributed is analysis of data from the 2012 flight of BLASTPol. Chapter 4 describes the work I did to characterize BLASTPol beam non-idealities. Chapter 5 contains my analysis of the polarization spectrum of the Carina Nebula, investigating the dependence on environment of dust grain fractional alignment in molecular clouds.

# Chapter 2

# The Spider Experiment

## 2.1 The Optics and Detectors

### 2.1.1 The Receivers

The Spider experiment consists of six monochromatic, millimetre-wavelength telescopes: the receivers [50, 51]. In the 2015 LDB flight configuration, there were three receivers with a band centre of 150 GHz, and three at 94 GHz. The receivers are also referred to as *optical inserts*, or just inserts, because they are designed to fit entirely within cylindrical ports in a large liquid helium cryostat (Section 2.2). The design of the receivers is based on the refractive telescope design used for the BICEP and BICEP2 instruments [52, 53]. Figure 2.1 shows a cutaway diagram of one of the inserts. The telescope aperture diameter is 25 cm, and its field of view (FOV) is 20° across the diagonal. The simple two-lens telecentric design ensures a flat focal plane geometry. The objective and eyepiece lenses are made from high-density polyethylene (HDPE) with a porous PTFE anti-reflection (AR) coating. The lens separation is 500 mm, and the effective focal length of the system is 583.5 mm, for a plate scale of 0.98°/cm.

The optical elements are cooled to 4 K to prevent excess in-band loading due to dielectric loss. The insert baseplate is thermally-connected to the 4 K cooling stage of

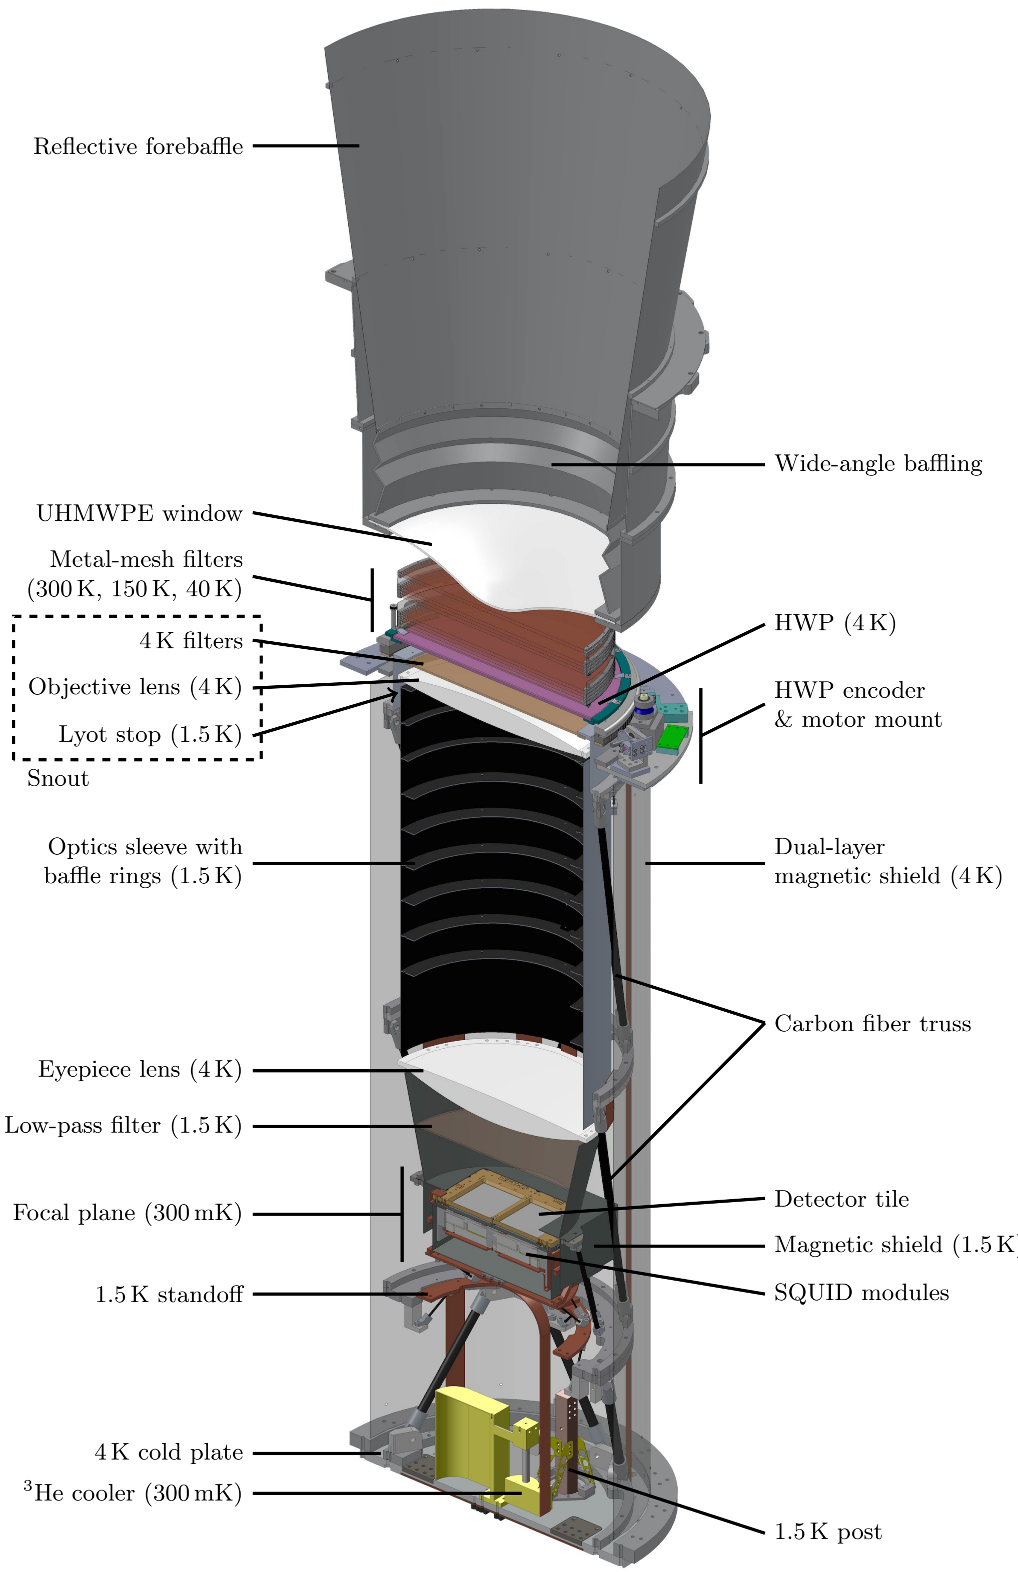


Figure 2.1: Cutaway diagram of a SPIDER optical insert. This is Figure 3 from [51].

the experiment. A $^3$He adsorption fridge (Section 2.2.3), for cooling the focal plane to 300 mK, is mounted to this baseplate. The focal plane unit (Section 2.1.2) is suspended above the baseplate by a rigid carbon fibre structure: the camera truss. The focal plane is surrounded by an Amuneal® Amumetal-4K magnetic shield that is cooled to 1.6 K (also referred to as the "spittoon"). The lenses are supported by another carbon fibre structure: the optics truss. The rigidity and low coefficient of thermal expansion of the carbon fibre prevents beam distortion. In between the eyepiece and objective lenses, a blackened sleeve with ring-shaped baffles prevents the detector beam power from spilling out onto warmer stages of the system. Since a significant fraction of the beam sidelobe power terminates on this sleeve, it is also thermally-connected to the 1.6 K cooling stage (Section 2.2.2). The top baffle ring of the sleeve provides a beam-defining Lyot stop, just behind the objective lens. A rotating half-wave plate (HWP, Section 2.1.3) is mounted in front of the objective. The vacuum window is a piece of ultra-high-molecular-weight polyethylene (UHMWPE) of 0.125" thickness. The window mounts to the base of a recessed bucket connected to a reflective forebaffle that prevents stray rays from entering the optical system.

A series of metal-mesh filters [54] and nylon dielectric filters at various stages in the optical system are used to attenuate infrared (IR) loading. These begin at 300 K at the vacuum window, and include filter stacks at the $\sim$120 K and $\sim$30 K radiation shields (Section 2.2.1), at 4 K at the insert "snout" (just above the objective), and a final metal-mesh low-pass filter at 1.6 K supported within the spittoon, just above the focal plane.

In summary, the Spider optics are optimized for degree-scale angular resolution on the sky, with small, easily-baffled apertures and a compact, on-axis design that can be entirely cooled. The design allows the HWP to be placed skyward of the primary optic, where it can modulate the polarization of the sky signal without affecting instrumental polarization systematics.

### 2.1.2 The Detectors

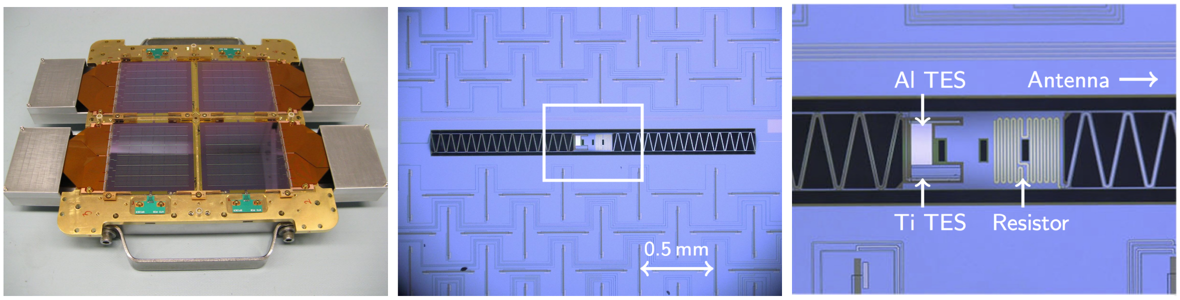


Figure 2.2: Images of a SPIDER detector array. *Left* – A photograph of the underside of a 150 GHz focal plane unit, showing the four Si tiles, each with a lithographed array of spatial pixels. *Middle* – A close-up of one of the spatial pixels, showing the arrays of slot dipole antennas, the microstrip lines, and the meandered SiN legs leading to one of the two bolometer islands. *Right* – A close-up of a bolometer island, showing the meandered Au resistor, the Ti TES, and the Al TES

Figure 2.2 shows images of a SPIDER detector array [50, 51, 55] at various scales. The focal plane unit consists of four silicon (Si) tiles, each with a lithographed array of spatial pixels: $8 \times 8$ pixels at 150 GHz and $6 \times 6$ pixels at 94 GHz. The lithographic process combines beam-defining and band-defining elements on a single planar surface, eliminating the need for external feed horns. Each spatial pixel consists of two interleaved arrays of slot dipole antennas, one sensitive to horizontal polarization, and one sensitive to vertical polarization. Radiation from the antennas is summed coherently by superconducting niobium (Nb) microstrip lines, the sum determining the overall antenna beam pattern. Filters in-line with the microstrips define the detector bandpass.

The antenna-coupled detectors are superconducting transition-edge sensor (TES) bolometers [56]. Power from each antenna array passes through the microstrip lines to a bolometer island suspended above the tile substrate by meandered silicon nitride (SiN) legs. Hence there are two bolometer islands per spatial pixel: one for each of the two orthogonal polarization states. The optical power is dissipated in a meandered gold (Au) resistor, heating up the island. The change in temperature causes a change in resistance of the TESs: thin films of titanium (Ti) and aluminum (Al) deposited on the SiN island substrate. During operations, one of these TESs is electrically-biased to operate at the transition between superconducting and normal resistance, resulting in a sharp, linear response $R(T)$. The Ti TES, which is used for in-flight observations, has a superconducting transition temperature $T_c \approx 500$ mK and a normal resistance $R_n \approx 30$ m$\Omega$. The Al TES has $T_c \approx 1.1$ K and $R_n \approx 0.1$ $\Omega$. Taking advantage of the lower optical loading in the stratosphere, the meandered leg design produces a very low[1] thermal conductance of $G \approx 20$ pW/K between the island and the 300 mK tile. The low $G$ allows the Spider arrays to have very low noise levels, at the cost of a low saturation power of 2 to 3 pW on the Ti transition. This value results in saturation under laboratory loading conditions. Therefore the receivers are biased on the Al transition during testing, allowing the instrument to be fully-characterized on the ground.

The TES bolometers are voltage-biased, resulting in negative electrothermal feedback. Increased optical power on the bolometers increases their temperature, in turn increasing the TES resistance, and hence reducing the electrical bias power dissipated in the device. The changing electrical bias current needed to keep the TES on transition is read out by a three-stage system of superconducting quantum interference devices (SQUIDs). These devices are highly-sensitive to changes in magnetic flux, and can be used for precise, low-noise current measurement. A schematic of the TES/SQUID bias and readout circuit is in Figure 1 of [57]. The system employs time-domain multiplexing. Each bolometer has its own first stage SQUID (SQ1), to which the bolometer signal is inductively-coupled. The SQ1 signal in turn is inductively-coupled to a summing coil that carries the SQ1 signals for all the detectors in a multiplexer column to a second stage SQUID (SQ2). The biases for each SQ1 (the row select lines) are turned on sequentially, meaning that the

[1] Relative to receivers of ground-based counterparts such as BICEP2, that receive much higher optical loading, and hence are designed to have higher-$G$ bolometers.

SQ2 only measures the signal from one bolometer in the column at a time. In this way, a given multiplexer (Mux) chip reads out the signals from 33 channels (32 bolometers + 1 dark SQUID channel to measure noise and magnetic pickup). The signals from the Mux chips then propagate to separate SQUID Series Array (SSA) modules, with 100 SQUIDs in series per channel, for amplification. Four Mux chips connect together on a Mux board, enough to read out one detector tile. There are four such boards per focal plane unit.

A major change from the BICEP2 focal plane architecture was made to accomodate the significant magnetic shielding required, given SPIDER's motion through Earth's magnetic field. As shown in the left panel of Figure 2.2, and in Figure 2.1, the Mux boards, encased in niobium, connect to the tiles through flexible aluminum superconducting circuits. These flexi-circuits allow the boards to be folded underneath the focal plane, where they are further shielded inside a niobium box. This version of the focal plane is known as Rev. X. For this reason, the individual SPIDER receivers are referred to as X1 through X6 herein, with odd-numbered receivers at 150 GHz and even-numbered receivers at 94 GHz.

### 2.1.3 The Half-Wave Plates

The SPIDER HWPs are described in great detail in [58] and [59]. The HWPs modulate the polarization of the sky signal, allowing each detector to independently measure the $Q$ and $U$ Stokes parameters for linear polarization (Section 1.6). Each HWP was a 330 mm diameter disc of sapphire: a birefringent material. The 150 GHz HWPs used an AR coating of Cirlex® polyamide material. The 94 GHz HWPs used a quartz AR coating. Each HWP was rotated by a cryogenic mechanism. This device consisted of a rotor mounted in a three-point roller bearing. The bearing is rotated by a worm gear driven by a stepper motor. The entire system operates at 4 K, including the motors, rather than using a shaft feed-through. Purpose-built optical encoders read out the rotor angle. The HWPs were periodically stepped through a sequence of angles during the flight (Section 2.9.1).

## 2.2 The Cryogenic System

The Spider optical inserts are completely enclosed within and cooled by a liquid helium (LHe) cryostat (Figure 2.3). The cryostat has a dry mass of 850 kg. The outer part of the cryostat is the vacuum vessel, an aluminum cylinder 2.11 m in diameter by 2.43 m in height. The vacuum vessel encloses vapour-cooled shields (VCSs), the LHe main tank, and a pumped superfluid tank (SFT). The basic operation of these elements is summarized in the following sections, with many more details given in [60] and [45].

### 2.2.1 The Main Tank

The main tank is an aluminum 5083 cylinder interrupted by cavities (that are contiguous with the vacuum space) for the telescope inserts. The tank interior has the capacity for 1284 L of $^4$He. The main tank provides a 4 K cooling stage for the experiment. The tank is suspended within the vacuum vessel by G-10 flexures. Surrounding the main tank are inner and outer vapour-cooled shields: VCS1 and VCS2 respectively. The flow of vapour from the boil off of the liquid cryogens passes through vent lines along which there are heat exchanger blocks, thermally-sunk to each VCS. Passage of the gas through the flow-restrictive heat exchangers cools the VCSs. The system thus regulates itself by negative feedback: increased thermal loading on the main tank increases the flow rate of boil off to the VCSs, increasing their cooling power, and reducing the loading. In flight, the temperatures of VCS1 and VCS2 were approximately 30 K and 118 K respectively. In between and outside of the VCSs, further radiative shielding is provided by multi-layer insulation (MLI).

The design goal of the cryogenic system was a 25-day hold time, to exceed the average

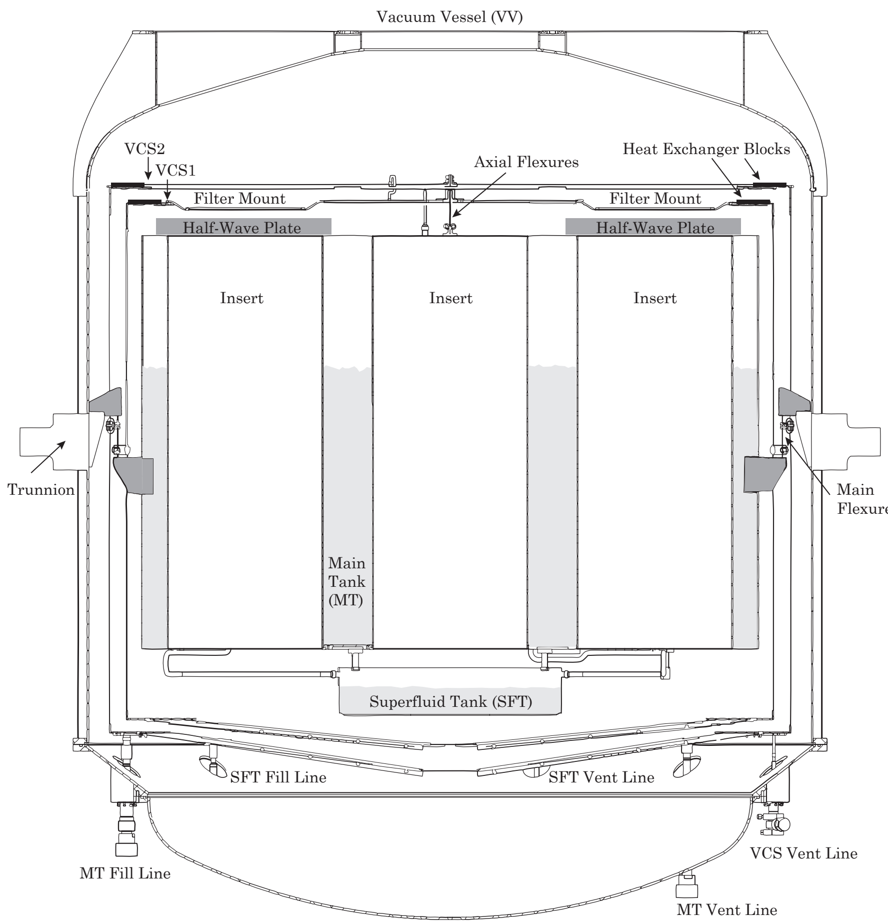


Figure 2.3: Cross-sectional view of the SPIDER cryostat. This is Figure 2 from [60].

LDB flight duration. The cryogenic performance measured in the days prior to launch indicated a 12-day hold time. However, the reduction in radiative loading at float altitude was expected to increase the hold time significantly, and the cryogenic hold time realized in flight was approximately 16 days.

### 2.2.2 The Superfluid Tank

Liquid from the main tank passes through a system of capillaries into a 16 L auxiliary helium tank: the SFT. The flow impedance of the capillaries determines their cooling power. During laboratory operations, the SFT volume is pumped down to between 2 and 5 Torr (267 and 667 Pa). In flight, it was vented to the ambient environment, whose mean pressure was $4.7 \pm 0.3$ Torr. At these pressures, the $^4$He becomes superfluid and reaches a temperature of 1.6 K. A cooling stage at this temperature is necessary for the operation of the closed-cycle fridges within each insert.

### 2.2.3 Closed-Cycle $^3$He Adsorption Refrigerators

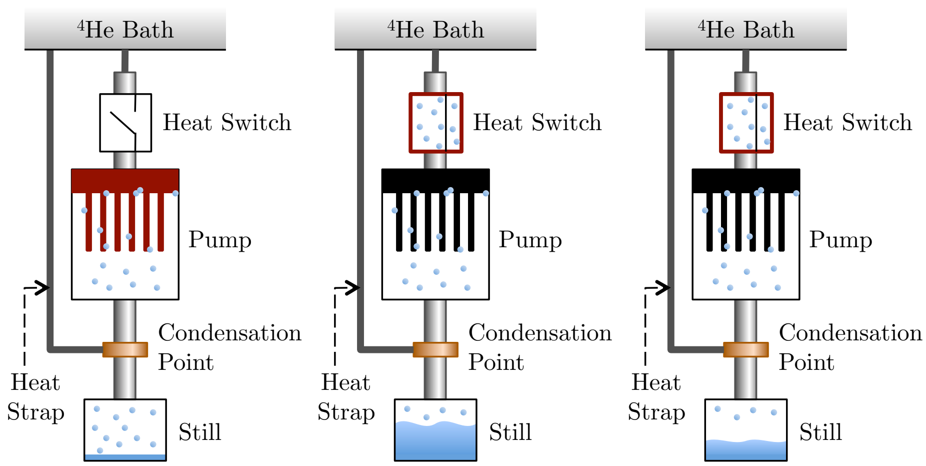


Figure 2.4: The steps in a SPIDER fridge cycle. *Left* – The pump is heated (red) releasing $^3$He atoms (blue) that condense and form liquid in the still. *Middle* – The heat switch (red) is closed, cooling the pump, which begins adsorbing $^3$He atoms. *Right* – The still begins to empty as the fridge provides cooling power and the $^3$He boils off.

The sub-kelvin temperatures necessary for operation of the detectors are reached using $^3$He adsorption refrigerators from Chase Research Cryogenics. One fridge unit is mounted at the base of each optical insert. The main components of a fridge are an activated charcoal pump, a pump heater, and a gas gap heat switch. During a fridge cycle, the pump heater is turned on long enough to release all the gaseous $^3$He adsorbed by the charcoal pump. The gas that is released forms liquid on the condensation point, which is connected to the 1.6 K bath of the SFT by a heat strap. The liquid drips down into a small vessel (the still). The pump heater is then turned off, and the heat switch closed. The heat switch thermally connects the pump to the 4 K stage, so that it cools and begins to adsorb helium atoms evaporating from the $^3$He bath. This pumping on the still lowers the vapour pressure, and hence the temperature of the liquid $^3$He. During operation, a still temperature of $\lesssim$ 300 mK can be achieved. The fridge units are designed to operate for 3–4 days between cycles. The steps of the fridge cycle are shown in Figure 2.4.

## 2.3 Computing and Data Acquisition

### 2.3.1 The Flight Logic Computers

Two Arcom/Eurotech Apollo (1.6 GHz Pentium M) single-board computers controlled all non-receiver-related operations on the experiment. These flight logic computers (FLCs) existed as a redundant pair. At any given time, one of them was in charge and writing data to the BLASTbus (Section 2.3.2). On SPIDER, the FLCs were named itsy and bitsy[2]. The FLCs were monitored by a Watchdog circuit [61]. The Watchdog, originally developed for BLAST, checks for the "tickle": a toggled digital signal on the parallel port of each FLC. If the tickle from a computer ceases, such as in the event of a cosmic ray strike causing the main control program to crash, the Watchdog reboots that computer and places the other computer in charge.

The computers used the Ubuntu 12.04 Linux operating system. Both FLCs ran pcm: multi-threaded C code that served as the SPIDER master control program. The architecture of pcm was based on that of mcp, the BLAST and BLASTPol master control program, which is described in great detail in [61].

[2]Collectively, this pair was referred to as "The Waterspout".

### 2.3.2 The BLASTbus

The FLCs interface to the rest of Spider's subsystems using the BLASTbus [62, 63], a custom synchronous RS-485 serial bus. The BLASTbus electronics were developed to provide general-purpose readout and control for balloon-borne experiments. On Spider, the BLASTbus system was used for everything except the bolometers, which were read out by the MCEs (Section 2.3.3). Each FLC had a PCI card, the BLASTbus controller, acting as a master node. The slave nodes are BLASTbus motherboards. A given BLASTbus readout crate would have a certain number of motherboards, depending on its specific application. Spider had two such crates. One was the Attitude Control System (ACS), which contained two BLASTbus motherboards used for fast motor control loops (Section 2.6.1.2), and for the readout and control of various pointing sensors, actuators, and ambient thermometry. The second BLASTbus crate was the Housekeeping Data Acquisition System (HK DAS) which, combined with an analogue preamplifier crate, provided for the biasing and readout of diodes and resistive temperature detectors (RTDs) in the cryostat. It also controlled the HWP motors.

### 2.3.3 The Multi-Channel Electronics

The three-stage time-domain multiplexed SQUID readout system described above is coupled to the ambient-temperature Multi-Channel Electronics (MCE) [57, 64]. The MCE system sets the detector biases, controls the multiplexers and SQUID amplifiers, and reads out the signals from the arrays. Each of the six independent Spider receivers had its own MCE readout crate. Each crate is connected by optical fibre to an MCE Control Computer (MCC). These Linux computers run the MCE Acquisition Software (MAS), and ultimately store the detector array data to disk. The MCE system includes a Sync Box, which generates a 5 MHz clock signal used to synchronize data acquisition amongst the six MCEs. This signal was used as the BLASTbus clock as well, to ease the synchronization of data acquired with the two different readout systems.

### 2.3.4 Data Storage

Past balloon flights raised some concerns about the reliability of solid-state drive (SSD) technology at float (Section 3.6). Therefore, SPIDER used both SSDs and spinning hard disk drives (HDDs) to store redundant copies of the data. Each FLC had one 500 gigabyte SSD and one 1 terabyte HDD. The SSDs were located in the flight computer box. Spinning drives require air to support the reader head above the platter, so the two FLC HDDs were placed inside a pressure vessel (PV).

The MCC data storage scheme was similar. Each MCC box contained two 500 gigabyte SSDs. Each MCC also had an additional 1 terabyte HDD. Four pressure vessels were used, divided into pairs. Each PV contained three drives. Therefore, each pair of two PVs had a full set (six receivers' worth) of data. One pair of PVs grouped the drives into the sets ( (X1, X3, X6), (X2, X4, X5) ), while the other pair of PVs used the grouping ( (X2, X3, X6), (X1, X4, X5) ).

## 2.4 Structural Overview

To support the massive flight cryostat while keeping the experiment below the mass limit for LDB flights, a light-weight carbon fibre[3] gondola structure was designed and built [47, 65]. The structure consisted of hollow carbon fibre tubes with aluminum inserts glued into the ends. The inserts had flanges with bolt circles, allowing them to be mounted to the planar faces of custom-machined aluminum joints. The geometry of the joints determined the angles at which the tubes met. The glue was 3M™ Scotch-Weld™ Epoxy Adhesive 2216 A/B Gray. Prior to the SPIDER flight, these materials and construction methods had not been used in a balloon-borne astrophysics experiment.

[3]More specifically: carbon-fibre-reinforced polymer

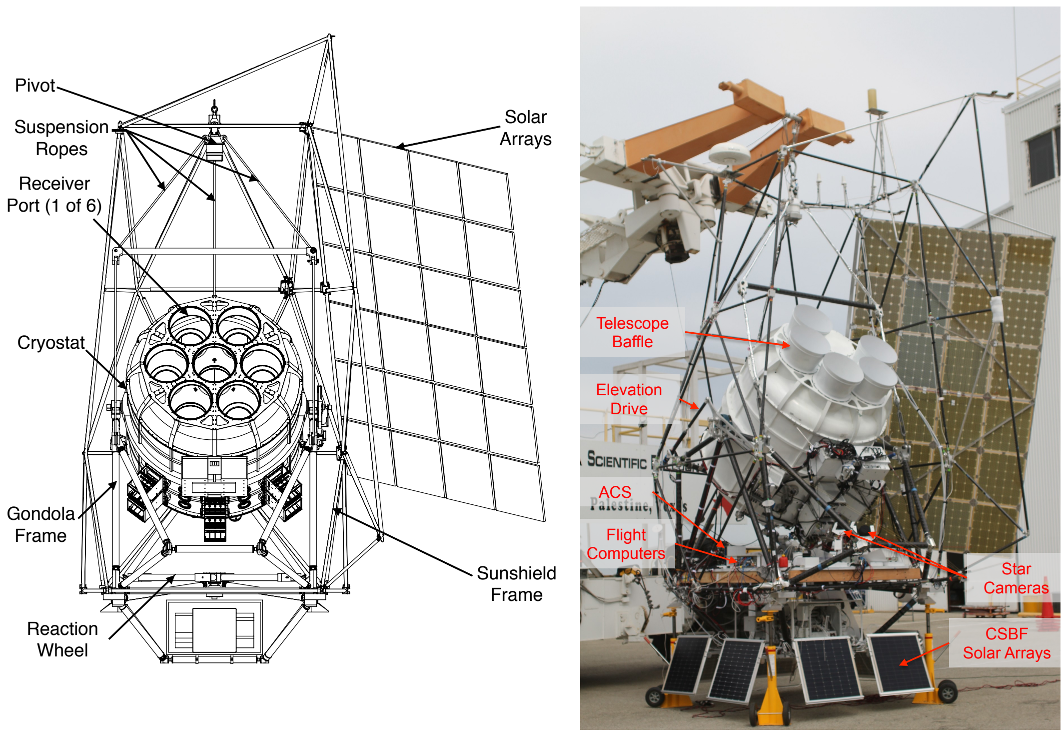


Figure 2.5: *Left* – A drawing of the Spider payload with major structural elements and components of the pointing system labelled. *Right* – A photograph of the assembled Spider payload hanging from the launch vehicle at the NASA Columbia Scientific Balloon Facility (CSBF) in Palestine, Texas.

Figure 2.5 shows a nearly fully-assembled SPIDER payload (see Figure 2.45 for the flight configuration). Following BLAST terminology, the payload can be divided functionally into the inner frame and outer frame. On SPIDER, the inner frame is the cryostat itself, along with everything mounted to it: the telescope baffles, bore-sight star camera (BSC), MCEs, housekeeping electronics, and one set of flight batteries. The outer frame (gondola) comprises the support structure itself, and an aluminum honeycomb deck to which the flight electronics are mounted. The deck includes the ACS, the FLCs, the MCCs and their power breakout, the Sync Box, the charge controllers and their relay boxes, the flight batteries, the elevation drive power supply, the PVs, the reaction wheel motor control box, the gyroscope box, the rotating star camera (RSC), and an RS-232 serial hub. Many of these subsystems are discussed in subsequent sections. The reaction wheel is also mounted to the outer frame, underneath the deck. Below that is a cage to house the Support Instrumentation Package (SIP): CSBF flight hardware used for payload tracking and telemetry.

The outer frame is suspended from three braided Technora® ropes. The aft rope connects directly from the gondola frame to the pivot (Section 2.6.2). The port and starboard fore ropes connect to a carbon fibre spreader bar, which reduces the horizontal component of the rope tension. Another set of ropes extends upward from there to the pivot.

The inner and outer frames are surrounded by the frame for the sunshield, which is also constructed from carbon fibre tubes with glued aluminum inserts connecting to aluminum joints. The sunshield frame uses tubes of a smaller outer diameter. For flight, the faces of the sunshield were covered with aluminized Mylar®. The port side of the sunshield frame included an extended wing, allowing SPIDER to scan closer to the sun in azimuth. The solar arrays (Section 2.5.1) mount to the port side of the sunshield frame as well.

## 2.5 The Power System

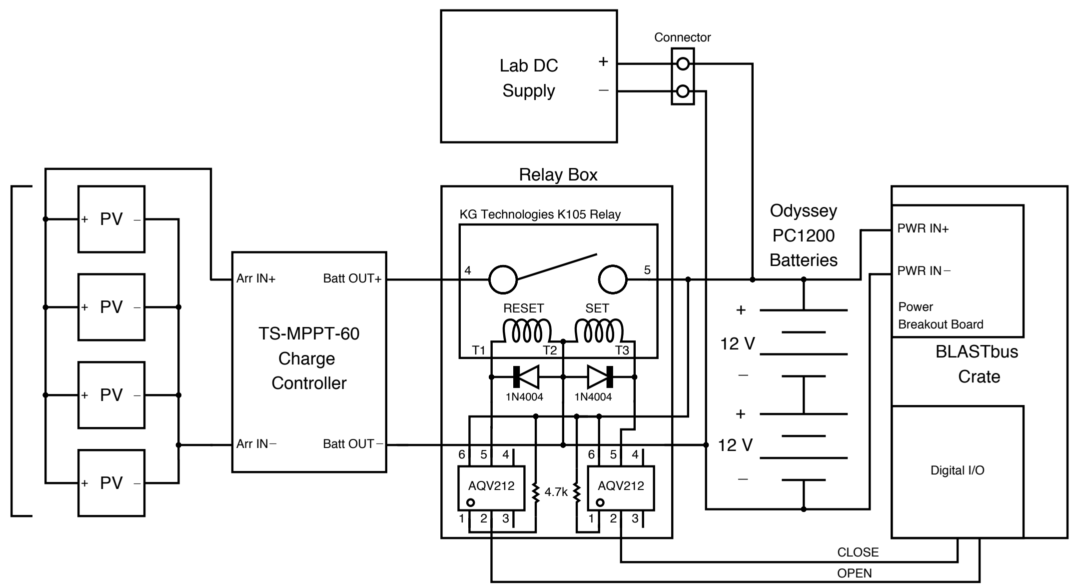


Figure 2.6: A schematic overview of a Spider power system. Two instances of this system were flown on the experiment.

Spider had two separate, identical power systems, one for the outer frame (gondola) systems, and the other for the inner frame (receiver and cryogenic) systems. This scheme allowed the inner and outer frame systems to be electrically-isolated. A schematic of one such power system is in Figure 2.6. As shown in the figure, the experiment is powered by solar arrays (Section 2.5.1). Their energy is stored by lead-acid batteries (Section 2.5.2) whose charging is regulated by an off-the-shelf charge controller (Section 2.5.3). The connection between the charge controller and batteries can be interrupted by the relay box, giving the ability to power cycle the charge controller in flight (Section 2.5.3.4).

For ground-based operation inside the high bay, where solar power is not available, a laboratory DC voltage supply is connected to the batteries in parallel with the charge controller. Inline connectors were used so that the two lab supplies could simply be disconnected prior to launch. Due to its length, the wiring between the lab supplies and batteries used 4 AWG welding wire. The rest of the power system wiring used 12 or

14 AWG stranded hook-up wire. All wiring on the experiment was PTFE-insulated, since there is risk of PVC insulation freezing and cracking at low temperatures. In each power system, the flight batteries connected to a Power Breakout Board (PBOB, Section 2.5.4), inside a BLASTbus electronics crate. The PBOBs collect all of the power switching and current sensing circuits for the experiment on two modular, custom PCBs. For the outer frame power system, the BLASTbus crate in question was the ACS, whereas for the inner frame power system, it was the HK DAS. A power budget for SPIDER based on in-flight measurements is in Section 2.5.5.1.

## 2.5.1 The Solar Arrays

### 2.5.1.1 Solar Cells and Panel Construction

The solar arrays use A-300 monocrystalline silicon solar cells manufactured by SunPower Corp®. Their characteristics are listed in Table 2.1.

Table 2.1: Electrical characteristics of a typical A-300 solar cell under standard testing conditions (STC): irradiance = 1000 W/m$^2$, $T$ = 25°C, and an AM 1.5G input solar simulation spectrum [66]

| Quantity | Value |
|---|---|
| Open Circuit Voltage [V] | 0.670 |
| Short Circuit Current [A] | 5.9 |
| Max Power Voltage [V] | 0.560 |
| Max Power Current [A] | 5.54 |
| Rated Power [W] | 3.1 |
| Efficiency [%] | Up to 21.5 |
| *Temperature Coefficients* | |
| Voltage [mV/°C] | −1.9 |
| Power [%/°C] | −0.38 |

These cells are assembled by SunCat Solar, LLC into panels constructed with balsa wood frames and a honeycomb backing material. Panels from this supplier have flown on many past balloon payloads, including BLASTPol. Figure 2.7 shows how the cells are grouped into larger units to make a full array. Within a SunCat panel, five cells are connected in series to make a string, and six strings are connected in series to complete the panel. Three panels are wired in series to form a module, or column. Given the power requirements of the experiment, as many columns as needed can be wired in parallel. In SPIDER, four columns were wired in parallel to form a $3 \times 4$ array of solar panels. These four columns are the four PV modules depicted in Figure 2.6. Two of these $3 \times 4$ arrays were flown: one for each power system.

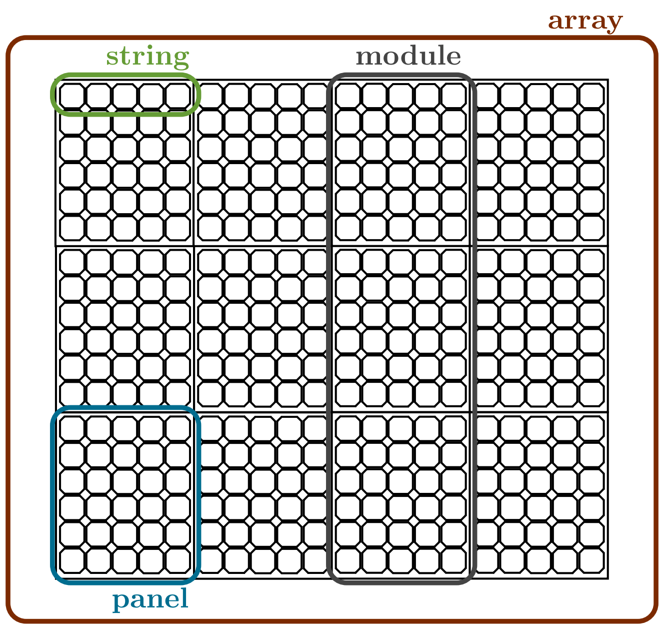


Figure 2.7: The arrangement of cells within a solar array

Given the information above, the peak power of an array at normal incidence can be computed.

$$\left(0.56\ \frac{\text{V}}{\text{cell}}\right)\left(90\ \frac{\text{cells}}{\text{module}}\right) = 50.4\ \frac{\text{V}}{\text{module}}$$

$$(50.4\ \text{V})\left(5.54\ \frac{\text{A}}{\text{module}}\right)(4\ \text{modules}) = 1.12\ \text{kW}$$

A more pessimistic scenario is to consider $T = 100$°C, for which the max power voltage is reduced to 0.384 V, resulting in 34.56 V per module. At this temperature, the max power current is 6.83 A, resulting in 971 W per array. These values also assume completely normal incidence. A more detailed analysis would account for the time-averaged incidence angle during a scan (Section 2.5.1.2). A higher solar irradiance is expected at float altitude, and therefore a higher peak power. Until SPIDER, the gain in power at float had not been measured, since past experiments such as BLASTPol had never drawn full power from the arrays for an extended period, not even during ascent. Section 2.5.5 presents the battery and solar array performance during the SPIDER 2015 flight.

#### 2.5.1.2 Optimal Array Angle

An aluminum frame was designed to mount the two $3 \times 4$ arrays of solar panels on the port side of the sunshield. During the design process, it was necessary to determine the optimal opening angle of the array frame from the bore-sight. This is the opening angle that maximizes the time-averaged incident power on the arrays during a scan. If $I_{\mathrm{solar}}$ is the solar irradiance at float, and $A$ is the total array area, then the time-averaged incident power is given by

$$\langle P \rangle = \frac{1}{T} \int_T A I_{\mathrm{solar}}(t) \cos(\eta(t))\, dt \tag{2.1}$$

Here, $\eta$ is the angle of incidence of the solar radiation on the arrays. In the simplified case where $I_{\mathrm{solar}}$ = constant, the average power is just $(AI_{\mathrm{solar}}/T) \int_T \cos(\eta(t))\, dt = AI_{\mathrm{solar}} \langle \cos \eta \rangle$. The goal, then, was to determine the array opening angle that maximizes $\langle \cos \eta \rangle$ given the typical SPIDER scan.

In a Cartesian coordinate system in which the array is in the $zy$–plane, the unit normal vector to the array is $\hat{\mathbf{x}}$. If $\hat{\mathbf{s}}$ is the direction vector to the sun, then $\cos \eta = \hat{\mathbf{s}} \cdot \hat{\mathbf{x}}$. In spherical coordinates, $\hat{\mathbf{s}}$ can be specified using an azimuthal angle $\alpha$ and an elevation

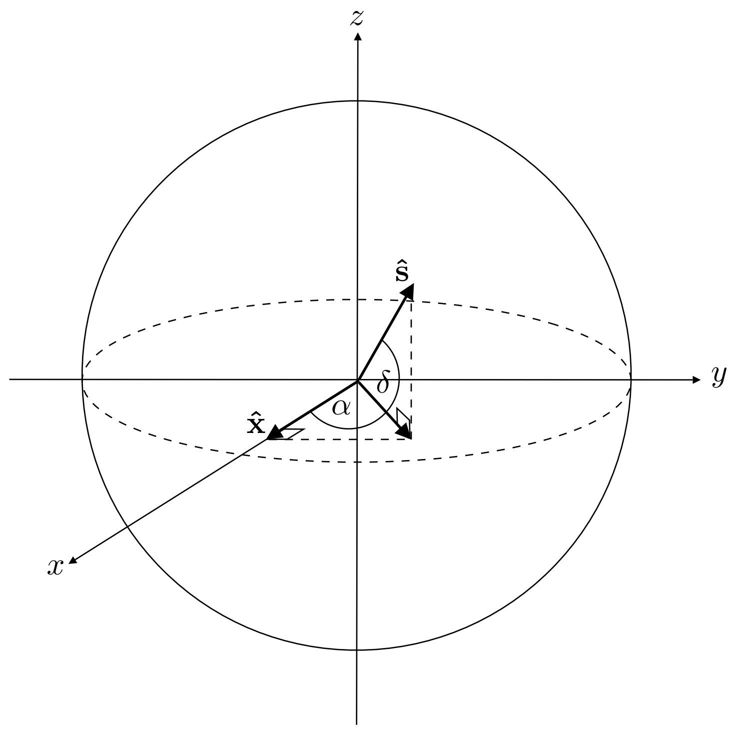


Figure 2.8: Sun angles $\alpha$ and $\delta$ relative to the array unit normal

angle $\delta$, measured from the $xy$–plane (Figure 2.8). We then have $\hat{\mathbf{s}} = (\cos\delta\cos\alpha)\hat{\mathbf{x}} + (\cos\delta\sin\alpha)\hat{\mathbf{y}} + (\sin\delta)\hat{\mathbf{z}}$. Therefore

$$\cos\eta = \hat{\mathbf{s}} \cdot \hat{\mathbf{x}} = \cos\delta\cos\alpha \tag{2.2}$$

A SPIDER bore-sight pointing simulator, `scan_sim` (Section 2.9) was run to determine $\alpha$, $\delta$, and hence $\cos\eta$ vs. time, over the course of 24 solar hours. These sun angles were computed as follows. The simulator produces the bore-sight azimuth $\phi(t)$. This can be expressed as an azimuth (Az) relative to the sun $\Delta\phi_{\mathrm{sun}} = \phi - \phi_{\mathrm{sun}}$ where $\phi_{\mathrm{sun}}$ is the sun Az, also computed by `scan_sim`. If the array opening angle from the bore-sight is $\Delta\phi_{\mathrm{array}}$, then the sun azimuth relative to the array normal is

$$\varphi = \Delta\phi_{\mathrm{sun}} - \Delta\phi_{\mathrm{array}} - 90^\circ \tag{2.3}$$

To see the geometry of Eq. 2.3, refer to Figure 2.29 in Section 2.8.1.1, which shows the specific case of $\varphi = 0^\circ$, $\Delta\phi_{\mathrm{sun}} = 144^\circ$, and $\Delta\phi_{\mathrm{array}} = 54^\circ$.

If the arrays were perfectly vertical, the problem would be solved, since $\alpha = \varphi$, and $\delta = \theta_{\text{sun}} = \vartheta$, the sun elevation, also computed by `scan_sim`. However, when mounted, the arrays were tilted by $\theta_{\text{array}} = 16°$ from vertical. This meant that $\varphi$ was not the azimuthal angle of the sun from the array normal. Rather, it was the azimuthal angle of the sun from the *projection* of the array normal onto the horizontal. An additional transformation to convert from horizon coordinates ($\varphi$, $\vartheta$) to "array coordinates" ($\alpha$, $\delta$) is required:

$$\sin\delta = \sin\vartheta\sin\Theta - \cos\vartheta\cos\Theta\cos\varphi \tag{2.4}$$

$$\sin\alpha = \frac{\sin\varphi\cos\vartheta}{\cos\delta} \tag{2.5}$$

$$\cos\alpha = \frac{\sin\vartheta - \sin\delta\sin\Theta}{\cos\delta\cos\Theta} \tag{2.6}$$

In Eqs. 2.4 through 2.6, $\Theta \equiv 90° - \theta_{\text{array}}$. These equations can be derived by applying the spherical sine and cosine rules to spherical triangle $PSZ$ in Figure 2.9. Position $P$ marks the pole of the array coordinate system, while $Z$ is the pole of the horizon coordinate system (i.e. the zenith). Position $S$ marks the coordinates of the sun.

To compute $\langle\cos\eta\rangle$, `scan_sim` was run over 24 hours, with typical values of the SPIDER scan parameters (Section 2.9). At each simulation time step, values of $\varphi$ and $\vartheta$ were computed as a function of opening angle $\Delta\phi_{\text{array}}$ over the range $0° < \Delta\phi_{\text{array}} < 90°$. These were then transformed to $\alpha$ and $\delta$, which were used to calculate $\cos\eta$ at each opening angle at each time step. The values of $\cos\eta$ for each opening angle were then integrated over the simulation. The result is shown in Figure 2.10. The maximum value of $\langle\cos\eta\rangle = 0.834$ occurs at an array angle of 31° (dashed grey lines). The original analysis assumed the arrays were tilted at $\theta_{\text{array}} = 20°$ from vertical, and obtained an optimal opening angle closer to 50°, possibly for a simulation scanning over a different range in right ascension. The final array frame construction was subject to practical constraints and ended up with an opening angle (projected onto the horizontal) of 54°

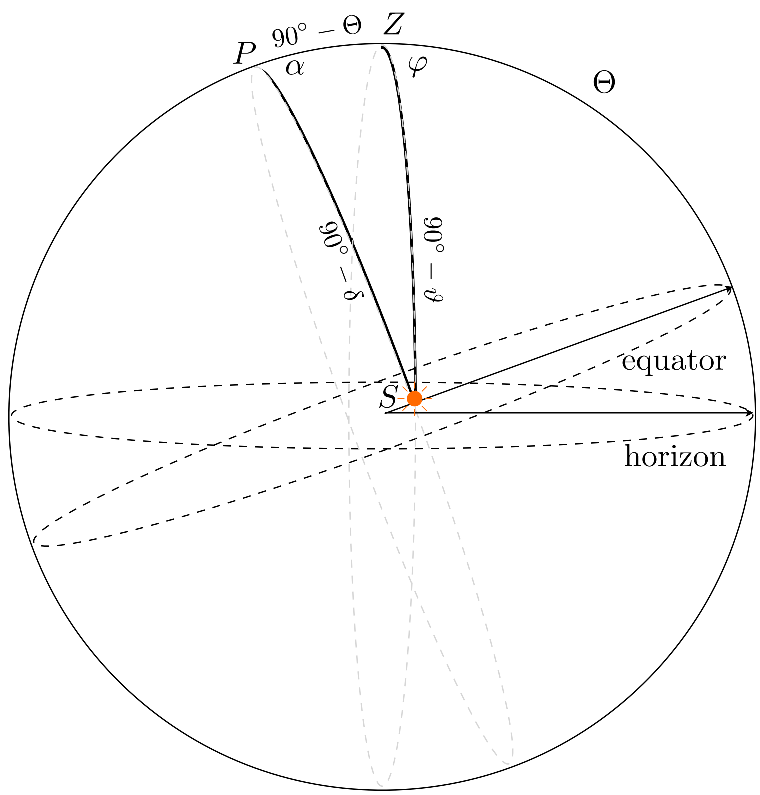


Figure 2.9: Geometry for the transformation from horizon coordinates to array coordinates

from bore-sight, and panels tilted at 16° from vertical. At 54°, $\langle \cos\eta \rangle = 0.780$ (dashed grey lines).

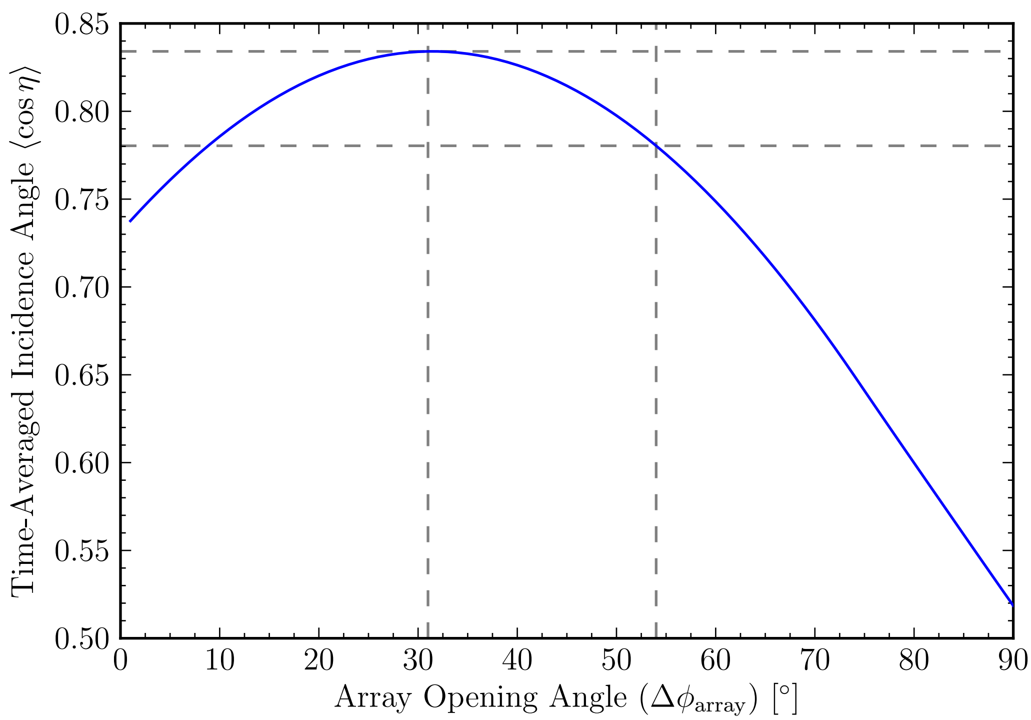


Figure 2.10: The time-averaged incidence angle $\langle \cos\eta \rangle$ vs. the array opening angle $\Delta\phi_{\mathrm{array}}$ for a particular scan simulation

### 2.5.2 The Flight Batteries

The flight batteries were ODYSSEY® Extreme Series™ PC1200 lead-acid batteries (Figure 2.11). These are 12 V, 40 A·h battery units, each with a mass of 17.4 kg. In each power system, two of these units were connected in series to make a 24 V battery pack. Therefore, a total of four units were required.

Other battery technologies such as lithium-ion or nickel–metal hydride (NiMH) offer a higher charge per unit mass. However, savings in weight come at the expense of greater complexity in charging and maintenance. For example, lithium-ion batteries typically require built-in electronics to balance the charge amongst individual cells and


Figure 2.11: The ODYSSEY® Extreme™ PC1200 lead-acid battery

to prevent thermal runaway. In lead-acid batteries, thermal runaway can be avoided if a standard sequence of charging stages (Section 2.5.3.1) is followed. This allows the use of commercially-available solar battery chargers.

Standard automotive lead-acid batteries with liquid electrolytes would not be suitable for flight. Sealed, "non-spillable" lead-acid batteries were required. These typically come with two different electrolyte types: gel, and Absorbed Glass Mat (AGM). Gel lead-acid batteries have silica dust added to the electrolyte, causing it to form a thick putty. AGM batteries use woven fibreglass mats within which the electrolyte is soaked. AGM batteries are available for both deep-cycle and starting purposes, and are suitable for use at low temperatures. They find use in motorcycles and all-terrain vehicles (where they can be mounted in any orientation) and for marine and aerospace applications.

The PC1200 batteries are AGM batteries that were recommended by CSBF, originally for BLASTPol. A replacement was needed for the NiMH batteries from Cobasys, LLC that were used in BLAST, which are no longer manufactured. At the time the BLASTPol power system was designed, the PC1200 batteries had already been used by a CSBF payload engineer in the CREAM cosmic ray experiment, establishing their suitabilty for use under vacuum. By the time Spider integration was underway, these batteries had seen use in both the 2010 and 2012 BLASTPol flights, and had proven to be robust.

#### 2.5.2.1 Peukert's Law for Battery Capacity

The nominal PC1200 capacity is specified to be 40 A·h. However for rechargeable batteries, the actual charge capacity depends on the discharge current, following a well-known empirical relation: Peukert's Law (Eq. 2.7) [67].

$$t(I) = t_0 \left( \frac{Q_0}{I t_0} \right)^k \tag{2.7}$$

The equation gives the discharge time $t$ (in hours) as a function of discharge current $I$ (in A). Battery manufacturers typically specify a nominal capacity $Q_0$ (in A·h) that corresponds to a specific discharge time $t_0$, and hence a specific discharge current $I_0 = Q_0/t_0$. If the Peukert constant $k = 1$, then the equation reduces to $t = Q_0/I$, with the discharge time simply varying inversely with discharge current, as expected for a constant battery capacity. However, if $k > 1$, then $t \propto I^{-k}$, corresponding to decreasing battery capacity with increasing discharge current. This result can perhaps be seen more clearly by reformulating the relation

$$Q(I) = It = I t_0 \left( \frac{Q_0}{I t_0} \right)^k = Q_0 \left( \frac{I t_0}{Q_0} \right) \left( \frac{Q_0}{I t_0} \right)^k = Q_0 \left( \frac{Q_0}{I t_0} \right)^{k-1} \tag{2.8}$$

From Eq 2.8, if $k = 1$, then $Q = Q_0$, but if $k > 1$, then $Q < Q_0$ when $I > I_0$. The ODYSSEY® datasheet gives a Peukert constant of $k = 1.106$ for the PC1200 model. This model's nominal charge of $Q_0 = 40$ A·h is specified for a discharge time $t_0 = 10$ h. At a discharge current of $I = 20$ A, which is the typical load on each power system, Eq. 2.7 yields $t = 1.686$ h. Therefore, the actual battery charge that can be provided at this discharge rate is only $Q = It = (20\text{ A})(1.686\text{ h}) = 33.7$ A·h. At full load, with no input from the solar arrays, the batteries last less than two hours before fully discharging. This is primarily a concern during ascent, when there is no attitude control and the solar arrays cannot be pointed towards the sun. However, in practice there is still enough intermittent solar irradiance plus ground albedo to prevent the batteries from discharging fully (Figure 2.16).

### 2.5.3 The Charge Controller

The TriStar MPPT-60 solar charge controller from Morningstar Corp. was used to regulate the charging of the batteries by the solar arrays. This model was recommended by CSBF for BLASTPol and had been flown on previous balloon payloads. The key feature of the controller was Maximum Power Point Tracking (MPPT). The controller could continuously detect the maximum power voltage of the solar arrays and receive input current at that voltage. It then acted as a switched-mode DC supply, stepping this voltage down to the appropriate battery-charging voltage (with $\gtrsim 97\%$ efficiency) before delivering output current. In this manner, all of the available array power at any given time could be used.

The MPPT-60 controller can supply a maximum of 60 A of continuous current to the batteries. It can be set to operate at a nominal system voltage of 12, 24, 36, or 48 VDC. The maximum input solar array voltage that can be supplied to it is 150 VDC. When set to operate with a system voltage of 24 VDC, the controller can handle a maximum solar array input power of 1600 W.

#### 2.5.3.1 Charging States

The charge controller uses the sequence of three charging states that are recommended for lead-acid batteries, including the flight batteries. These are Bulk, Absorption, and Float (Figure 2.12).

The Bulk charging state occurs when the batteries are not at 100% state of charge. The charge controller delivers the maximum available solar array power during this charging state. Once the battery terminal voltage reaches the Absorption voltage setpoint, the controller begins using constant voltage regulation to maintain it there. Charging current decreases with time as the batteries approach 100% state of charge. The batteries must remain at the Absorption voltage setpoint for a cumulative 150 to 180 minutes (depending on battery type) before the charge controller switches to the Float state. In

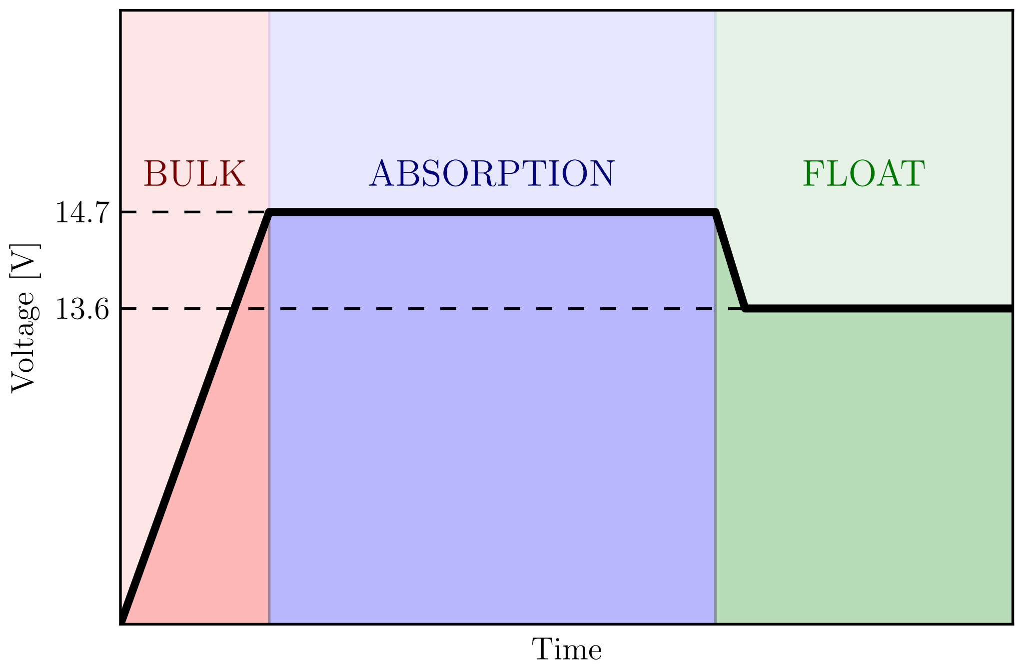


Figure 2.12: Voltage vs. time during the three charge controller charging states

the Float state, the batteries are fully-charged and no further chemical reactions take place. The Float voltage setpoint is a lower setpoint that maintains battery charge while reducing heating and gassing. If the battery voltage drops below the Float setpoint for more than 60 minutes, the charge controller re-enters the Bulk state. Two additional states (not shown) are Night, and Equalize. The controller enters the Night state when no input solar array power is available. The Equalize state is an extra boost to an even higher voltage setpoint for 60 to 120 minutes during Absorption. Equalization is carried out once every two to four weeks. Its purpose is to stir liquid electrolytes and level cell voltages. Equalization is not useful for sealed lead-acid batteries, and can be detrimental, causing excessive heating and gassing. Therefore, this state was never used with the flight batteries.

The charging voltage levels for the Absorption and Float states shown in Figure 2.12 are the maximum levels recommended for a single ODYSSEY® PC1200 12 V battery unit. The actual voltage setpoints used by the charge controller in each state were determined by setting the battery type using DIP switches on the control board. Various options for sealed (gel or AGM) or flooded (liquid) lead-acid batteries could be selected. The battery type selected for flight had Absorption and Float setpoints of 29.2 V ($2 \times$ 14.6 V) and 27 V ($2 \times 13.5$ V) respectively. Both setpoints were within the acceptable range for the PC1200 batteries.

#### 2.5.3.2 Hardware Modifications

The TriStar controller body has conduit knockouts through which wires can be routed, where they can connect to screw terminals on the control board. For both the array input and battery output, $+$ and $-$ screw terminals are provided that can acccomodate a single wire between 2 AWG and 14 AWG. This interface presented a practical difficulty, since the array input circuit consisted of four 12 AWG conductors, one for each array column. The battery charging circuit also used four 12 AWG conductors. A reliable way was needed to connect these to the controller.

The problem was solved by milling the ends of four small copper rods to fit securely inside each of the screw terminals. One or more through holes were drilled along the length of each rod, allowing the four 12 AWG wires to be attached to the rods by crimping ring terminals to the ends of the wires and bolting them to the rods (Figure 2.13). The conduit knockouts were not used. The two sets of eight wires instead connected to the contacts of two 8W8 Combination D-subminiature (D-sub) connectors, one for battery power and one for array power. These were mounted to the controller's cover plate. Holes were milled into the plate for these 8W8 connectors and for two 9-pin D-sub connectors, one for the RS-232 serial communication link, and the other for the battery sense connection. The battery sense connection (which is omitted from Figure 2.6) consisted of 24 AWG wires connected directly from sense terminals on the controller to the battery terminals, allowing the controller to measure the true battery voltage without

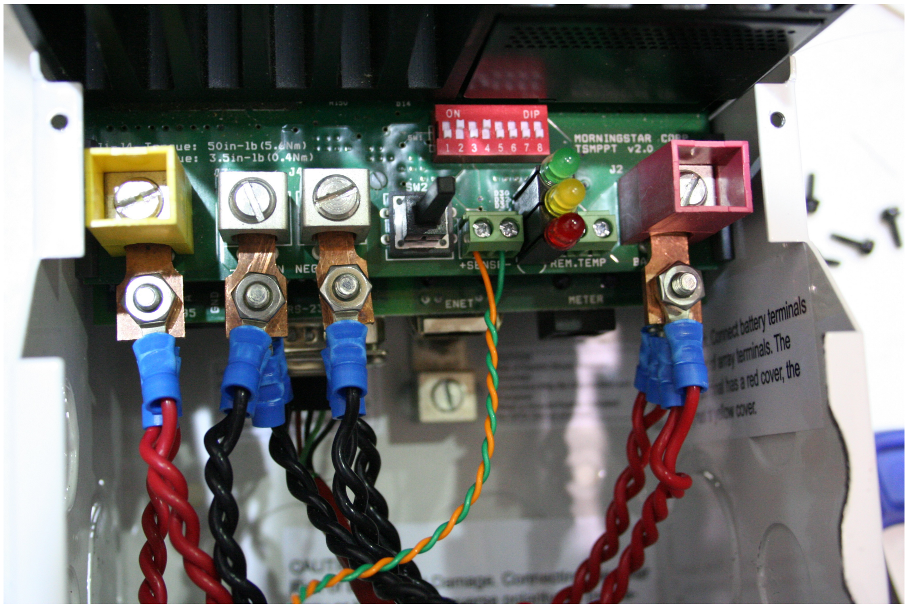


Figure 2.13: Hardware modifications to the TriStar charge controller. This photo shows the inside of the controller, with the copper rods used as bus bars for power wiring. Also visible are the battery sense terminals (orange and green wires) and the DIP switches.

the voltage drops across the power lines. The improvised front plate with connectors provided a reliable interface to the controller for flight, making it unncessary to open the cover plate or to loosen the screw terminals once they had been tightened down. A final modification was that the charge controller heat sinks were painted white to increase their emissivity, and a painted aluminium radiator plate was attached to each of them, to further increase radiative cooling.

#### 2.5.3.3 Software and Communications Protocol

A dedicated pcm thread was written to read data from each charge controller over RS-232. The thread periodically queried a controller for the values of the array voltage, array current, battery voltage, battery current, target battery charging voltage, charging state, heat sink temperature, and alarm and fault bitfields. Communication with the charge controllers used the Modbus serial communications protocol. The serial communication used a baud rate of 9600 bits/s, no parity, 8 data bits, 2 stop bits, and no flow control. It is important to note that the charge controllers required a straight through 5-wire RS-232 cable rather than the more standard 3-wire connection. In addition to TX/RX, RX/TX, and GND (pins 2, 3, and 5 on the standard 9-pin D-sub connector), the Data Terminal Ready (DTR, pin 4) and Request To Send (RTS, pin 7) signal lines must be connected in order to supply power to the TriStar's opto-isolated serial port.

#### 2.5.3.4 Charge Controller Relay Box

The charge controllers are powered from their battery output side. For the 2010 flight of BLASTPol, since the TriStar controller had never been flown before, a relay box was developed to interrupt the connection between the controller and the flight batteries (Figure 2.6). Two modified versions of this relay box design were built for Spider. These relay boxes added the ability to power cycle the charge controllers in flight, in case they entered a bad state or a serial communications error occurred. Ideally, this feature would never be used and the relays would remain closed throughout the flight. Indeed, power cycling was never necessary in any of the three flights in which the TriStar MPPT controller was used.

A KG Technologies Inc. K105 series mechanical latching relay rated for 100 A and 24 V was installed in each relay box. The relay control board included two solid-state relays to drive the SET and RESET coils of the latching relay. This design is essentially the same as the high-current power switching circuit that was used for the pointing motors, which is described in detail in Appendix A.1.3. The only difference was that the motor power switching circuits used the K100 relay, rather than the K105 model. The K105 relay includes a small auxiliary switch that is actuated when the main relay changes states. In the Spider relay boxes, this switch was wired to a blue LED that was mounted to the side of the box to indicate the relay state: closed when illuminated.

### 2.5.4 The Power Breakout Board (PBOB)

The power breakout board (PBOB) was a custom printed circuit board (PCB) whose purpose was to collect all of the power switching and current sensing electronics for the experiment in one place. The goal was to improve upon the somewhat haphazard and decentralized power breakout schemes of past balloon experiments. The design was intended to be somewhat modular, with enough switching circuits of various types to control power to the typical suite of subsystems on an experiment such as BLASTPol or Spider. The PBOB had to meet a number of design challenges. First, it had to consist entirely of analogue or low-density digital electronics that would not be susceptible to a cosmic ray upset. Second, different switching circuits needed to be controlled by different external interfaces, either TTL voltage levels, or voltage pulses. Finally, different switching circuits needed to provide switched output power at different voltages and with different current carrying capacities, depending on the subsystem whose power was being switched. Schematics and detailed descriptions of the three types of switching circuits that were designed to meet these requirements are in Appendix A.1.

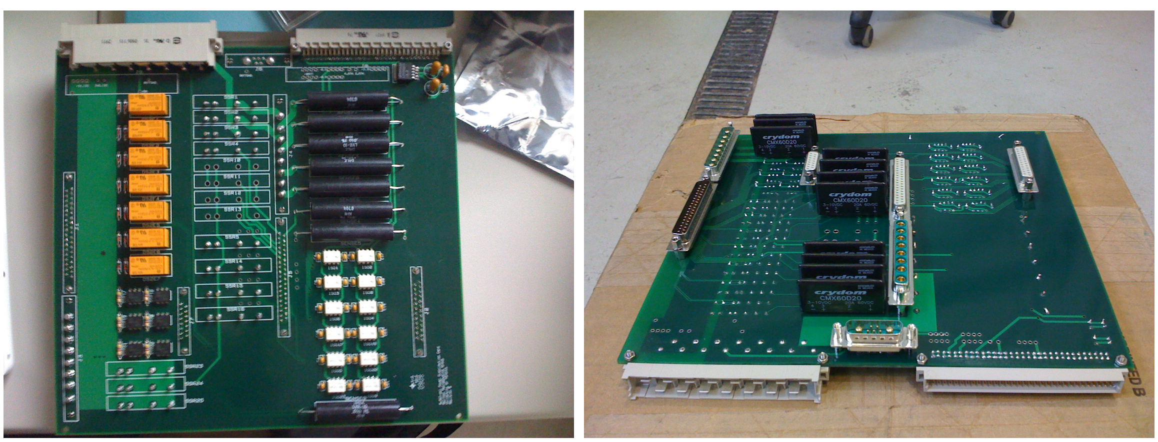

Figure 2.14: Photographs of an assembled PBOB

Two instances of the PBOB were installed in Spider, one for each of the two power systems. Figure 2.14 shows two photographs of a finished board. The board was designed to plug into the backplane of an existing BLASTbus crate. The PCB was designed with four copper layers (two surface and two internal). To carry more current, the copper layers had a weight of 2 oz/ft$^2$, which is twice the default value used by board manufacturers. The open-source KiCad software was used for the PCB layout.

#### 2.5.4.1 MCC Power Breakout Box (MCC PBOB)

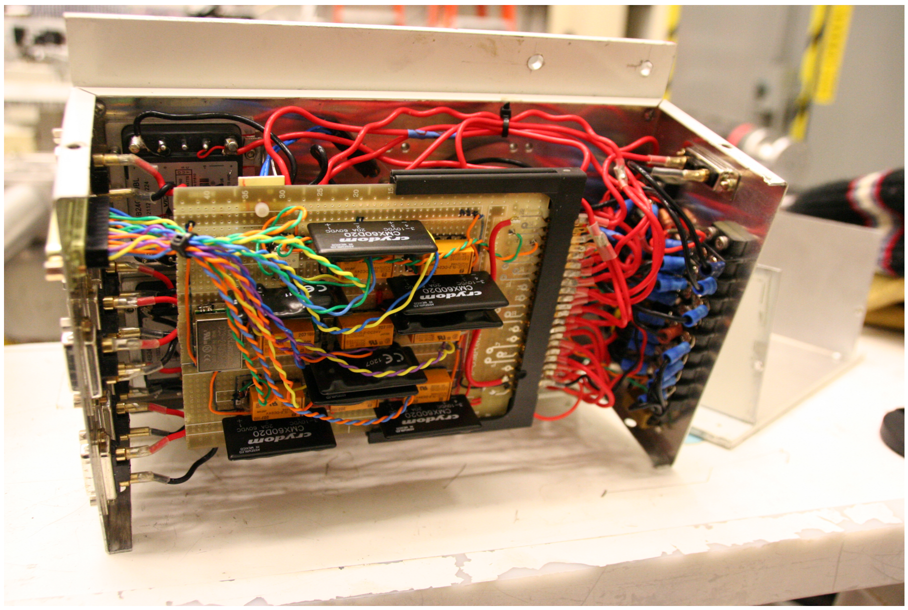

Figure 2.15: The MCC Power Breakout Box

The two PBOBs had enough switching circuits for all SPIDER subsystems except for the six MCCs and Sync Box. Therefore, a separate MCC power switching box was constructed as a late addition to the power system. The box received input power directly from the flight batteries. This was supplied to three 100 W Vicor DC-to-DC converters (DC-DCs), each with a 24 V output. These were used to isolate battery ground from MCC power/case ground. Eight PBOB switching circuits of Type 1 (Appendix A.1.1) were assembled on a protoboard. One of them was a master switch, switching the input battery power to the box. The other seven switched the Vicor 24 V outputs, allowing the power to each of the MCCs and the Sync Box to be individually switched. Two computers were supplied by each Vicor. The assembled MCC Power Breakout Box is shown in Figure 2.15.

### 2.5.5 In-Flight Performance

#### 2.5.5.1 Power Budget

Table 2.2: The power budget for the SPIDER 2015 flight

| Subsystem | Current [A] | Power [W] |
|---|---|---|
| *Outer Frame* | | |
| ACS | 2.40 | 64.8 |
| FLCs | 2.19 | 59.1 |
| MCCs | 4.88 | 131.7 |
| RW | 1.50 | 40.4 |
| Pivot | 1.84 | 49.7 |
| El Drive | 1.30 | 35.0 |
| Star Cameras | 1.70 | 45.8 |
| **Subtotal** | **15.80** | **426.5** |
| *Inner Frame* | | |
| MCEs | 16.17 | 472.1 |
| HK & Misc. | 1.99 | 58.1 |
| **Subtotal** | **18.16** | **530.2** |
| **TOTAL** | **33.96** | **956.8** |

Table 2.2 lists the measured power consumption of major SPIDER subsystems in flight. The current values are measured by the PBOB sense circuits and read in over the BLASTbus. The data in the table are averaged over a time period of almost one day (~72 000 s, beginning Jan. 4, 2015 at 14:37:36 UTC) during which SPIDER was continuously in `spider_scan` mode (Section 2.9). Motor current measurements therefore indicate average consumption while scanning. Power is estimated by multiplying these

current values by the mean battery voltage over this time interval. The inner frame charge controller remained at the Absorption setpoint of 29.2 V throughout the flight, while the outer frame charge controller was at the Float setpoint of 27 V.

Outer frame subsystems not in the tally include the differential GPS (dGPS), which was not operational, the line-of-sight (LOS) transmitters (powered off), and the heaters (powered off). Low power pointing subsystems such as the magnetometer and inclinometers are included in the ACS measurement. The RS-232 serial hub current is included in the FLC measurement. The mean battery charging current reported by the outer frame charge controller was 17.46 A during this time, which corresponds to 471.4 W supplied to the outer frame batteries. The inner frame charge controller reported supplying 17.08 A (498.8 W) to the inner frame batteries, which is less than the measured load. This discrepancy could be due to a gain calibration on the isolation amplifiers used in the inner frame current sense circuits, which could be uncertain by more than 10%.

Since the motors (Section 2.6) sometimes supply battery current rather than consuming it, the average reaction wheel (RW) and pivot current does not represent the sustained power consumption during significant portions (several seconds) of a scan. During scan turn-arounds, battery current to the motor drives consistently reached peak values of 4 A for the pivot and 6 A or higher for the RW, corresponding to 108 W and 162 W respectively.

#### 2.5.5.2 Power During Ascent

Figure 2.16 shows outer frame charge controller data and pointing data from ascent. During the time interval shown, Spider's altitude increased from 17.2 km to 18.7 km. The payload was drifting freely in azimuth during this time, since motor control was not established until ~24.5 km altitude. As a result, the available solar array power varied continuously between maximal and minimal. The charge controller switched charging states accordingly. During the grey shaded intervals in the figure, the charge controller

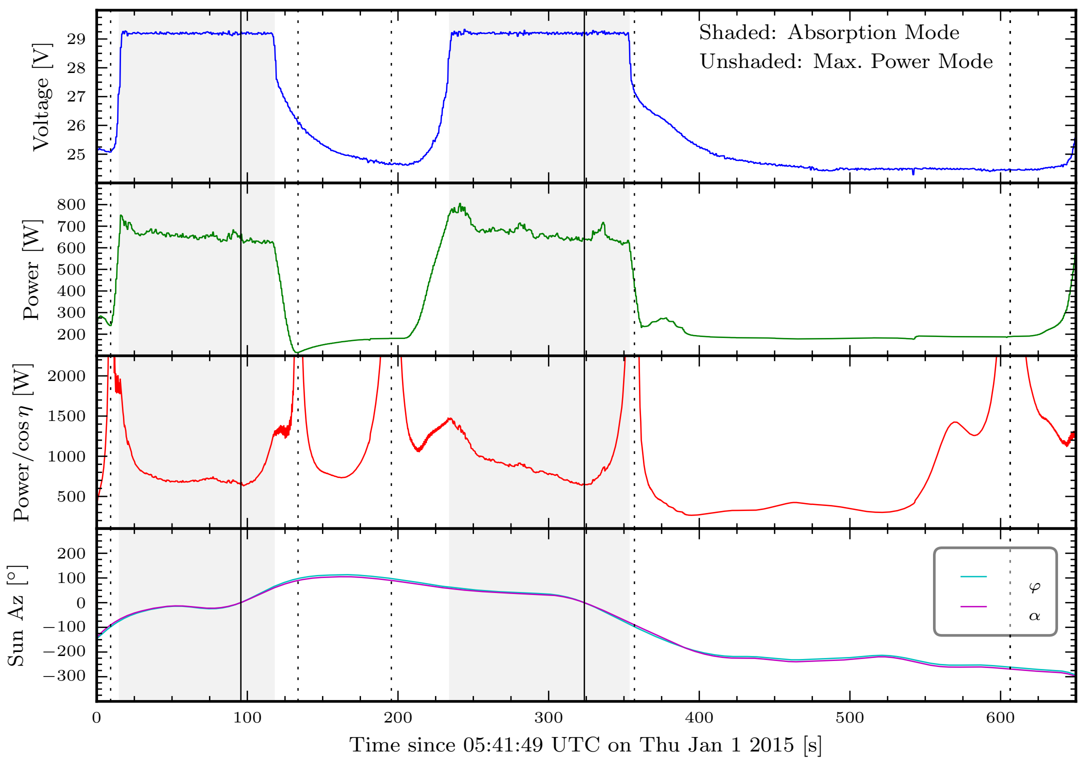


Figure 2.16: Power delivered to the outer frame batteries by the charge controller during ascent. *Top to Bottom* – Battery voltage (blue), power delivered to batteries: the product of the battery voltage and charging current reported by the controller (green), power corrected for the cosine of the incidence angle $\eta$ (red), and sun azimuth relative to the array normal ($\varphi$ and $\alpha$ from Section 2.5.1.2, cyan and magenta).

was in Absorption mode. During the unshaded intervals, the controller was in Bulk (MPPT) mode.

The sun azimuth relative to the horizontal projection of the array normal, $\varphi$, is plotted in the lowermost panel. As described in Section 2.5.1.2, this angle and the sun elevation[4] $\vartheta$, are used to compute the sun azimuth $\alpha$ and elevation $\delta$ relative to the array normal, over this time interval. These angles are then used to compute $\cos\eta = \cos\alpha\cos\delta$, where $\eta$ is the angle of incidence. The solid vertical lines are times of nearly normal incidence

[4]The sun elevation was approximately 25° during the period shown.

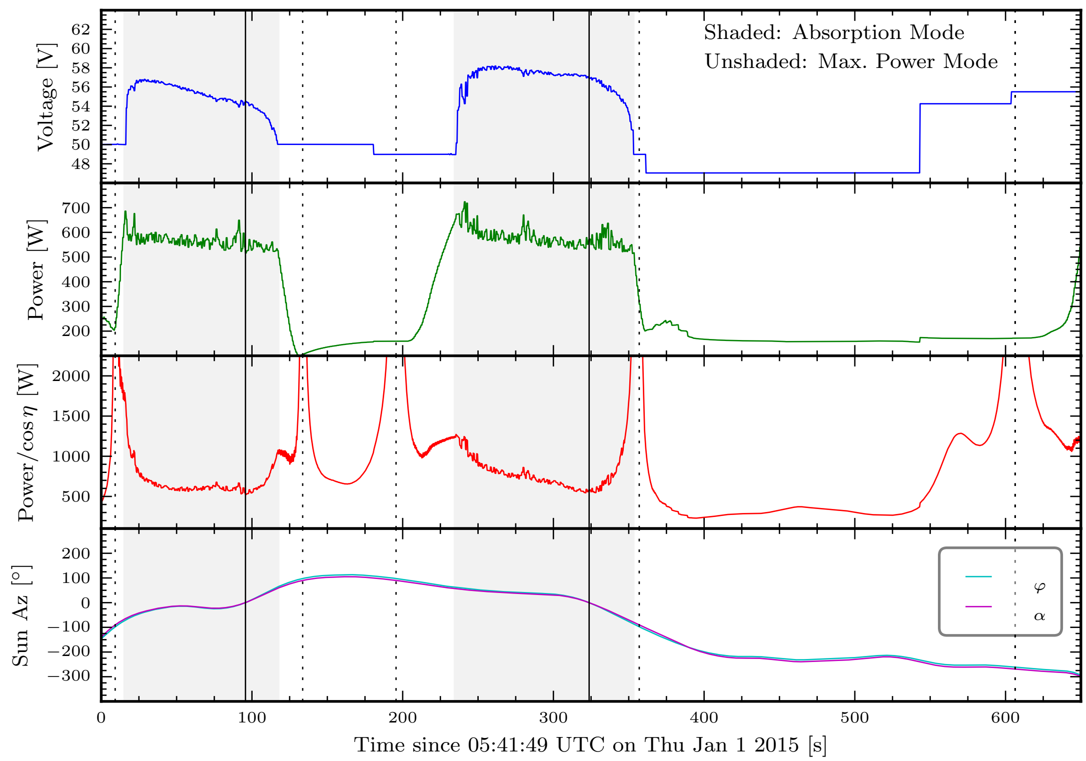


Figure 2.17: Power received from the outer frame array by the charge controller during ascent. *Top to Bottom* – Array voltage (blue), power delivered to controller: the product of the array voltage and input current reported by the controller (green), power corrected for the cosine of the incidence angle $\eta$ (red), and sun azimuth relative to the array normal ($\varphi$ and $\alpha$ from Section 2.5.1.2, cyan and magenta).

(times when $\alpha = 0°$) and the dotted vertical lines are times of parallel incidence ($|\alpha| = 90°$ or $270°$). In the second lowest panel (red curve), the power delivered to the batteries is divided by $\cos\eta$, to correct for the angle of incidence. In the unshaded regions, this should provide an estimate of the peak power of the arrays. In unshaded areas that are between two dotted lines, the sun is shining on the (translucent) back side of the arrays ($\cos\alpha < 0$). In these intervals, $|\cos\alpha|$ was used in the correction factor, so as to provide a positive power estimate. The corrected peak power does appear to be lower for reverse incidence than forward, as might be expected. Overall, however, the correction does not work very well, producing singularities near $\alpha = 90°$. Presumably this is due to a contribution to the power from ground albedo, which does not vary as $\cos\eta$. This contribution is at a level of ∼200 W.

Figure 2.17 is the corresponding set of plots to Figure 2.16, except showing the outer frame power calculated using the input array voltage and current measured by the charge controller, rather than the output battery voltage and current. The inner frame power received during this time interval was similar to the outer frame power, and is not shown.

#### 2.5.5.3 End-of-Flight Discharge Test

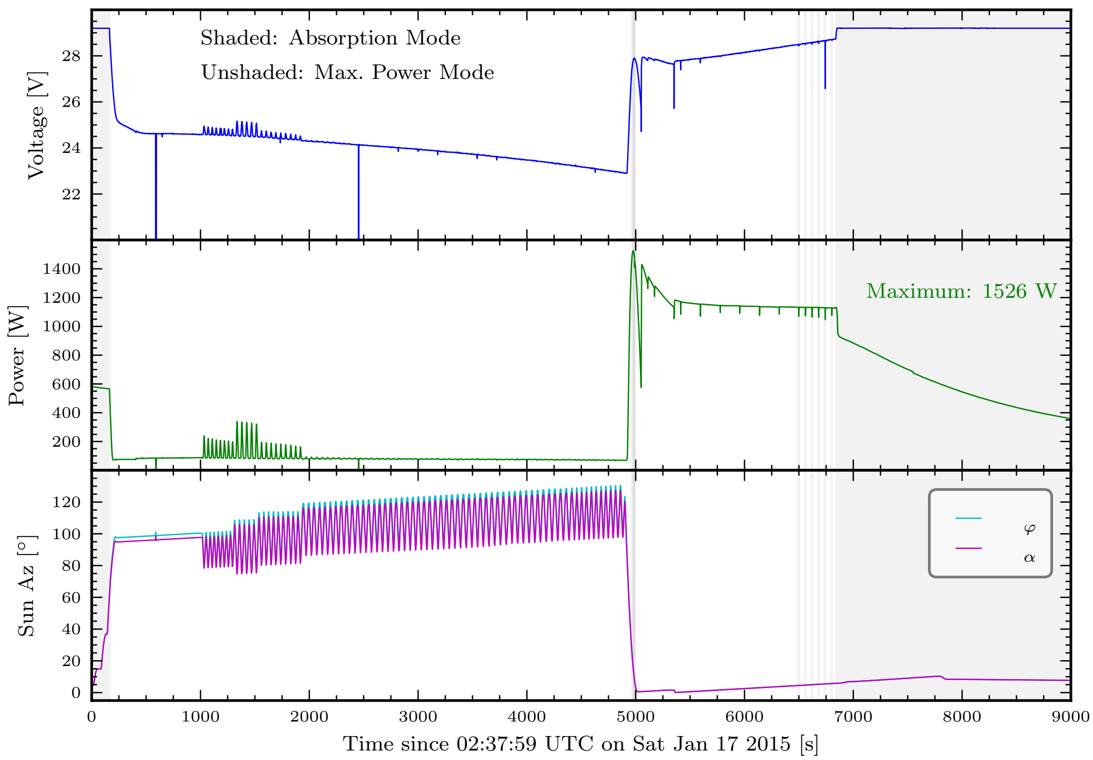


Figure 2.18: Power delivered to the inner frame batteries during the discharge test. *Top to Bottom* – Battery voltage (blue), power delivered to batteries from the charge controller (green), and sun azimuth angles $\varphi$ and $\alpha$ relative to the array normal (cyan and magenta).

Near the end of flight, the solar arrays were pointed edge-on to the sun, minimizing the available power. Inner frame charge controller data and pointing data acquired during this test are shown in Figure 2.18. During the grey shaded intervals, the charge controller was in Absorption mode. During unshaded intervals, it was in Bulk (MPPT) mode. As shown in the figure, SPIDER rotates the panels away from the sun and then carries out a sinusoidal Az scan. Initially, power is produced at the low Az end of the scan, until the scan centre is adjusted upward. Over the course of 80 minutes, the batteries discharge to 22.9 V. SPIDER is then rotated to a position of nearly normal incidence. The charge controller intermittently enters Absorption mode once array power is restored, even though the battery voltage is still below the Absorption setpoint of 29.2 V. This is normal behaviour for the TriStar model.

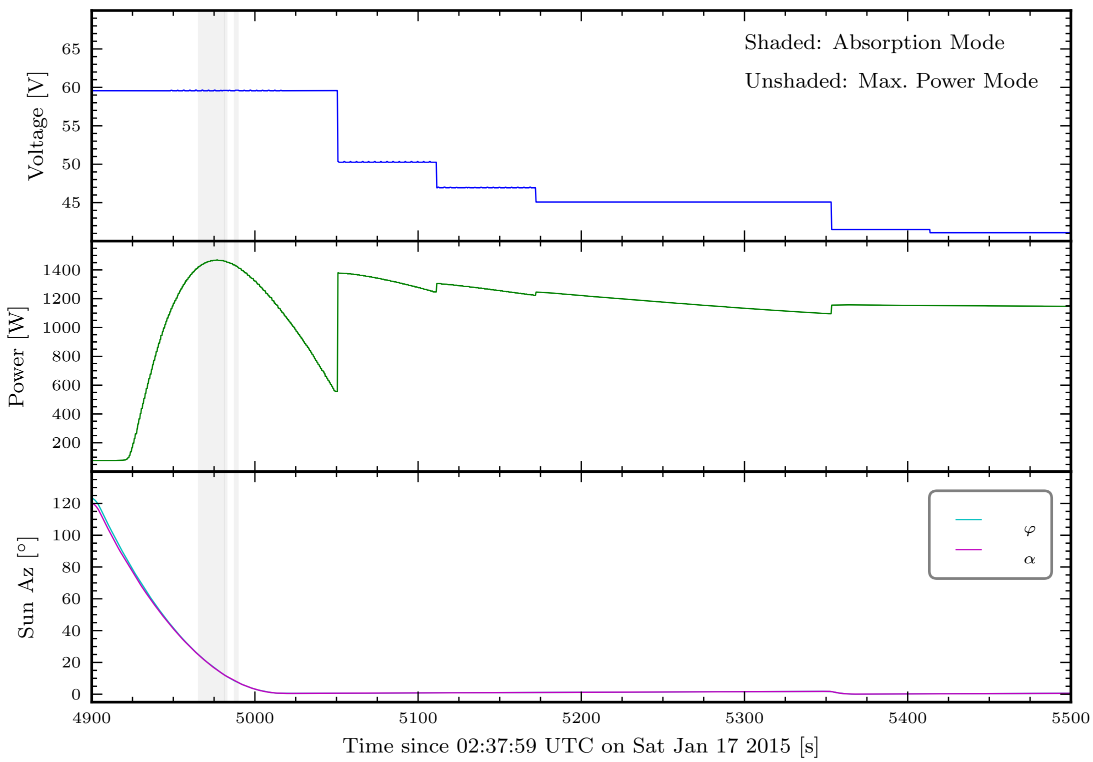


Figure 2.19: Power received from the inner frame solar array during the discharge test. *Top to Bottom* – Array voltage (blue), power received by the charge controller from the array (green), and sun azimuth angles $\varphi$ and $\alpha$ relative to the array normal (cyan and magenta).

Another curious feature is that the output power to the batteries rises to a maximum as the arrays approach normal incidence, but then appears to drop sharply, before rising again. This latter behaviour is explained by the time delay required for the charge controller to find the peak power point of the array. As shown in Figure 2.19, the angle of incidence is changing relatively rapidly, and presumably the peak power point with it. Yet, the array voltage remains near 60 V throughout most of the slew. The system goes beyond the peak power point of the arrays as the payload continues to rotate, and the power drops. Then the array voltage undergoes a step change to the peak power point at the new angle of incidence, causing the input power to rise sharply again. This process is repeated several times until azimuth is changing only slowly and the input power becomes steady. Increasing solar cell temperature has the dominant effect here, rather than the increasing irradiance. As the cell temperatures increase, the peak power voltage and overall peak power both decrease, explaining why the array power jumps up to a slightly lower value after each step change. In conclusion, at the steady state temperature, under full illumination, the solar arrays each provide 1100 W of power at float altitude, providing a factor of $\sim$2 margin over the SPIDER power budget (Table 2.2).

## 2.6 The Azimuth Drive

The SPIDER Az drive has two primary components: the reaction wheel and the motorized pivot. The reaction wheel is supported by four carbon fibre struts extending upward and inward from the bottom corners of the gondola frame (Figures 2.5 and 2.20). It sits just below the main aluminum honeycomb floor of the gondola, where the electronics are mounted. The reaction wheel works by conservation of angular momentum: a torque applied to the reaction wheel by its motor results in a torque of opposite sign on the gondola. The pivot is a motorized joint located at the top of the payload, where the three suspension cables meet (Figure 2.5). Extending upwards from the pivot rotor is the flight train: the steel cabling that connects to the balloon. The pivot stator is rigidly connected to the suspension cables and the rest of the payload underneath. The pivot is able to provide additional torque in Az by twisting the flight train, which acts as a torsional spring. This torque aids in scanning the gondola, and prevents the reaction wheel from reaching its saturation speed: the speed at which the motor's back-EMF[5] prevents further current through the motor windings. Without the pivot, saturation inevitably occurs because the reaction wheel must absorb the angular momentum produced by external torques caused by wind shear and other disturbances. On long timescales, the pivot is able to dump this excess angular momentum to the balloon, through the flight train.

## 2.6.1 The Reaction Wheel

### 2.6.1.1 Mechanical Design

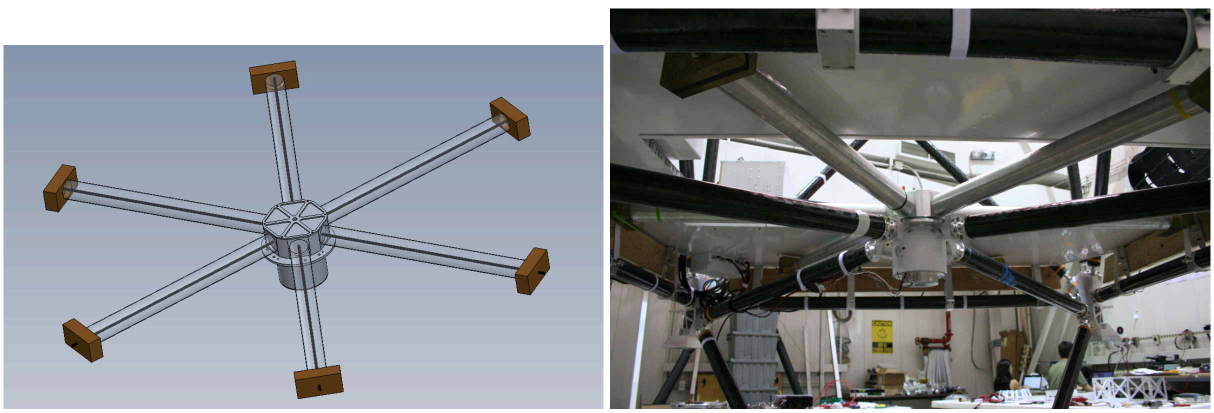

Figure 2.20: *Left* – A SolidWorks rendering of the Spider reaction wheel, showing the bricks, spokes, hub, and motor. The spokes have been made transparent in order to reveal the threaded rods. *Right* – a photograph of the reaction wheel as it is mounted on the underside of the gondola.

The reaction wheel consists of six 7 kg brass bricks connected to a central hub by 1 m long spokes (Figure 2.20). These spokes are 6061-T6 aluminum pipes with a 4" outer diameter (OD) and a 0.125" wall thickness. Structural support is provided by 3/8 - 16

[5]For these motors, the back-EMF $V_b = K_b\omega$, where $K_b$ is the voltage constant (Table 2.3)

threaded rods running through the bricks and spokes, and terminating inside the central hub. Mass is concentrated around the outside of the wheel in order to maximize its rotational inertia at a given weight. The reaction wheel's moment of inertia around its spin axis is 44.77 kg·m$^2$, and its mass is 50.34 kg, including the motor, which extends below the wheel.

#### 2.6.1.2 The Motor Torque Control Loop

The reaction wheel is driven by a K178200-6Y1 brushless DC motor from Parker Bayside Motion (Table 2.3). This frameless motor was installed in a custom-designed motor housing that included a resolver for feedback sensing. The resolver is an analogue rotary position sensor that determines the angular position and speed of the rotor shaft relative to the stator windings. This sensor is necessary for correct commutation of the motor. Motor current is driven by a DPRALTR-060B080 digital servo amplifier from Advanced Motion Controls (AMC). This servo drive carries out digital commutation based on the resolver feedback. It can drive current through the motor windings up to a maximum of 60 A. The servo drive regulates the output current, executing the lowest-level control loop in the reaction wheel control system.

A voltage is supplied to one of the servo drive's analogue inputs by a ±5 V Digital-to-Analogue Converter (DAC). The drive interprets this as an output current request using a proportionality constant of 3.6 A/V. This determines the motor torque[6] (Figure 2.22). The input DAC level is determined by a Proportional, Integral (PI) negative feedback control loop (Figure 2.21).

As shown in Figure 2.21, the feedback sensors for this control loop are SPIDER's yaw axis gyroscopes. SPIDER has six KVH DSP-3000 digital fibre optic rate gyroscopes: two for each of the yaw, pitch, and roll rotation axes. The gyros measure the payload angular velocity in each axis, relative to an inertial frame. Ignoring pendulations, rotation around

[6]For these motors, $\tau = K_t I$, where $K_t$ is the torque constant (Table 2.3)

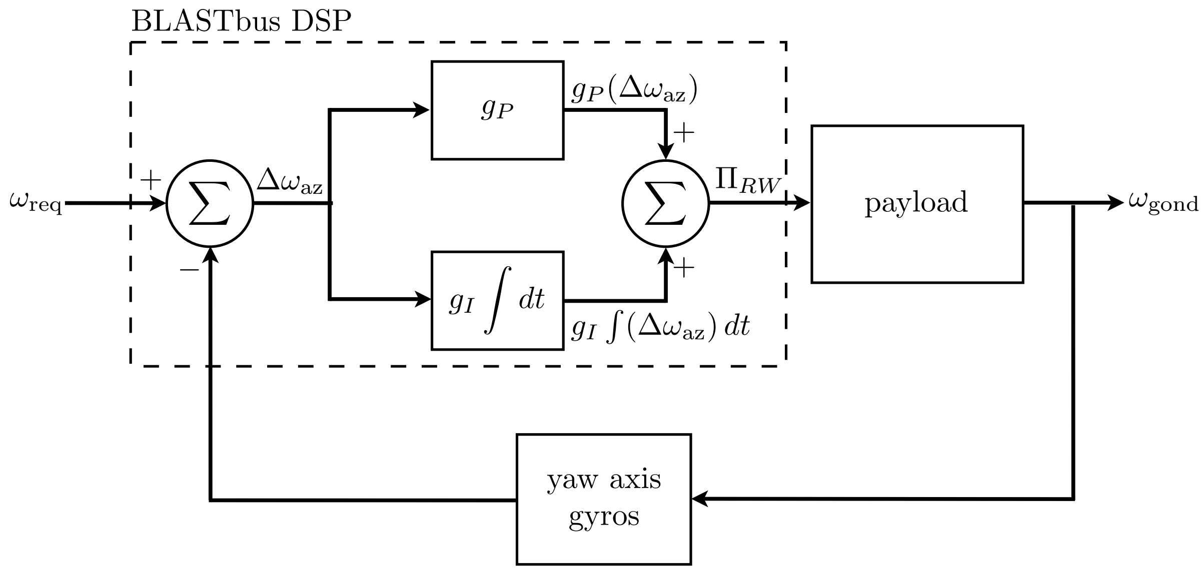


Figure 2.21: A block diagram of the reaction wheel (RW) motor torque control loop

the yaw axis corresponds to motion of the payload in azimuth. In the control system, the yaw axis gyros are sampled at approximately 1 kHz by a digital signal processor (DSP) on one of the BLASTbus motherboards in the ACS. The gyro signal is then subject to further digital filtering. A real-time process on the DSP subtracts this measured angular velocity, $\omega_{gond}$, from the requested azimuthal angular velocity $\omega_{req}$, to produce the Az velocity error:

$$\Delta\omega_{az} \equiv \omega_{req} - \omega_{gond}. \tag{2.9}$$

The DSP then computes the two control terms. The $P$ term is proportional to $\Delta\omega_{az}$ , so that the reaction wheel motor torque increases in proportion to the velocity error that it is trying to correct. The $I$ term is proportional to the integral of $\Delta\omega_{az}$, which helps reduce steady-state error:

$$\Pi_{RW} = g_P(\Delta\omega_{az}) + g_I \int (\Delta\omega_{az})\,dt. \tag{2.10}$$

The result, $\Pi_{RW}$ is the 16-bit DAC digital input level. Therefore, based on the velocity error, the DSP control loop ultimately determines the reaction wheel motor torque $\tau_{RW}$ (Figure 2.22). The values of the reaction wheel gains $g_P$ and $g_I$ could be commanded in flight, and were tuned starting shortly before the payload reached float altitude ($\gtrsim$ 120 000 ft), during the period when LOS commanding was possible.

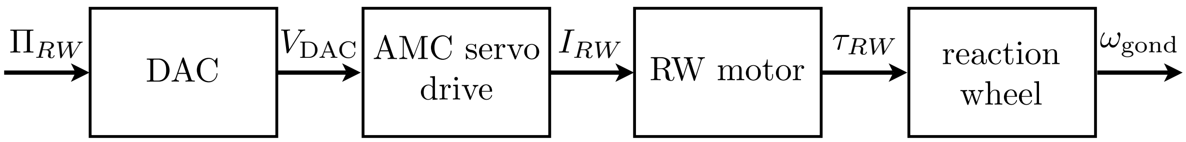


Figure 2.22: Expanded diagram of the block labelled "payload" in Figure 2.21, illustrating how the DSP control loop ultimately controls reaction wheel torque by setting the DAC level $\Pi_{RW}$ based on the Az velocity error

With this control loop in operation on the DSP, the payload is a velocity-commandable system, from the point of view of the flight computers. On the flight computers, pcm computes $\omega_{\mathrm{req}}$ at approximately 120 Hz, and its value is communicated to the ACS over the BLASTbus. More complex pointing and scanning motions are achieved by programming pcm to vary $\omega_{\mathrm{req}}$ with time (see Section 2.8).

## 2.6.2 The Pivot

### 2.6.2.1 Mechanical Design

As shown in the left panel of Figure 2.23, the pivot design incorporates three bearings. At the top of the pivot is a steel section of the casing that has welded tabs to which the ropes that suspend the gondola are attached. The pivot rotor shaft (red) is supported within this section of the casing by an SKF 51218 thrust ball bearing. This bearing supports the entire weight of the payload. Just above this bearing is an SKF NK 90/25 needle roller bearing. Another identical needle roller bearing supports the rotor farther down, in the cylindrical section of the casing where the motor windings are located. The primary purpose of the needle bearings is alignment. The resolver is visible as a small cylindrical protrusion at the very bottom of the casing, extending partially into the control box.

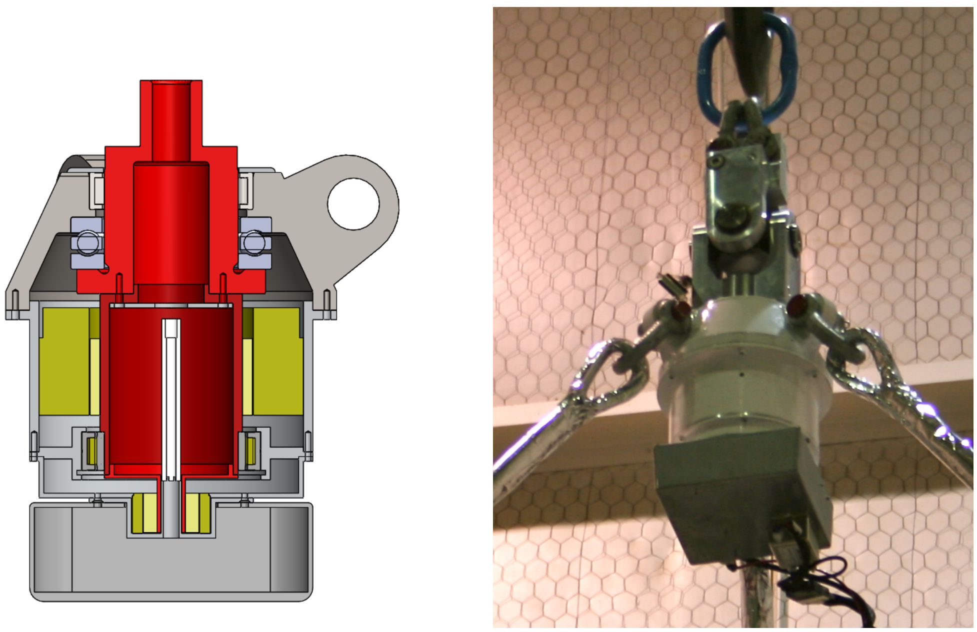

Figure 2.23: *Left* – A cross-sectional view of a SolidWorks model of the pivot motor. *Right* – A photograph of the pivot showing, from top to bottom, the universal joint that connects to the flight train, the motor casing (white), and the pivot control box. The three cables from which the gondola is suspended are also visible.

Table 2.3: Properties of the reaction wheel and pivot brushless DC motors from Parker Bayside Motion. The physical parameters are the same for both models: only the winding constants differ. To ensure proper operation, the correct winding constants must be stored in the firmware of the each motor's servo drive.

| | Reaction Wheel | Pivot |
|---|---|---|
| Model № | K178200-6Y1 | K178200-8Y1 |
| *Physical Parameters* | | |
| Maximum Mechanical speed [rpm] | 6000 | |
| Stall Torque Continuous [N·m] | 25.74 | |
| Maximum Winding Temperature [°C] | 155 | |
| Rotor Inertia [kg·m$^2$] | $1.8 \times 10^{-3}$ | |
| Number of Rotor Magnet Poles | 18 | |
| Mass [kg] | 6.34 | |
| *Winding Constants* | | |
| Stall Current Continuous [$A_{rms}$] | 12.9 | 8.15 |
| Peak Current [$A_{rms}$] | 40.9 | 25.76 |
| Voltage Constant $K_b$ [V/(rad·s$^{-1}$)] | 1.639 | 2.595 |
| Torque Constant $K_t$ [N·m/$A_{rms}$] | 2.007 | 3.178 |
| Resistance [Ω] | 0.6857 | 1.7 |
| Inductance [mH] | 6.118 | 15.3 |

#### 2.6.2.2 The Motor Velocity Control Loop

The pivot was driven by a K178200-8Y1 brushless DC motor from Parker Bayside Motion. This motor has the same physical dimensions as the reaction wheel motor (178 mm OD), but with a different set of windings (Table 2.3). The system driving the pivot motor was similar to the reaction wheel system described in Section 2.6.1.2, with two important differences. First, the pivot control loop was implemented on the flight computers as part of pcm, instead of on a DSP in the ACS. Therefore, it operated at a lower rate (∼120 Hz). Second, the pivot's servo drive[7] was programmed to be able to interpret the pivot DAC voltage as a *velocity request*, rather than a current request. In this mode of operation, velocity mode, the analogue input scaling was approximately 1.6 rpm/V.

In velocity mode, the drive attempts to servo the pivot rotor velocity to a requested value, based on resolver feedback. Operating the pivot in velocity mode rather than torque mode was a significant departure from the pivot control in previous balloon payloads. During early laboratory tests in torque mode, abrupt pivot motions were found to drive pitch and roll pendulation modes of the payload, especially during scans as wide and fast as SPIDER's (Section 2.8). Since static friction is considerably larger than rolling friction in the pivot, motor current would build up until the static friction was overcome, resulting in rapid motion. This problem can be avoided with velocity control, which ensures smooth and continuous motion of the pivot.

In BLASTPol, whose pivot operated in torque mode, the control terms included the following

$$\tau_{\text{piv}} = g_1 \Delta\omega_{RW} + g_3 \Delta\omega_{\text{az}}, \tag{2.11}$$

where $\Delta\omega_{RW} \equiv \omega_{RW} - \omega_{SP}$. The first pivot control term attempts to servo the reaction wheel rotation speed to a setpoint value, $\omega_{SP}$, by providing a torque proportional to the

[7]The pivot was originally going to use a different AMC servo drive model from the same series: the DPRALTR-020B080, which had a peak current rating of 20 A, vs. 60 A for the reaction wheel's servo drive. However, the pivot's original servo drive was replaced with the spare reaction wheel servo drive on the ice.

error between this setpoint and the measured reaction wheel speed $\omega_{RW}$[8]. The second control term helps the payload to scan by providing a torque proportional to the Az velocity error. In early SPIDER testing, the second term of Eq. 2.11 was found to cause random pivot motions due to amplification of the gyro noise present in $\Delta\omega_{\rm az}$. As a first attempt at a solution, this term was replaced with $g_3\alpha_{\rm req}$, where $\alpha_{\rm req}$ is the requested or "theoretical" payload angular acceleration, in azimuth. During a scan, $\alpha_{\rm req}$ varies deterministically with azimuth.

When this measure failed to prevent sudden pivot motions, velocity control was implemented. It was necessary to translate the terms of Eq. 2.11 into equivalent expressions for $\omega_{\rm piv}$, the pivot rotation rate. The pivot provides torque by twisting the flight train, which can be considered a torsional spring. Therefore, $\tau_{\rm piv} = k\theta_{\rm piv}$, where $k$ is the effective flight train torsional spring constant[9], and $\theta_{\rm piv}$ is the rotation angle of the pivot, measured relative to the angle of zero twist in the flight train. Differentiating both sides of the equation yields $\dot{\tau}_{\rm piv} = k\omega_{\rm piv}$. Substituting in the expression $g_3\alpha_{\rm req}$ for $\tau_{\rm piv}$, we obtain $\omega_{\rm piv} = (g_3/k)\dot{\alpha}_{\rm req}$ as an equivalent of the second control term in velocity mode. During a scan, SPIDER's Az varies sinusoidally with time. Therefore, so do all of its derivatives. In particular, for a sinusoidal scan profile, $\dot{\alpha}_{\rm req} \propto -\omega_{\rm req}$. The scan control term can therefore be implemented by setting $\omega_{\rm piv} \propto -\omega_{\rm req}$ with some proportional gain.

Two additional control terms were implemented in velocity mode to servo the reaction wheel to the setpoint speed. Note that in SPIDER, $\omega_{SP}$ is typically set to zero. The width and angular acceleration of the scan necessitate a large swing in reaction wheel velocity, centred on 0 deg/s. Therefore $\Delta\omega_{RW} = \omega_{RW}$, and the two control terms are given by

$$\omega_{\rm piv} = g_1\omega_{RW} + g_2\tau_{RW}. \tag{2.12}$$

[8]The measured reaction wheel speed comes from the resolver, and is read by the flight computers from the reaction wheel motor's servo drive over RS-232.

[9]The effective spring constant of the flight train used by CSBF has been determined theoretically to have a value of 0.4 N·m/deg, and this agrees with measurements done during the 2012 flight of BLASTPol [68].

The effect of these terms can be understood by considering the simple case in which the payload in is *stop mode*, meaning that the pointing system is trying to servo to zero speed in azimuth. Under these conditions, there is no net torque on the payload from the pointing motors: $\tau_{RW} + \tau_{\text{piv}} = 0$. Differentiating both sides, we obtain $\dot{\tau}_{\text{piv}} = k\omega_{\text{piv}} = -I\ddot{\omega}_{RW}$ where $I$ is the reaction wheel moment of inertia. Substituting in the expression for $\omega_{\text{piv}}$ from Eq. 2.12 and re-arranging, we obtain

$$\ddot{\omega}_{RW} + kg_2\dot{\omega}_{RW} + \frac{k}{I}g_1\omega_{RW} = 0. \tag{2.13}$$

The dynamical equation for the reaction wheel angular velocity is that of a damped harmonic oscillator. The strength of the damping depends on $k$ and the gain $g_2$. The undamped frequency of oscillation is higher for stiffer $k$ and lower for larger $I$. A non-zero net torque on the payload acts as a driving term. Figure 2.24 below shows a measurement of the reaction wheel speed taken during lab testing that exhibits this behaviour.

Combining the two reaction wheel control terms (Eq. 2.12) with the scanning term results in an overall pivot control loop with three proportional gain terms: $\omega_{\text{piv}} = g_1\omega_{RW} + g_2\tau_{RW} - g_3\omega_{\text{req}}$. This equation is in terms of the dynamical quantities of interest. However, as with the reaction wheel control loop on the DSP, what the pivot control loop in pcm actually computes is $\Pi_{\text{piv}}$, the 16-bit pivot DAC level. This is translated into $\omega_{\text{piv}}$ by the combination of the pivot DAC and the servo drive. Furthermore, $\tau_{RW}$ is not measured directly, but estimated using the *reaction wheel's* DAC level $\Pi_{RW}$, which can be regarded as the *commanded* reaction wheel torque. Therefore, what pcm computes is more properly expressed as

$$\Pi_{\text{piv}} = g_1\omega_{RW} + g_2\Pi_{RW} - g_3\omega_{\text{req}}. \tag{2.14}$$

The parameters $g_1$, $g_2$, $g_3$, and $\omega_{SP}$ could be commanded in flight. The pivot gains were tuned starting shortly before the payload reached float altitude. Pivot velocity control showed excellent performance during scan tests conducted with a simulated flight train that has realistic dynamical properties. However, before the Spider 2015 flight, this

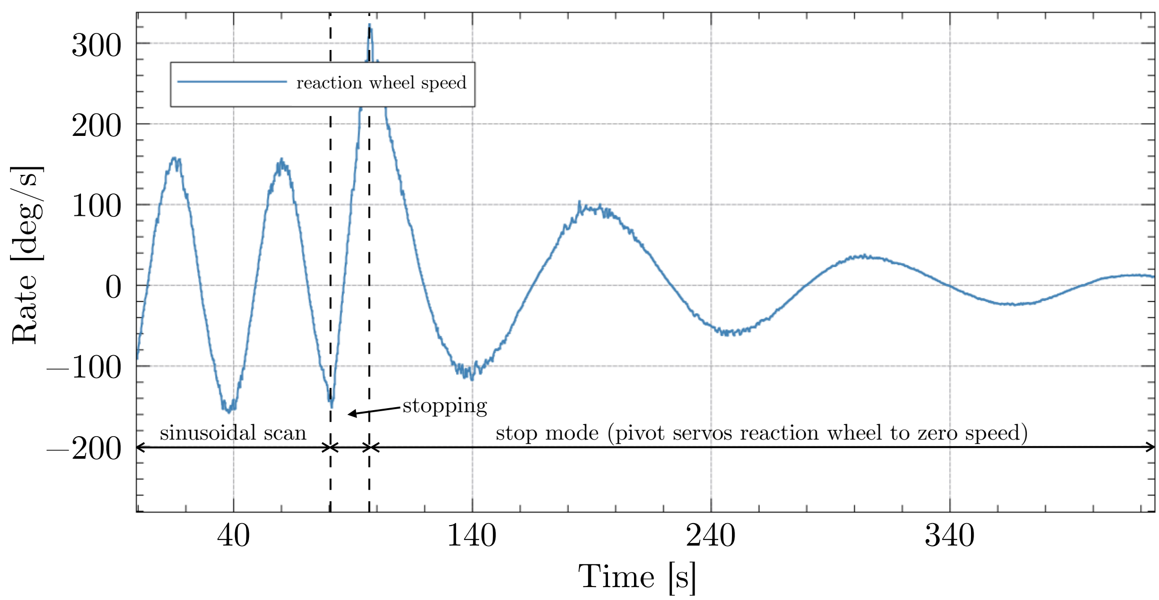


Figure 2.24: Reaction wheel speed vs. time during scan tests conducted with the payload suspended from a simulated flight train. In this plot, the gondola has just gone into stop mode from a sinusoidal scan, one period of which is visible on the left. The large angular acceleration involved perturbs the system of Eq. 2.13, causing an increase in reaction wheel speed that then damps out, returning to the setpoint value.

control system had never successfully flown before. Therefore, as a contingency against the inability to tune the system at float, a command was implemented that reverted the pivot's servo drive to torque mode, and switched the pivot control terms in pcm over to those in Eq. 2.11. Due to difficulties experienced with pivot motor operation on the ice (Section 2.10.1), the pivot was actually operated in torque mode for most of the flight, including during all CMB observation. Unlike during the early pointing system tests, torque mode had adequate performance in the final payload configuration. Velocity mode was only tested at the end of the flight, but proved successful (Section 2.10.4).

## 2.7 The Elevation Drive

### 2.7.1 Mechanical Design

The cryostat mounts to the gondola frame as shown in Figure 2.25. Trunnions mounted to the sides of the cryostat engage with SKF FSYE-3-NH pillow block bearings that rest atop the gondola joints at the centre of the port and starboard sides of the frame. The pillow blocks provide an axis for the cryostat to rotate in elevation (El). Aluminum rocker arms bolt to the outside ends of the trunnions, such that they are rigidly attached to the cryostat. Linear actuators mounted farther back on the gondola frame push on the rocker arms. Extension of the linear actuators lowers the cryostat's elevation, while retraction raises it. This system provides the mechanical advantage necessary to rotate the 3500 lb cryostat, even if it is unbalanced. The centre of mass of the cryostat shifts as cryogens boil off.

The linear actuators are VecTac E-Drive VT209-12 ball screw actuators rated for 900 $lb_f$ of thrust. This model has a 12" throw. The input shafts of these actuators are each connected with a flex coupling to a Stober P221SPR0070MT ServoFit™ Precision Planetary Gearhead with a 7:1 gear ratio. This gearbox couples to a NEMA 23 inline brake that is spring-actuated to prevent motion when the system is not powered by 24 VDC.

The brake couples to a Cool Muscle CM1-C-23L20C stepper motor from Myostat Motion Control Inc. The couplings between the gearbox and the inline brake, and between the brake and the stepper motor, are secured with shaft collars.

Elevation feedback sensing is provided by two Encoder Technology EA58-S absolute electro-optical encoders. These are rigidly mounted to the gondola frame, and their output shafts connect with flex couplings to shafts protruding from the sides of each rocker arm, directly inline with the elevation axis. Thus, the encoders measure the cryostat elevation angle relative to the gondola. The encoders have a 16-bit resolution, resulting in a minimum El step measurement of 0.0055°.

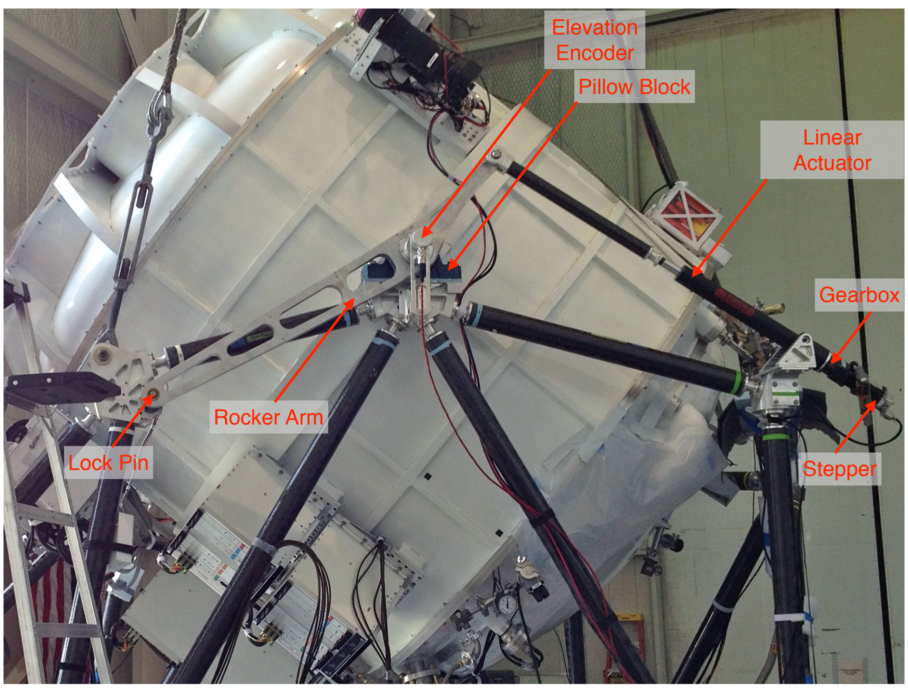


Figure 2.25: A photograph of the port side of the cryostat and gondola, with elements of the elevation drive labelled. These elevation drive components are repeated on the starboard side.

### 2.7.2 Motor Drive and Power

Each Cool Muscle stepper unit includes a built-in motor controller and a magnetic encoder for high resolution feedback on the rotor position, making each unit a fully-integrated servo system. By operating closed-loop with fast feedback, the motors avoid the drawbacks of other steppers, such as vibration and missed steps. Steppers can also be reliably commanded to any velocity simply by controlling the step rate. This feature enables a position control loop in which velocity is commanded based on the position error (Section 2.7.3). This type of control is suitable given that the only feedback sensing in elevation is the position angle from the encoders. An alternative would have been to use brushless DC torque motors controlled with pulse-width modulation (PWM). However, in the absence of fast velocity feedback, it is difficult to tune such a system to servo to position reliably, especially given the static friction in the system and the unbalanced load.

The Cool Muscles' controllers are commanded in pulse-per-step mode with two opto-isolated step and direction inputs. The step input receives a square wave pulse train whose frequency determines the step rate. The pulse frequency is computed by pcm as described in Section 2.7.3. This frequency is communicated to the ACS over the BLASTbus at ∼120 Hz. The square wave output is generated by an Altera FPGA on one of the BLASTbus motherboards in the ACS. The system imposes a maximum pulse frequency of 10 kHz.

The power supply to the elevation drive required some careful consideration. Most 24 V components on the experiment accept an input range of 18 to 36 VDC, and can therefore be powered directly from the flight batteries. The Cool Muscle steppers are an exception, requiring an input of 24 V ± 10%. As a result, each motor is powered by a Vicor DC-DC. Since Vicors cannot handle reverse current, which is produced by inductive loads such as the motors, a 200 Ω Dale 25 W power resistor is placed across the output of each DC-DC, in parallel with the motor, biasing the output current to be positive.

The output 24 VDC from each Vicor is provided in parallel to both a Cool Muscle stepper and its inline brake. Powering the brakes causes them to disengage, enabling motion. The Cool Muscle controllers take time to power on and energize the motor windings, providing holding torque. Therefore, a time delay was introduced between powering the motors and disengaging the brakes, in order to prevent the cryostat from falling. The two components remain powered by the same 24 VDC circuit, and the delay is implemented in hardware, so that the brake is never disengaged when the motor is unpowered. At the motor, the connection to the inline brake is interrupted by a Crydom CMX60D20 solid state relay whose input is controlled by a simple RC circuit (Appendix A.2). When 24 VDC is applied, it reaches the motors immediately, and the inline brakes approximately 200 ms later.

### 2.7.3 Position Control Algorithm

In every scan mode, pcm computes a requested elevation angle $\theta_{\rm req}$. At the BLASTbus data rate of $\sim$120 Hz, the elevation position control routine in pcm computes a rotation rate $\omega_{\rm el}$ for the elevation drive as follows

$$\omega_{\rm el} = {\rm sgn}(\Delta\theta) g_{\rm el} \sqrt{|\Delta\theta|}, \tag{2.15}$$

where the position error $\Delta\theta \equiv \theta_{\rm req} - \theta_{\rm enc}$, with $\theta_{\rm enc}$ being the mean of the starboard and port elevation encoder readings. The square root velocity-position profile corresponds to a constant negative acceleration to zero speed as the measured elevation approaches the target value. The magnitude of the acceleration depends on the gain $g_{\rm el}$, whose value could be commanded in flight, but had already been tuned during laboratory tests to be 0.2. When $|\Delta\theta| \leq 0.02°$, pcm simply sets $\omega_{\rm el}$ to zero.

The computed rotation rate is limited to $|\omega_{\rm el}| \leq \omega_{\rm max}$, where $\omega_{\rm max}$ is the rate of change of elevation corresponding to the maximum motor pulse frequency of 10 kHz. The gain is set high enough that the pulse rate saturates to this value during large elevation slews. To convert between rotation rate in El and motor pulse frequency, pcm must first compute

$$\frac{dx}{dt} = \frac{dx}{d\theta}\frac{d\theta}{dt} = \frac{dx}{d\theta}\omega_{\text{el}}, \tag{2.16}$$

where $x$ is the linear actuator extension. Due to the geometry of the system (Figure 2.26), $\frac{dx}{d\theta}$ varies with $\theta$ (Section 2.7.4) and pcm computes it at ∼120 Hz using the measured elevation $\theta_{\text{enc}}$. Having computed the rate of change of linear actuator extension, pcm can determine the pulse frequency as follows:

$$\frac{dx}{dt} = (\text{gear ratio}) \cdot \left(\frac{\text{linear actuator thrust (mm)}}{\text{rotation}}\right) \cdot \left(\frac{\text{stepper rotations}}{\text{pulse}}\right) \cdot f_{\text{el}}, \tag{2.17}$$

where $f_{\text{el}}$ is the pulse frequency. The gear ratio is 7:1, the actuator undergoes 5 rotations per inch of extension, and the number of motor steps per full rotation is set to 5000 in the firmware of the Cool Muscle controllers. Given these numbers, the maximum pulse frequency of 10 kHz results in motion of the linear actuators at 1.45 mm/s. Given the time it would take to cover the full 12" range of the linear actuator at that rate, the *average* elevation speed over the full mechanical El range of 11.6° to 56.4° is $\langle\omega_{\text{el}}\rangle \approx 0.2$ deg/s at the maximum pulse rate.

### 2.7.4 Elevation Angle Versus Linear Actuator Extension

Figure 2.26 shows the geometry of each side of the elevation drive system. The relevant triangle has edges $\ell_c$, $\ell_r$, and $L + x$. The distance $\ell_c$ is between the cryostat elevation axis and the linear actuator's trunnion-mount axis. The quantity $\ell_r$ is the length of the rocker arm (between its connection points to the actuator, and to the cryostat). Finally $L$ is the total length of the fully-retracted actuator, consisting of the length of the carbon fibre push rod $\ell_p$, and the length $\ell_l$ of the linear actuator itself. This side of the triangle has the additional length $x$, which is the extension of the linear actuator. The goal is to

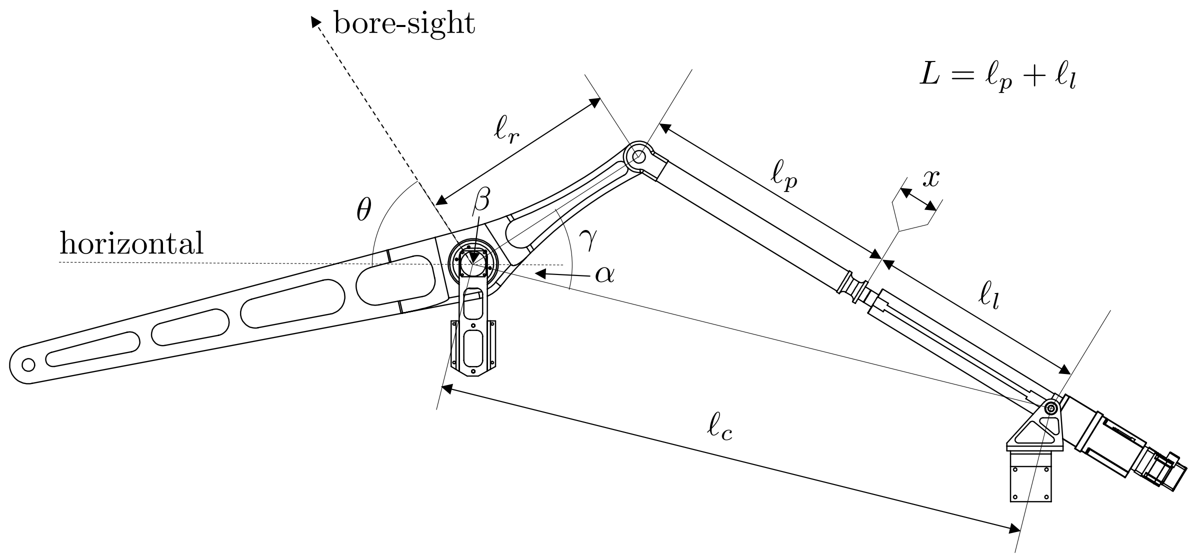


Figure 2.26: The geometry of the elevation drive system. The dimensions are as follows: $\ell_c = 1207.79$ mm, $\ell_r = 400$ mm, $\ell_l = 449.961$ mm, $\ell_p = 558$ mm, $\alpha = 14.06°$, and $\beta = 90.76°$. The original dimension of $\ell_p$ was 528.975 mm, but this increased due to the addition of strain gauges in-line with the linear actuators.

determine how the elevation angle of the bore-sight, $\theta$, varies with $x$. From the diagram:

$$(\gamma - \alpha) + \beta + \theta = 180°. \tag{2.18}$$

Applying the cosine rule to the aforementioned triangle gives an expression for the the interior angle $\gamma$

$$(L + x)^2 = \ell_c^2 + \ell_r^2 - 2\ell_c\ell_r \cos\gamma \tag{2.19}$$

Substituting Eq. 2.19 into Eq. 2.18 and re-arranging, we obtain:

$$\theta(x) = 180° + \alpha - \beta - \arccos\left(\frac{\ell_c^2 + \ell_r^2 - (L + x)^2}{2\ell_c\ell_r}\right) \tag{2.20}$$

This relation results in the variation of $\theta$ with $x$ shown in Figure 2.27. The green curve shows the case of the original push rod length, while the blue curve is for the push rod whose length has been increased by the addition of in-line strain gauges. The latter configuration was flown. The figure also shows the deviation of Eq. 2.20 from a simple linear relation given by

$$\theta_{\mathrm{lin}}(x) = \theta_{\mathrm{max}} - \left(\frac{\theta_{\mathrm{max}} - \theta_{\mathrm{min}}}{x_{\mathrm{max}} - x_{\mathrm{min}}}\right) x \tag{2.21}$$

## 2.8 Observing Modes

As discussed in Section 2.6.1.2, the fast reaction wheel motor torque control loop on the DSP servos the payload's azimuthal angular velocity to a requested value $\omega_{\mathrm{req}}$. Higher-level control algorithms, implemented in pcm on the flight computers, produce various scanning and pointing motions by varying $\omega_{\mathrm{req}}$ as a function of time or other observables. The observing modes used in flight are discussed below, with the exception of the main strategy that was developed for CMB observations during flight. This mode is covered in more detail in Section 2.9.

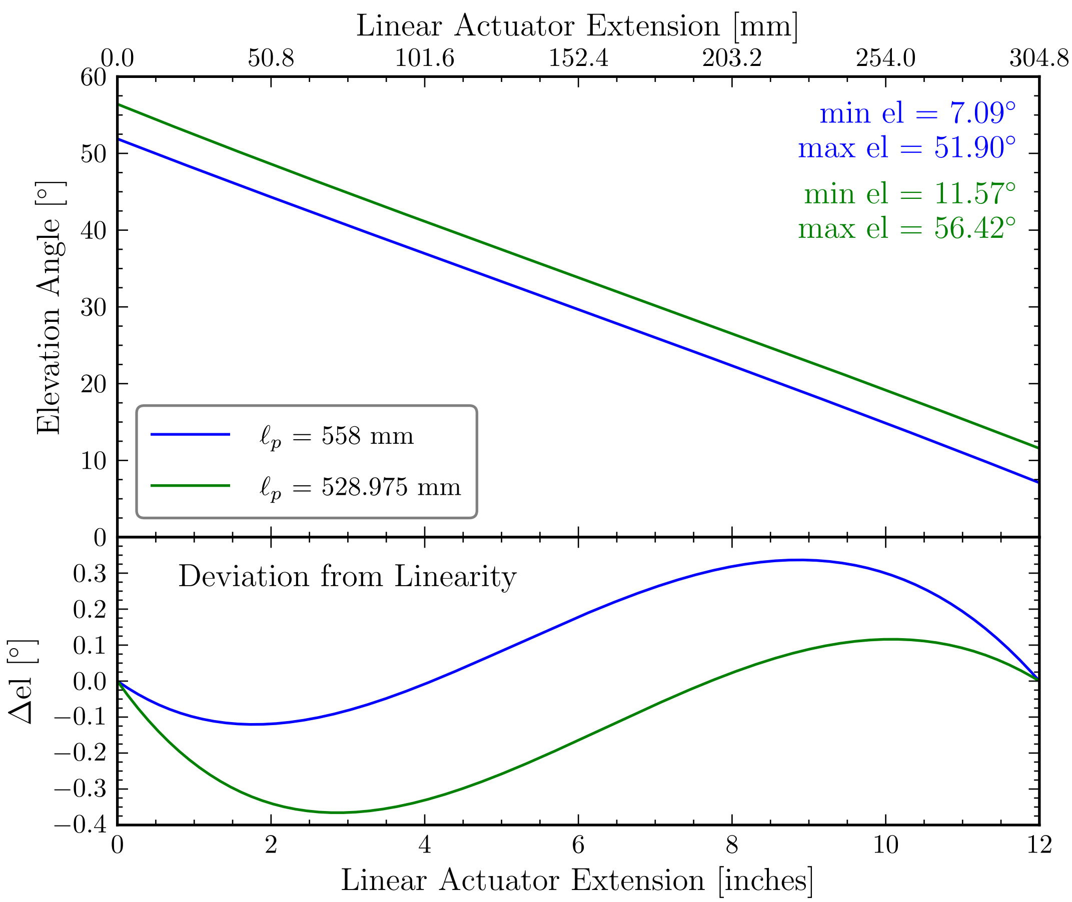


Figure 2.27: *Top* – The variation of elevation angle $\theta$ with linear actuator extension $x$ for the two cases of the original push rod length (*green*), and the modified length after the inclusion of strain gauges (*blue*). *Bottom* – The deviation of the $\theta(x)$ function from linearity for the two cases.

### 2.8.1 Az-El Goto

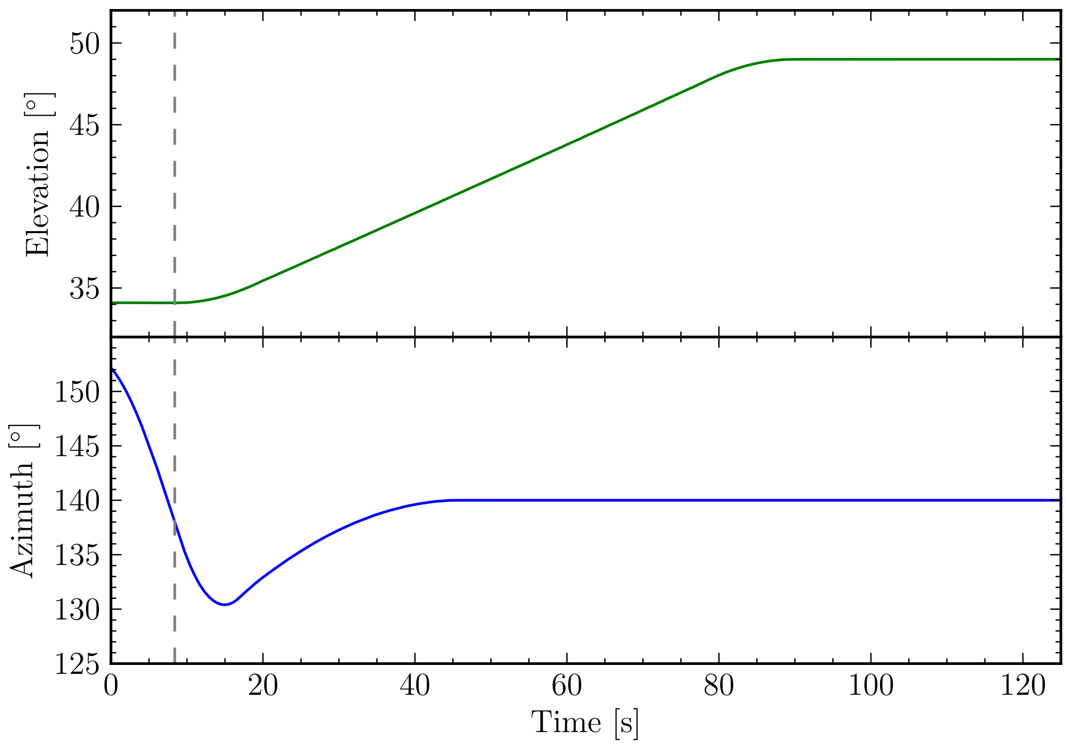


Figure 2.28: Az and El from the in-flight pointing solution, responding to the command `az_el_goto 140 49`. The dashed grey vertical line indicates when the command was received.

A command to point to a particular position in horizon coordinates was carried over from BLASTPol. The `az_el_goto` command is sent with the input parameters being the requested Az and El to go to: $\phi_{\mathrm{req}}$ and $\theta_{\mathrm{req}}$. The Az velocity request is computed from the error between the requested and measured positions:

$$\omega_{\mathrm{req}} = \mathrm{sgn}(\Delta\phi)\frac{g_{\mathrm{pt}}}{10\ 000}\sqrt{|\Delta\phi|} \tag{2.22}$$

where $\Delta\phi \equiv \phi_{\mathrm{req}} - \phi$, with $\phi$ being the current Az of the payload, according to the in-flight pointing solution. The pcm algorithm chose to slew between $\phi$ and $\phi_{\mathrm{req}}$ by either increasing or decreasing Az: whichever direction did not cross the sun azimuth during the slew. The pointing gain $g_{\mathrm{pt}}$ was a 16-bit integer whose value could be commanded in flight. Given the input $\theta_{\mathrm{req}}$, the elevation position servo operated as described in Section 2.7.3. Figure 2.28 shows an example of an Az-El Goto move from the 2015 flight.

#### 2.8.1.1 Anti-Sun

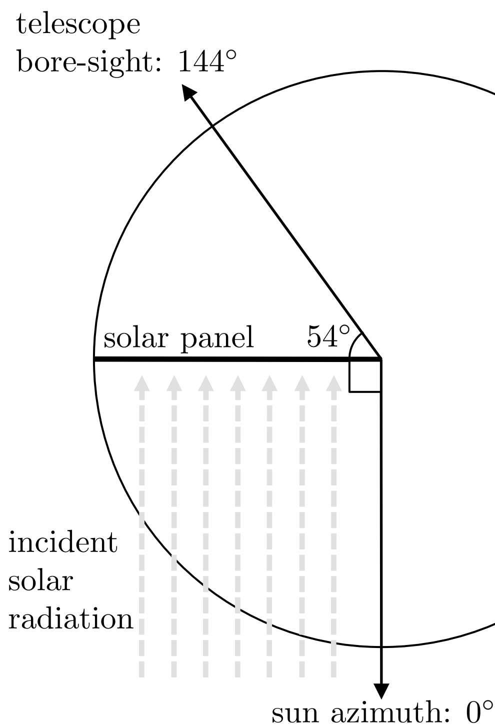


Figure 2.29: The telescope bore-sight azimuth angle necessary for normal incidence of radiation on the solar panels

A special case of Az-El Goto, also carried over from BLASTPol, was Anti-Sun. When the `antisun` command was sent, pcm would servo the payload to a specific azimuth position in the roughly anti-sun direction. The azimuth was chosen to make solar radiation be incident as close to normal to the solar panels as possible. As shown in Figure 2.29, the panel opening angle, projected onto the horizontal plane, was measured to be 54°. As a result, the bore-sight azimuth relative to the sun was 144° in Anti-Sun mode, rather than 180°. This pointing mode was used near the end of ascent during motor tuning, and again near the end of the flight, when additional pointing and power system tests were conducted. A limitation of this implementation is that Az-El Goto servos to a constant Az position $\phi_{\text{req}} = \phi_{\text{sun}} + 144°$, computed at the time the command is sent. Keeping the payload anti-sun by this definition for an extended period of time would have required periodically re-sending the command to update the goto position, to compensate for the sun's drift.

### 2.8.2 RA-Dec Goto

A command to point to a particular position in equatorial coordinates was also carried over from BLASTPol. The `ra_dec_goto` command is sent with the input parameters being the requested right ascension (RA) and declination (Dec) to go to: $\alpha_{\text{req}}$ and $\delta_{\text{req}}$. At the BLASTbus frame rate of ~120 Hz, given the payload latitude and LST, pcm transformed these (RA, Dec) coordinates into (Az, El), and then servoed the bore-sight direction to these (continuously updating) coordinates in the same manner as in Az-El Goto mode. This mode was less of a true RA-Dec Goto mode in Spider than in BLASTPol. The Spider elevation drive would only update the elevation if it became more than 0.02° different from the target elevation, whereas the BLASTPol elevation drive would servo continuously to the target elevation. Therefore, in this mode Spider would drift slowly away from the requested RA and Dec until the El error became large enough to trigger a correction. This behaviour was deemed acceptable, since the ability to track a point on the sky very precisely was far less necessary for Spider than it had been for BLASTPol. RA-Dec Goto mode nevertheless found use early in the Spider flight, during line of sight. The command was used to point at RCW38, a bright source in the Galactic plane that was useful for detector tuning and calibration. Once the telescopes had slewed to the Az and El of the source, a sinusoidal Az scan with El stepping (Section 2.8.4), centred on these coordinates, was initiated to raster over the source.

### 2.8.3 Drift

The `drift` command causes the payload to be servoed to a constant azimuthal angular velocity, which is an input parameter. It is therefore the simplest pointing mode in pcm, with $\omega_{\text{req}}$ = constant. Most often, the `stop` command was used, which commands the special case of Drift mode for which $\omega_{\text{req}} = 0$. Stop mode is typically the first pointing mode entered when the Az motors are tested for the first time after assembling the payload. It was also used near the end of ascent, when the pointing motor control loops were enabled and tuned for the first time at float.

### 2.8.4 Sinusoidal Azimuth Scan

A scan mode in pcm that carries out a sinusoidal scan in Az with steps in El at scan turn-arounds was developed for laboratory tests of the pointing system. It was also used as a subcomponent of the science scan mode (Section 2.9). A sinusoidal profile was chosen in order to make it possible, within the limitations of motor torque, to carry out scans of the required speed and width. In particular, SPIDER must scan fast enough to modulate the sky signal into frequencies above the $1/f$ knee of the detector noise spectra. A scan with a sufficiently-fast, *constant* speed would perhaps be ideal, but would require more acceleration at scan turn-arounds than the pointing motors can provide.

The sinusoidal scan is completely defined by three parameters: the scan centre, $\phi_c$, the scan amplitude $A$, and the peak Az angular acceleration, $\alpha_{\text{pk}}$. During a scan, the variation of Az angle $\phi$ with time is of the form

$$\phi(t) = \phi_c - A\cos\left(\frac{2\pi}{T}t\right), \tag{2.23}$$

where $T$ is the scan period, and the phase of the scan has been chosen arbitrarily. The corresponding expressions for angular velocity and acceleration are

$$\dot{\phi}(t) = \omega(t) = \frac{2\pi}{T}A\sin\left(\frac{2\pi}{T}t\right) \tag{2.24}$$

and

$$\ddot{\phi}(t) = \alpha(t) = \frac{4\pi^2}{T^2} A \cos\left(\frac{2\pi}{T} t\right). \tag{2.25}$$

The angular velocity and acceleration can also be written as $\omega(t) = \omega_{\mathrm{pk}} \sin\left(\frac{2\pi}{T} t\right)$ and $\alpha(t) = \alpha_{\mathrm{pk}} \cos\left(\frac{2\pi}{T} t\right)$. Comparing these two equations with Eqs. 2.24 and 2.25 above, it can be shown that

$$T = 2\pi \sqrt{\frac{A}{\alpha_{\mathrm{pk}}}} \tag{2.26}$$

and

$$\omega_{\mathrm{pk}} = \sqrt{\alpha_{\mathrm{pk}} A}. \tag{2.27}$$

Since the available pivot and reaction wheel motor torque ultimately limits the maximum payload angular acceleration, $A$ and $\alpha_{\mathrm{pk}}$ are the base parameters of the scan mode. The values of $\omega_{\mathrm{pk}}$ and $T$ result from the values of these base parameters, rather than being specified directly. This reduces the probability of a sinusoidal scan with an unattainable value of $\alpha_{\mathrm{pk}}$ being commanded. For the purposes of testing, $\alpha_{\mathrm{pk}} = 0.8$ deg·s$^{-2}$ and $A = 45°$ were long considered the canonical "flight-like" scan parameter values, resulting in $T = 47$ s and $\omega_{\mathrm{pk}} = 6$ deg/s. However, during flight, in the science scan mode, $A$ varies with time based on other parameters, and $\alpha_{\mathrm{pk}}$ had to be lowered for practical reasons (Sections 2.9 and 2.10).

When developing the sinusoidal scan routine for pcm, the goal was to implement the scan described above in as stateless a way as possible. In order to determine $\omega_{\mathrm{req}}$ at any time, pcm needs to know only the current payload azimuth, and the direction of motion. The latter is measured by the rate gyroscopes, while the former is computed in the in-flight pointing solution [69]. Since the velocity request is computed based on position, it is necessary to determine the velocity vs. position profile for the sinusoidal scan. Taking $\omega(t) = \omega_{\mathrm{req}}$, Eqs. 2.23 and 2.24 can be combined to produce

$$\frac{(\phi - \phi_c)^2}{A^2} + \frac{\omega_{\mathrm{req}}^2}{\omega_{\mathrm{pk}}^2} = 1. \tag{2.28}$$

The velocity vs. position curve is an ellipse with semi-major and semi-minor axes given by $A$ and $\omega_{\mathrm{pk}}$.

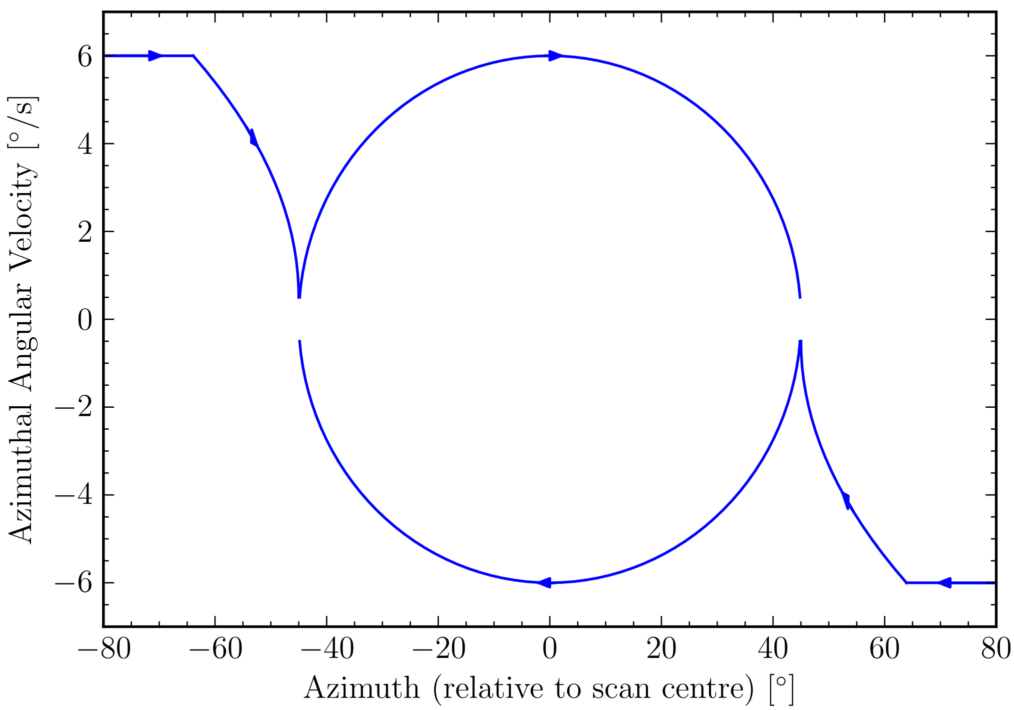


Figure 2.30: An example of $\omega_{\mathrm{req}}$ vs. Az (relative to the scan centre) as determined by the sinusoidal scan algorithm in pcm. The canonical values of $A = 45°$ and $\omega_{\mathrm{pk}} = 6$ deg/s were used.

Figure 2.30 depicts velocity vs. position in the six regimes used by the scan algorithm in pcm. These are also summarized in equation form in Table 2.4. In this table, $\phi_l \equiv \phi_c - A$ is the left Az scan endpoint, and $\phi_r \equiv \phi_c + A$ is the right Az scan endpoint. When the payload azimuth is beyond the scan endpoints on either side, the velocity follows a square root profile with acceleration $\alpha_{\mathrm{pk}}$, up to a maximum speed of $\omega_{\mathrm{pk}}$. Within the scan endpoints, the velocity follows the ellipse corresponding to a sinusoidal scan (Eq. 2.28). The algorithm follows the elliptical profile until $|\omega_{\mathrm{req}}| < \omega_{\mathrm{min}}$, which occurs at a distance of $\Delta\phi_{\mathrm{turn}}$ before the turn-around. Within this distance of a scan endpoint, the speed is constant at $\omega_{\mathrm{min}}$; $\omega_{\mathrm{req}}$ simply flips sign from $\pm\omega_{\mathrm{min}}$ to $\mp\omega_{\mathrm{min}}$. Therefore, the algorithm jumps from the top to the bottom elliptical branch, or vice versa. In pcm, $\omega_{\text{min}} = 0.05$ deg/s, which is considered to be the smallest reliably-measurable speed given gyro noise and offsets. In Figure 2.30, $\omega_{\text{min}}$ has been exaggerated to 0.5 deg/s for clarity. Not shown in the figure is a commandable overshoot $\delta\phi$ which the payload can travel beyond the scan endpoints without pcm switching over from the elliptical to the square root profile.

Table 2.4: The variation of $\omega_{\text{req}}$ with payload position and velocity in the sinusoidal scan algorithm

| Velocity Request ($\omega_{\text{req}}$) | Az Position ($\phi$) & Velocity ($\omega_{\text{gond}}$) | Scan Regime |
|---|---|---|
| $+\sqrt{2\alpha_{\text{pk}}[(\phi_l - \delta\phi) - \phi]} + \omega_{\text{min}}$ | $\phi < (\phi_l - \delta\phi)$ | beyond left scan endpoint |
| $+\omega_{\text{min}}$ | $\phi_l < \phi \leq (\phi_l + \Delta\phi_{\text{turn}})$ | in left turn-around zone |
| $+\omega_{\text{pk}}\sqrt{1 - (\phi - \phi_c)^2/A^2}$ | $(\phi_l + \Delta\phi_{\text{turn}}) < \phi < (\phi_r - \Delta\phi_{\text{turn}})$; $\omega_{\text{gond}} > 0$ | in scan range and moving from left-to-right |
| $-\omega_{\text{pk}}\sqrt{1 - (\phi - \phi_c)^2/A^2}$ | $(\phi_l + \Delta\phi_{\text{turn}}) < \phi < (\phi_r - \Delta\phi_{\text{turn}})$; $\omega_{\text{gond}} < 0$ | in scan range and moving from right-to-left |
| $-\omega_{\text{min}}$ | $(\phi_r - \Delta\phi_{\text{turn}}) \leq \phi < \phi_r$ | in right turn-around zone |
| $-\sqrt{2\alpha_{\text{pk}}[\phi - (\phi_r + \delta\phi)]} - \omega_{\text{min}}$ | $\phi > (\phi_r + \delta\phi)$ | beyond right scan endpoint |

This scan mode was entered by sending the `sine_scan` command, which takes as input parameters $A$, $\phi_c$, $\theta_{\text{req}}$ (the starting El), $n$ (the number of Az half-scans per El step),

$\delta\theta$ (the El step size), and $N$ (the total number of steps before resetting the elevation). Other parameters such as $\alpha_{\rm pk}$, $\delta\phi$, and a tunable phase delay $\Delta t$ in units of BLASTbus frames, were specified using the `set_scan_params` command. The phase delay was used to propagate the present value of the Az solution $\phi_0$ forward in time before using it to compute $\omega_{\rm req}$: $\phi = \phi_0 + \omega_{\rm gond}\Delta t$. This correction made the scan algorithm somewhat predictive, compensating for latency in the system. The delay was set to 2.75 frames in flight. The parameters of `set_scan_params` were grouped into a separate command because they are common to both `sine_scan` and `spider_scan`. The latter is discussed in the next section.

## 2.9 The SPIDER Observing Strategy

The scan strategy to be used for CMB observations has the sinusoidal Az scan from Section 2.8.4 as a basis, but includes the following additional elements. A quadrangular region (hereafter simply "the box"), within which observations are to be confined, is defined on the sky. This box is the thick blue outline in Figure 2.31. The coordinates of the four corner points of the box in RA and Dec are specified. These corner points are connected by great-circle arcs. This method makes it easy to avoid observing too close to the Galactic plane; the defined scan region simply does not encompass these areas. The box also determines sky coverage, subject to additional constraints from the sun azimuth and mechanical elevation limits (Figure 2.31).

A specific point within the box, known as the "track point", is chosen in RA and Dec. The instantaneous sinusoidal Az scan is constrained to pass through this point, and to terminate on the edges of the box. Therefore, at the BLASTbus frame rate of $\sim$120 Hz, pcm transforms from equatorial to horizon coordinates to determine the track point elevation, and the scan endpoints in azimuth $\phi_r$ and $\phi_l$ of a line at this elevation that is confined to the box. The requested elevation angle $\theta_{\rm req}$ (Section 2.7.3) is set to

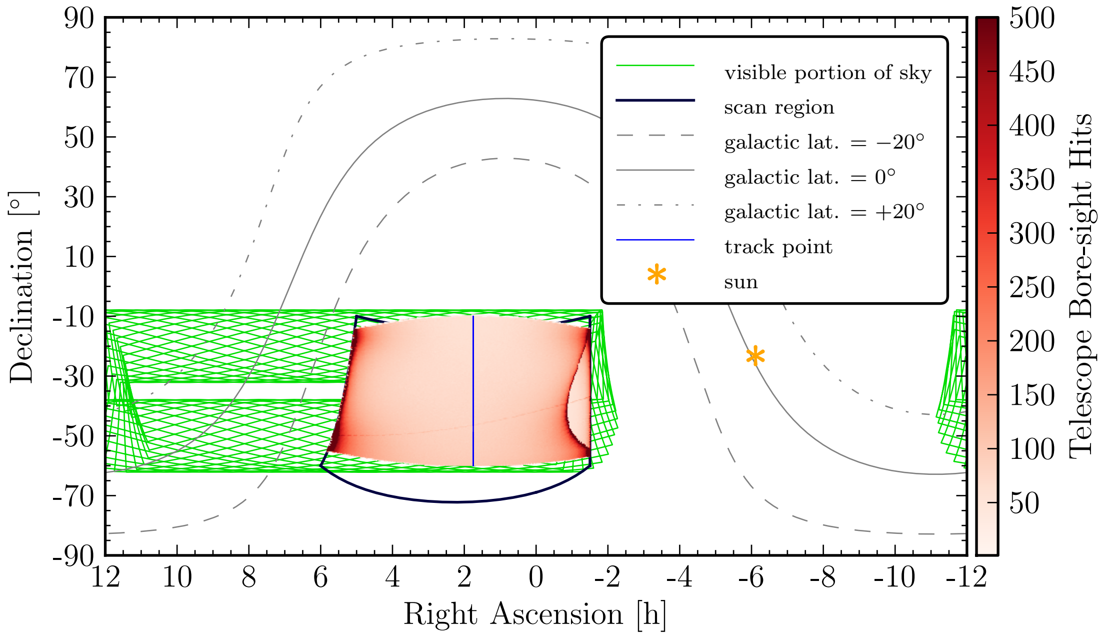


Figure 2.31: Output from a simulation of the SPIDER scan strategy spanning 24 solar hours beginning Dec. 20, 2014 at 00:00:00 UTC. This simulation occurs at lat. = 77.85° S, lon. = 166.67° E, which are the coordinates of McMurdo Station, Antarctica. The dark blue outline is the defined scan region. The blue vertical line traces out the path of the track point. Every hour, a green box has been drawn whose edges are SPIDER's El limits and Az limits relative to the sun. The union of these green boxes encompasses the total area of sky visible to SPIDER, in principle, from this location. This simulation used El limits of $20° \leq \theta \leq 50°$. The Az limit is $|\phi - \phi_{\rm sun}| \geq 70°$ on the port side, and $|\phi - \phi_{\rm sun}| \geq 90°$ on the starboard side. The intensity map shows the number of hits for the *telescope bore-sight* in pixels of HEALPix [70] $N_{\rm side} = 256$, indicating the area of sky actually observed. The grey curves are lines of constant Galactic latitude.

the computed track point elevation. The results of these coordinate transformations will change with latitude and with local sidereal time (LST). Therefore, the science Az scan is like the one described in Section 2.8.4, but with time-variable scan endpoints $\phi_r$ and $\phi_l$, and hence a time-variable amplitude ($A$) and scan centre ($\phi_c$).

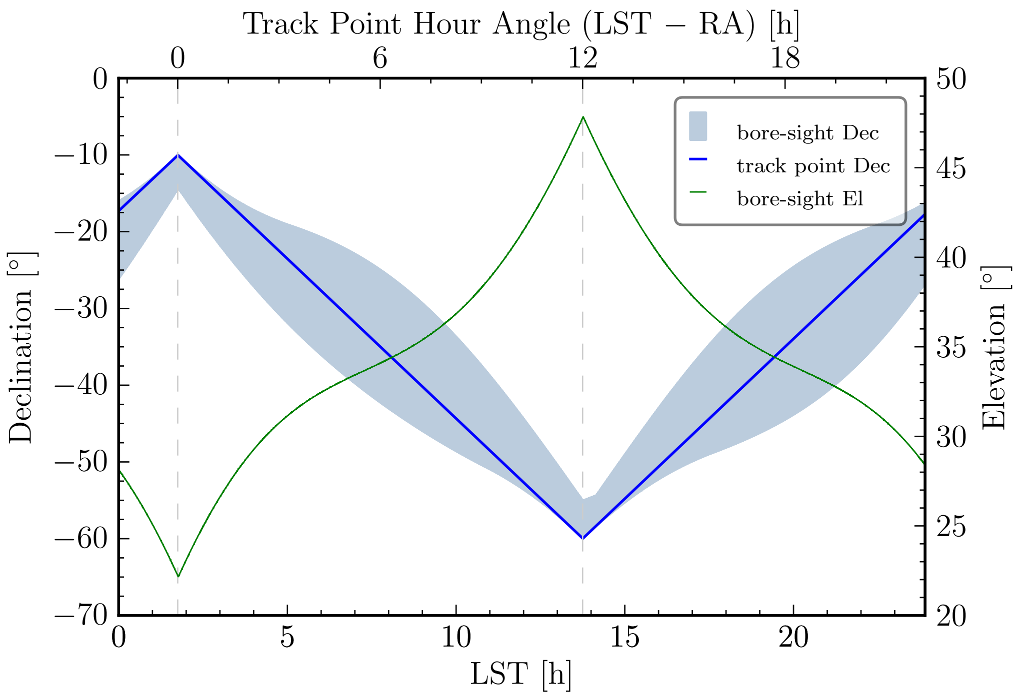


Figure 2.32: Output from a scan simulation showing the variation of the track point and bore-sight Dec with LST. Also shown is bore-sight El vs. LST (note the different vertical scale). Variation of Dec at a constant rate with LST requires variation of El at a variable rate. In this simulation, the track point RA was 1.75 h. Therefore, at LST = 1.75 h (hour angle = 0 h), the scan box has risen to its highest El. The lowest-El (highest-Dec) portion of it is observed then. At LST = 13.75 h (hour angle = 12 h) the scan box has set to its lowest El. The highest-El (lowest-Dec) portion of it is observed then. This phasing maximizes sky coverage.

The RA of the track point is kept constant, while its Dec is varied back and forth between the top and bottom edges of the box at a constant rate with LST. It reaches its highest declination when the hour angle (LST − RA) of the track point is zero, meaning that the centre of the box reaches its highest point, crossing the meridian. The declination of the track point reaches its lowest value 12 sidereal hours later. This phase of the Dec vs. LST variation maximizes sky coverage (Figure 2.32). The purpose of moving the track point is to ensure even coverage of the box in Dec, filling in the gaps between rows of detector beams on the sky. Rather than drifting continuously, Spider's elevation angle is updated at every $n^{\mathrm{th}}$ scan turn-around in order to keep up with the motion of the track point.

In flight, the science scan mode was entered by sending the `spider_scan` command, which takes as input parameters the box corners $(\alpha_i,\ \delta_i)$, $i \in [1, 4]$, the track point RA $\alpha_{\mathrm{track}}$, the Dec limits $\delta_{\mathrm{top}}$ and $\delta_{\mathrm{bot}}$, and $n$, the number of Az scan turn-arounds per El step. The pcm code was also adapted into a stand-alone program, `scan_sim`, which produced simulated Spider bore-sight pointing timestreams. The program used the above input parameters, along with latitude, longitude, simulation start date, and duration, to simulate observations. The `scan_sim` code was used for debugging the scan control algorithm and investigating available sky coverage. It was also essential for flight planning, determining the scan box and other `spider_scan` parameters to be used for a given launch date and latitude range. Figures 2.31 and 2.32 were produced using `scan_sim` output.

### 2.9.1 HWP Stepping

A strategy for HWP stepping is given in Table 2.5. In this table, integer sidereal days begin when the bore-sight is at the bottom of the box, moving upwards in Dec. Half-integer sidereal days begin once the scan has reached the top of the box and begins moving back downward. Three sequences of HWP angles were used. Receivers whose HWPs were stepped using the sequence in the leftmost column are labelled + receivers, based on their starting polarization angles. In the first four sidereal half-days, they have

Table 2.5: A strategy for stepping the SPIDER HWPs

| Sidereal Day | HWP Angle [°] X3 & X6 (+) | HWP Angle [°] X1 & X4 (×) | HWP Angle [°] X2 & X5 |
|---|---|---|---|
| 0.0 | 0.0 + | 22.5 × | 22.5 × |
| 0.5 | 22.5 × | 0.0 + | 0.0 + |
| 1.0 | 45.0 + | 67.5 × | 45.0 + |
| 1.5 | 67.5 × | 45.0 + | 67.5 × |
| 2.0 | 22.5 × | 45.0 + | 45.0 + |
| 2.5 | 45.0 + | 22.5 × | 22.5 × |
| 3.0 | 67.5 × | 90.0 + | 67.5 × |
| 3.5 | 90.0 + | 67.5 × | 90.0 + |

vertical and horizontal polarization sensitivity on the sky during up-going scans, and $+45°$ and $-45°$ sensitivity during down-going scans. The opposite is true for $\times$ receivers (middle column). For both types of receivers, the mapping of $+$ and $\times$ to up-going and down-going scans is swapped in the next four sidereal half-days.

The two detectors of orthogonal polarization sensitivity in a given spatial pixel are referred to as $A$ and $B$ detectors. In a $+$ receiver, an $A$ bolometer (represented by red line segments in Table 2.5) will measure $+Q$ during the up-going scan in the first half of day 0. The 22.5° HWP step will then switch its sensitivity to $+U$ on the down-going scan. The same bolometer will undergo the sequence $-Q$ and $-U$ on sidereal day 1. During the same time period, a $B$ bolometer (blue line segments in the table) will undergo the sequence ($-Q$, $-U$, $+Q$, $+U$). Therefore, after two sidereal days, every bolometer in the $+$ receiver should have made an independent measurement of the Stokes polarization vector in each sky pixel. This statement is true for the the $\times$ receivers as well, differing only in the order of the measurements.

In sidereal days 2 and 3, the same sequence of HWP steps occurs as in days 0 and 1, but shifted by $+22.5°$, so that if a given Stokes parameter was measured by a given bolometer during up-going scans in the first set of four half-days, it is measured during

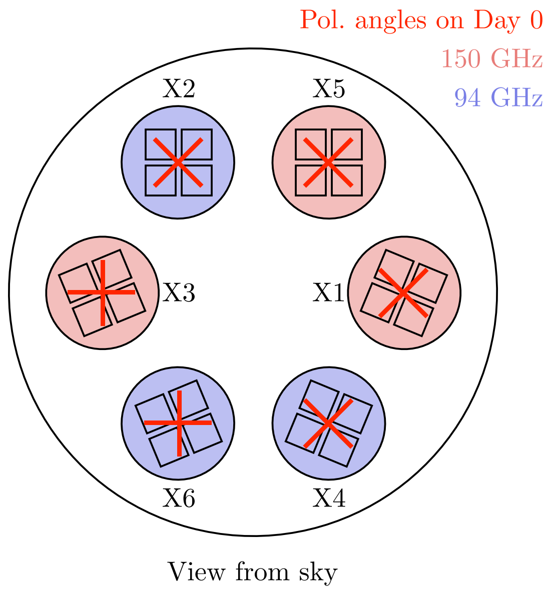


Figure 2.33: The starting polarization orientations of each receiver on the sky, denoted by orthogonal red lines

down-going scans in this set, and vice versa. In sidereal days 4–7, the HWP angles are shifted by +90° from the sequence in days 0–3. This new set of positions (90°, 112.5°, 135°, and 157.5°) may not be perfectly equivalent to the first set of positions due to HWP non-idealities. Therefore, it is after *eight* sidereal days that both up-going and down-going scans of the box exist for each unique HWP position.

Referring to the middle and bottom rows of inserts in Figure 2.33, the assignment of the + and × sequences to receivers resulted in a (+, ×) pair in each row. There was one such row per band: X3 and X1 in the middle row at 150 GHz, and X6 and X4 in the bottom row at 94 GHz. The assignment of one of each polarization type per row was done in case the inserts became inoperative sequentially by row from top to bottom as cryogen liquid level decreased during flight. With each band having one receiver of each type, the question arose of what HWP stepping sequence to use for the third receiver in each band. Neither + nor × would suffice. For example, if X2 and X5 were simply made + receivers, then in days 0 and 1, the scans with preferentially more + oriented receivers would always be up-going. Therefore, a third HWP stepping sequence (rightmost column of Table 2.5) was devised for X2 and X5, so that in every block of four half-days, receivers of each polarization type were evenly-distributed among the scan directions.

The strategy described above was generally followed in flight, with only minor changes resulting from practical considerations. Due to the properties of individual rotation mechanisms, some HWPs were stepped only in one direction, resulting in long turns, or were stepped in a manner that avoided passing through certain positions. Also, the assignment of the integer-day moves from Table 2.5 to up-going scans and the half-integer-day moves to down-going scans was not strictly necessary; they could be swapped in principle. The actual phasing that occured in flight was a matter of convenience, since HWP sidereal day 0 could not begin until after line of sight testing was complete and CMB observations had begun.

### 2.9.2 Observations in Flight

In flight, the scan parameters differed from the nominal case discussed above. The El range[10] was determined empirically to be from 20° to 49°. Below 20°, the cryostat bottom dome made contact with the rear suspension rope. Above 49°, the spreader bar entered the BSC field of view. It was also determined during testing just prior to launch that the peak Az scan angular acceleration ($\alpha_{\rm pk}$) would have to be lowered from 0.8 deg·s$^{-2}$ to between 0.4 and 0.5 deg·s$^{-2}$. Near the end of ascent, when the pointing motors were first activated and tuned, $\alpha_{\rm pk} \leq 0.3$ deg·s$^{-2}$ was found to be the practical upper limit. These changes are explained in Section 2.10.

Another difference in flight was that the actual scan region differed from the box shown in Figure 2.31. Six different scan boxes were used in flight (Table 2.6). Since launch was not until January 1, 2015, the sun had moved eastward across the sky (leftward in

[10]While this is significantly narrower than the purely mechanical El limits (Figure 2.27), it is nearly the same as the limits of 20° and 50° expected to be imposed by ground albedo and the balloon respectively.

Table 2.6: The unique SPIDER scan boxes from the 2015 flight, listed by the time each was first commanded. All RA ($\alpha$) values are listed in hours and Dec ($\delta$) values in degrees. The Dec limits remained between $-13°$ and $-56°$ throughout the entire flight, and the number of Az half-scans per El step was always 3.

| Scan | Time (UTC) | $(\alpha_1, \delta_1)$ | $(\alpha_2, \delta_2)$ | $(\alpha_3, \delta_3)$ | $(\alpha_4, \delta_4)$ | Track Pt. $\alpha$ |
|---|---|---|---|---|---|---|
| 1 | 2015-01-01 15:54:58 | (0.25, 0) | (5, 0) | (6, $-60$) | (0.25, $-60$) | 2.75 |
| 2 | 2015-01-03 20:08:02 | (1.25, 0) | (6.25, 0) | (7.25, $-60$) | (1.25, $-60$) | 3.75 |
| 3 | 2015-01-04 02:42:33 | (1.25, 0) | (5.33, 0) | (7.25, $-60$) | (1.25, $-60$) | 2.75 |
| 4 | 2015-01-06 18:45:31 | (1.5, 0) | (5.58, 0) | (7.5, $-60$) | (1.5, $-60$) | 2.75 |
| 5 | 2015-01-10 18:01:11 | (2.5, 0) | (4.75, 0) | (5, $-60$) | (2.5, $-60$) | 3.625 |
| 6 | 2015-01-15 00:47:39 | (3.5, 0) | (5.75, 0) | (6, $-60$) | (3.5, $-60$) | 3.625 |

Figure 2.31). Therefore, the areas of sky with RA $<$ 0 h shown in that figure were no longer visible. Furthermore, the minimum azimuth relative to the sun on the port side, designed to be 70° with the sunshield wing, was measured on the ice to be 78°. At smaller angles, sunlight began shining on the rims of the telescope baffles. This statement is also dependent on the sun elevation and cryostat elevation. To be conservative, the port sun Az limit was set to 80° before launch. Even with this limit in effect, there were concerns in flight that “clipping” seen in some detector timestreams was a result of pointing too close to the sun, perhaps from the sun glinting around the sunshield edges, or due to a systematic offset in the Az coarse sensor calibration. The sun continued its eastward motion over the course of the flight as well. For all of these reasons, the box was adjusted eastward (and hence Galaxy-ward) twice (Figure 2.34, scan boxes 2 and 4).

Once a combined eight sidereal days had been completed on the third and fourth scan boxes, it was decided to narrow the box in RA by a factor of $\sim$2. It was unclear how many days’ worth of cryogens remained, and doing this would concentrate the remaining integration time in the best non-Galaxy portion of the available sky. This SPIDER deep scan was carried out for four sidereal days and then shifted leftwards (scan boxes 5 and

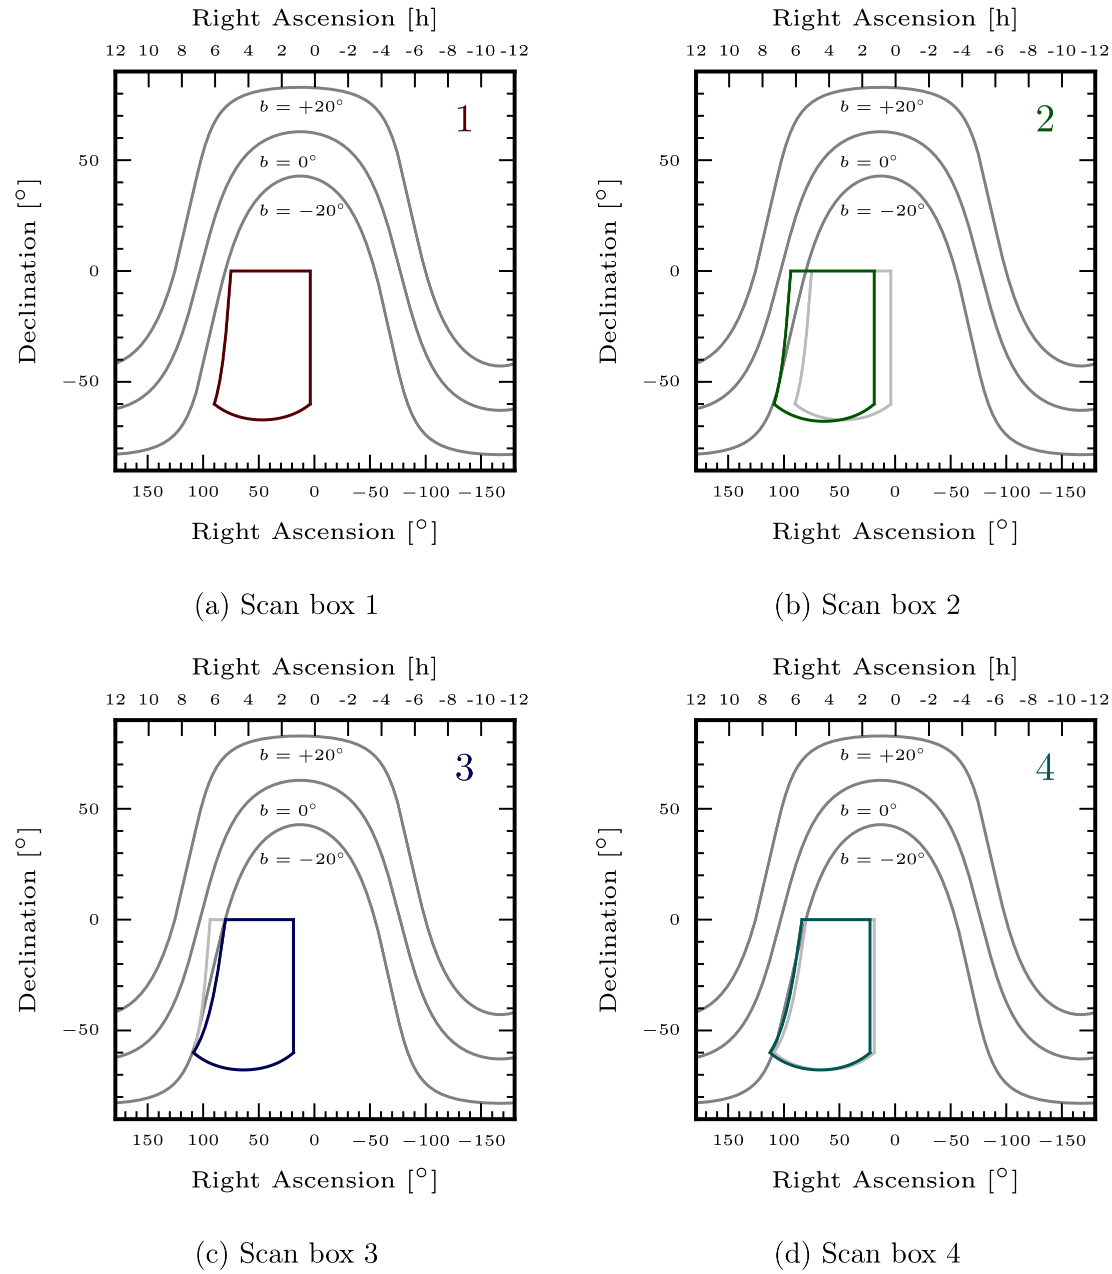


(a) Scan box 1

(b) Scan box 2

(c) Scan box 3

(d) Scan box 4

Figure 2.34: The six scan boxes used during flight. In each panel, the scan box indicated by number is coloured, and the previous scan box is light grey. (*Continued below*).

6 in Figure 2.34). Together, the two deep scan boxes covered most of the area of the wide box. "Dithering" the deep scan box within the wide one in this manner was done to produce more uniform coverage in the central area of the wide scan. In the wide scan

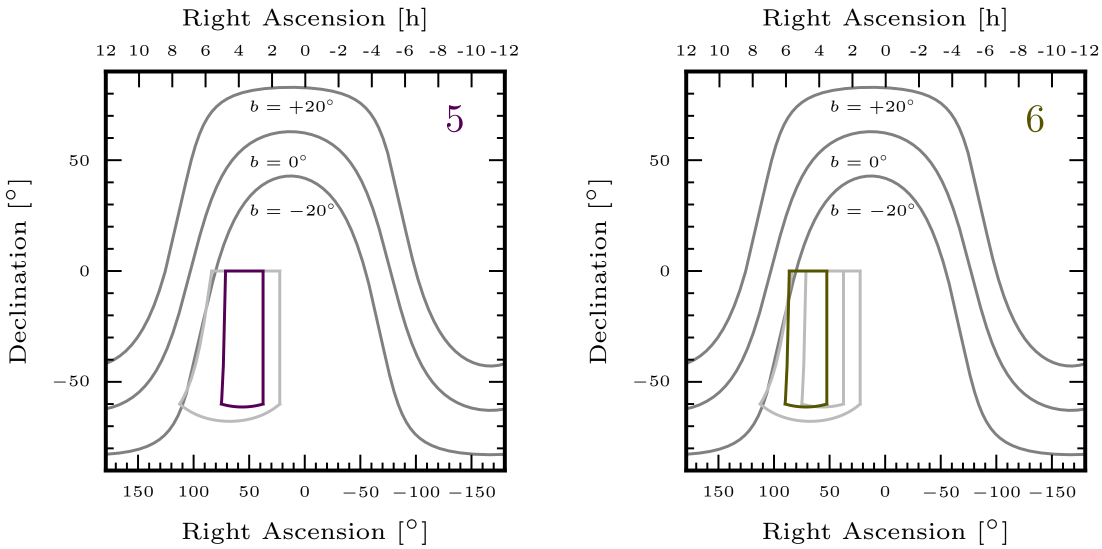


(e) Scan box 5

(f) Scan box 6

Figure 2.34: (*Continued from above*) The six scan boxes used during flight. The panel for scan box 6 has the previous *two* scan boxes in light grey, to show where the two SPIDER deep scans were located relative to the wide scan.

alone, integration time tended to be concentrated around the edges, due to the sinusoidal scan profile.

The resulting hits maps for the telescope bore-sight in flight are shown in Figures 2.35 and 2.37. These figures are intended only to provide a general sense of where SPIDER was pointed in flight. Figure 2.36 more accurately represents the combined coverage and integration time for all SPIDER focal planes. It shows the results of running the in-flight pointing timestreams through the qpoint simulator [71], which takes into account the on-sky detector beam widths and offsets within the FOV of each focal plane, and the nominal on-sky focal plane orientations at each pointing. The results of this simulation show a geometric sky fraction $f_{\text{sky}} = 10\%$, and a hits-weighted $f_{\text{sky}} = 6.5\%$.

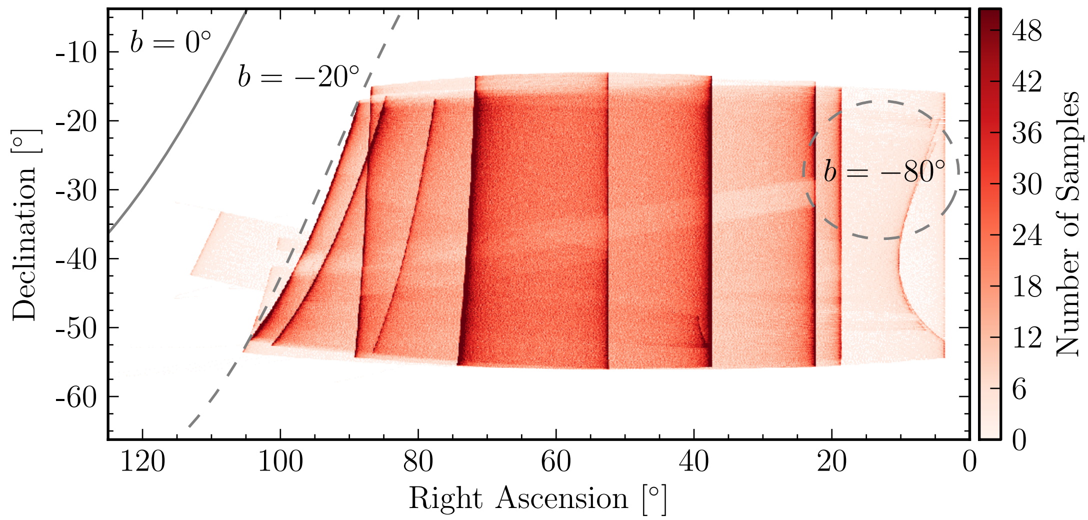


Figure 2.35: A map showing bore-sight hits, produced by binning RA and Dec from the in-flight pointing solution into 0.1° pixels. This is a close-up of the SPIDER scan region, and only data acquired when the experiment was in `spider_scan` mode are displayed.

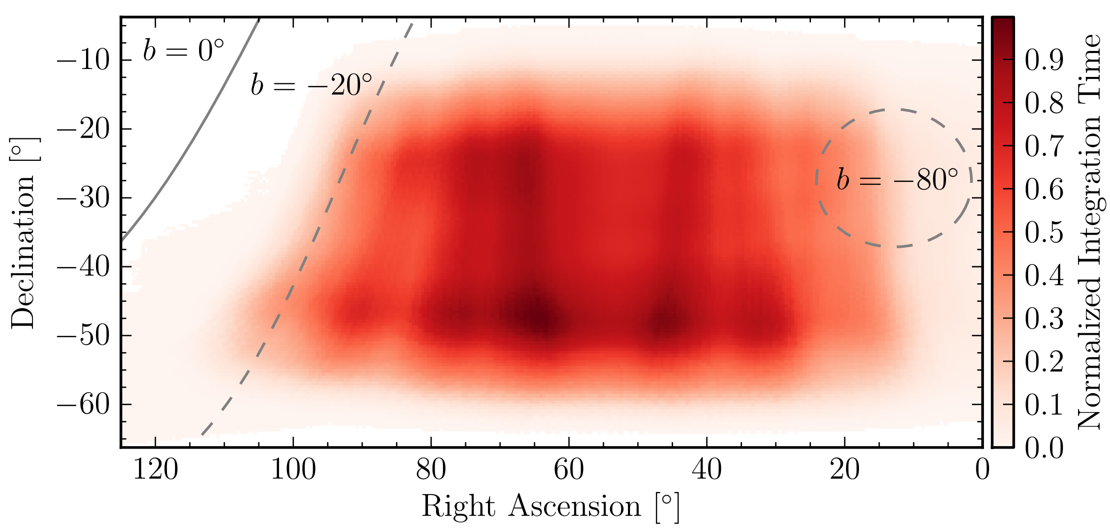


Figure 2.36: The normalized integration time resulting from running the pointing data shown in Figure 2.35 (excluding fridge cycles) through the qpoint simulator, using nominal values for the detector beam widths, offsets, and focal plane orientations. This figure is adapted from [71].

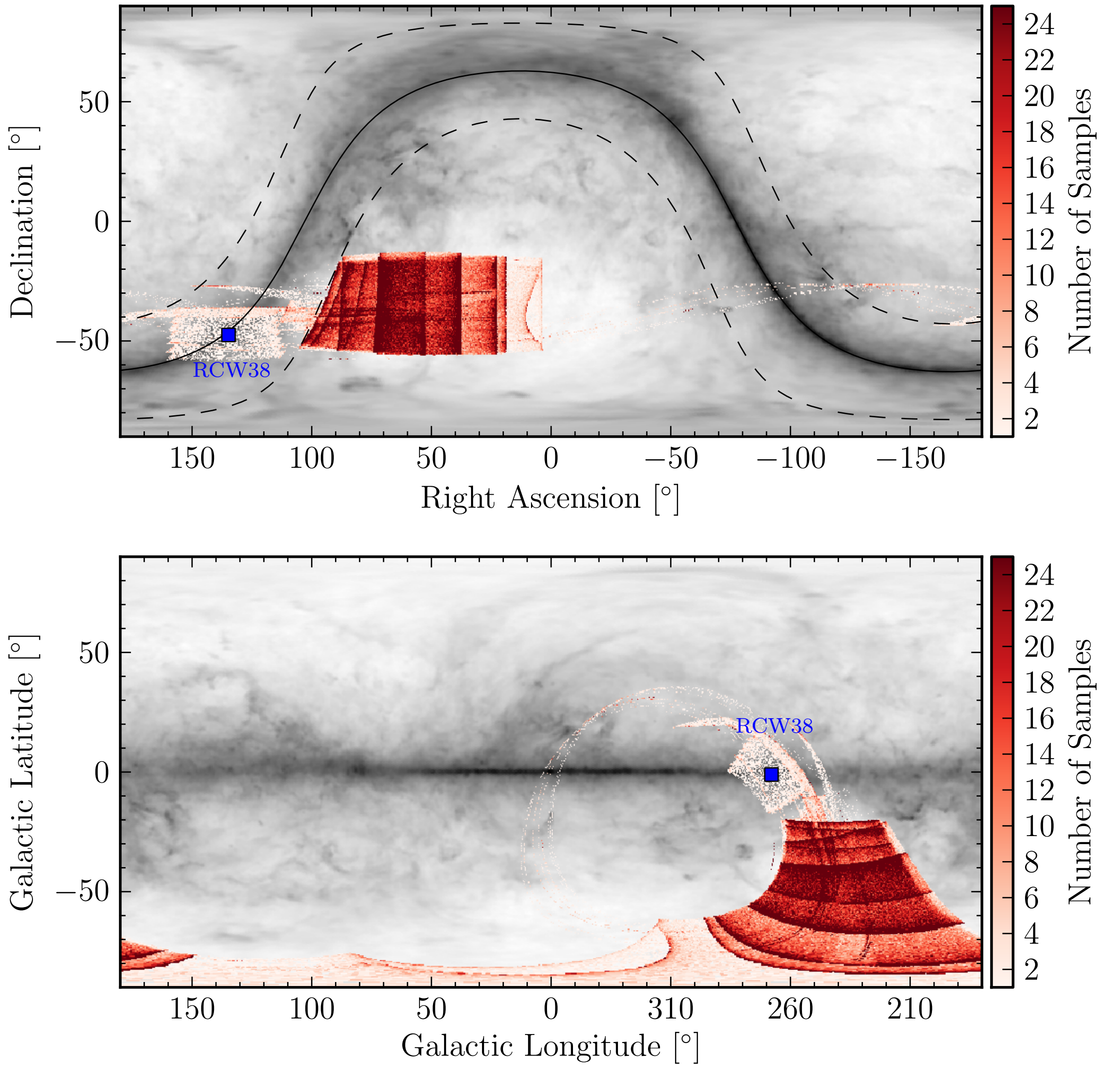


Figure 2.37: Maps showing bore-sight hits over the whole sky, produced by binning RA and Dec from the in-flight pointing solution into 0.1° pixels (subsequently re-binned into pixels of HEALPix $N_{\mathrm{side}} = 256$). These hits maps are overlaid on *Planck* 353 GHz all sky dust emission maps. All pointing data acquired after the last time a flight computer was rebooted in flight are displayed. In addition to the main CMB map, a secondary scan of RCW38, a bright calibration source in the Galactic plane, is visible. *Top* – Equatorial coordinates. *Bottom* – Galactic coordinates.

## 2.10 Pointing System Performance

This section focuses on evaluating pointing system performance based on flight data. The performance of the pointing system during laboratory testing is summarized in [49].

### 2.10.1 Motor Performance

During pointing system tests on the ice prior to launch, drops in pivot torque as a function of pivot motor winding current were observed and measured. The problem could temporarily be fixed by re-running the auto-commutation procedure on the pivot servo drive, but would reoccur during scan testing. Eventually, it was determined that the adhesive mounting the pivot stator windings to the pivot casing had failed, causing the stator windings to slip during operation, thus changing the commutation. In the process of stator slippage, the wire for one of the three motor phases was damaged, breaking at a pinch point just outside its connection to the motor windings. This caused an internal short of that motor phase to the case. It is extremely fortunate that the 12 AWG wire broke at a point *outside* the windings, where it could be accessed and soldered without causing irreparable damage to motor. After pivot re-assembly, the stator windings were secured with three set screws. The screws were installed into tapped holes that had already been placed around the outside of the casing for this purpose. The set screws were torqued down to 70 in·lbs and potted with Miller-Stephenson epoxy.

Given this situation, it seemed prudent not to drive the pivot motor as hard as originally planned. The peak current limit of the pivot motor was reduced in the servo drive firmware. The peak gondola angular acceleration was reduced from $\alpha_{\text{pk}} = 0.8$ deg·s$^{-2}$ to between 0.4 and 0.5 deg·s$^{-2}$. The decison was also made to operate the pivot in torque mode, rather than velocity mode (see Section 2.6.2.2). In torque mode, the pivot motor current could be controlled directly. Furthermore, torque mode had fewer and more intuitive gains. With the pivot operating in torque mode, the system performed surprisingly well on the ice, with no signs of the driven pendulations that had been observed during past testing (with a different and incomplete payload configuration).

When the motors were tuned for the first time at float, during LOS testing, only $0.3\ \mathrm{deg}\cdot\mathrm{s}^{-2}$ could be achieved. Further gain tuning to optimize the system for the flight train dynamics was not pursued at that time. The peak scan speed of ~3.6 deg/s resulting from this acceleration was deemed acceptable. For the duration of scientific observations, the payload operated in pivot torque mode at this reduced $\alpha_{\mathrm{pk}}$. Velocity mode testing was reserved for the end of flight, once cryogens had run out. During this testing, stable scanning at $\alpha_{\mathrm{pk}} = 0.5\ \mathrm{deg}\cdot\mathrm{s}^{-2}$ was achieved (Section 2.10.4), demonstrating the potential of the pivot control mode that was developed specifically for SPIDER.

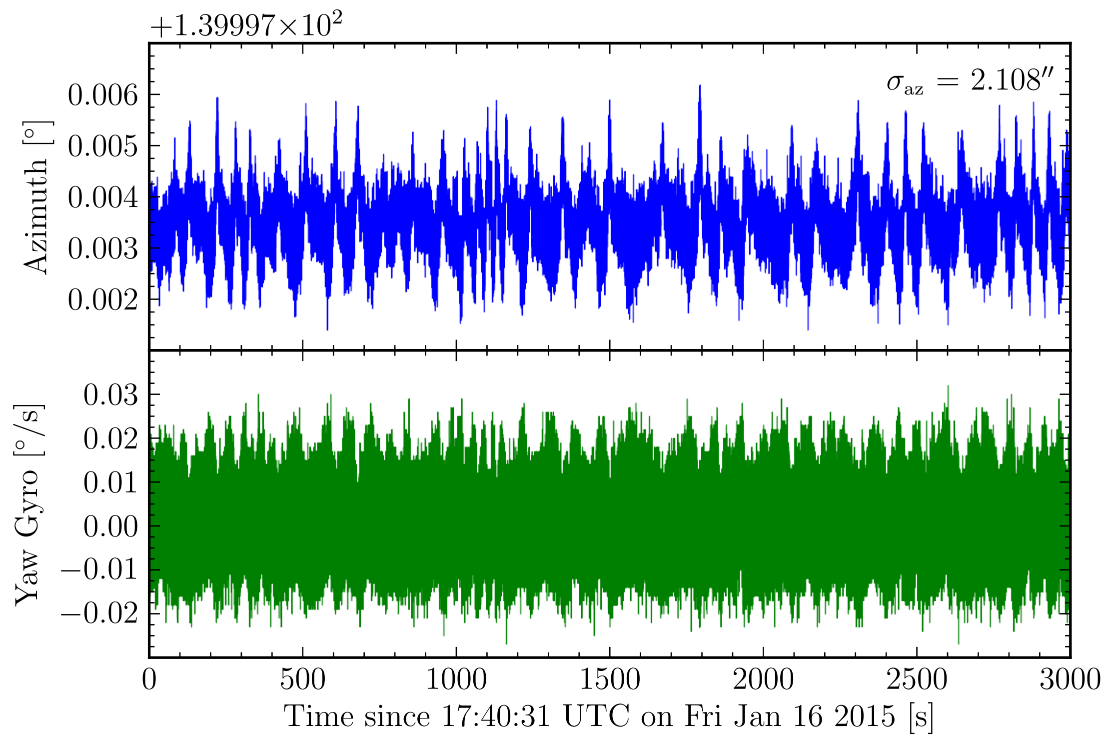


Figure 2.38: Az pointing timestreams during an Az-El Goto, with the pivot in torque mode. *Top* – Az according to the in-flight pointing solution. *Bottom* – the yaw gyro timestream.

## 2.10.2 Pointing Stability: Az-El Goto Mode

Az-El Goto mode was commanded near the end of flight, to test pointing stability. This test was first carried out with the pivot in torque mode, and later repeated in velocity mode.

### 2.10.2.1 Az-El Goto in Pivot Torque Mode

Figure 2.38 shows 50 minutes of data taken in Az-El Goto mode with the pivot in torque mode. The Az stability is $\sigma_{\rm az} = 2.108''$ during this interval, although the timestream has spikes with a nearly $7''$ amplitude.

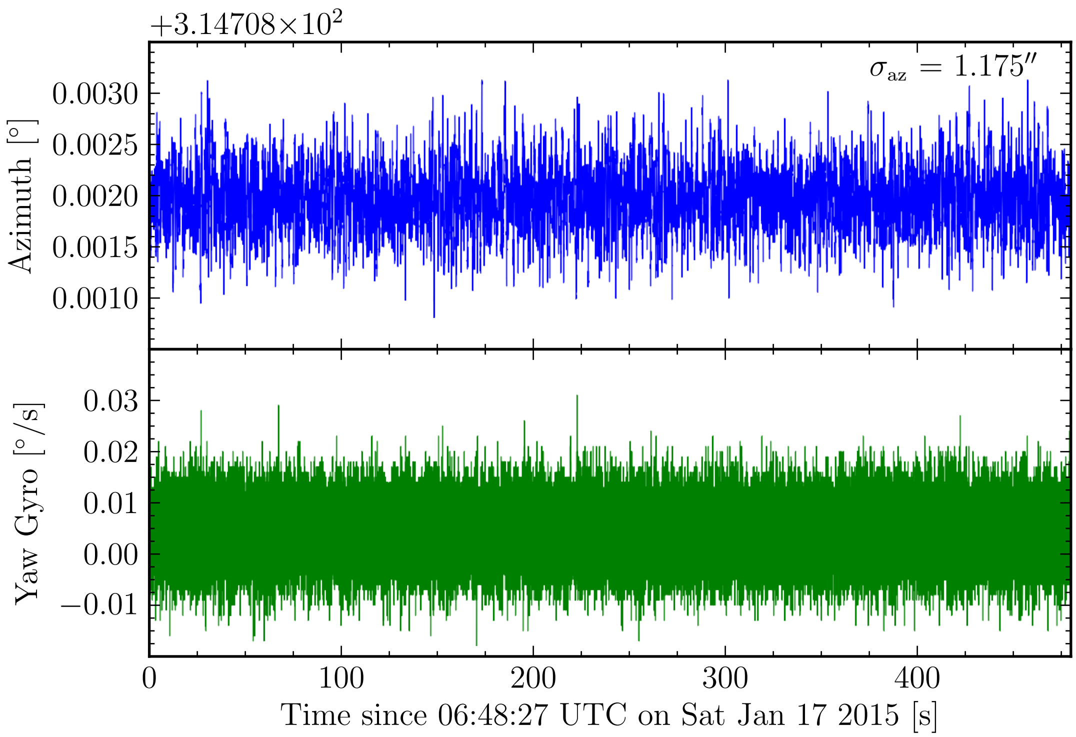


Figure 2.39: Az pointing timestreams during an Az-El Goto, with the pivot in velocity mode. *Top* – Az according to the in-flight pointing solution. *Bottom* – the yaw gyro timestream.

#### 2.10.2.2 Az-El Goto in Pivot Velocity Mode

Figure 2.39 shows 8 minutes of data taken in Az-El Goto mode with the pivot in velocity mode. The Az stability is $\sigma_{\mathrm{az}} = 1.175''$ during this interval, lower than for the torque mode test. Although this plot shows less data than the torque mode test above, the timestreams in torque mode appear to have ∼2 min. oscillations that are not present in the velocity mode timestream, suggesting that the pointing stability actually is better in velocity mode.

Figure 2.40 shows the power spectral density (PSD) of the yaw gyro timestream. The white noise level is visible, and at lower frequencies, some lines driven by the control system. Below that, there is a supression of power due to the control system servoing out gyro noise.

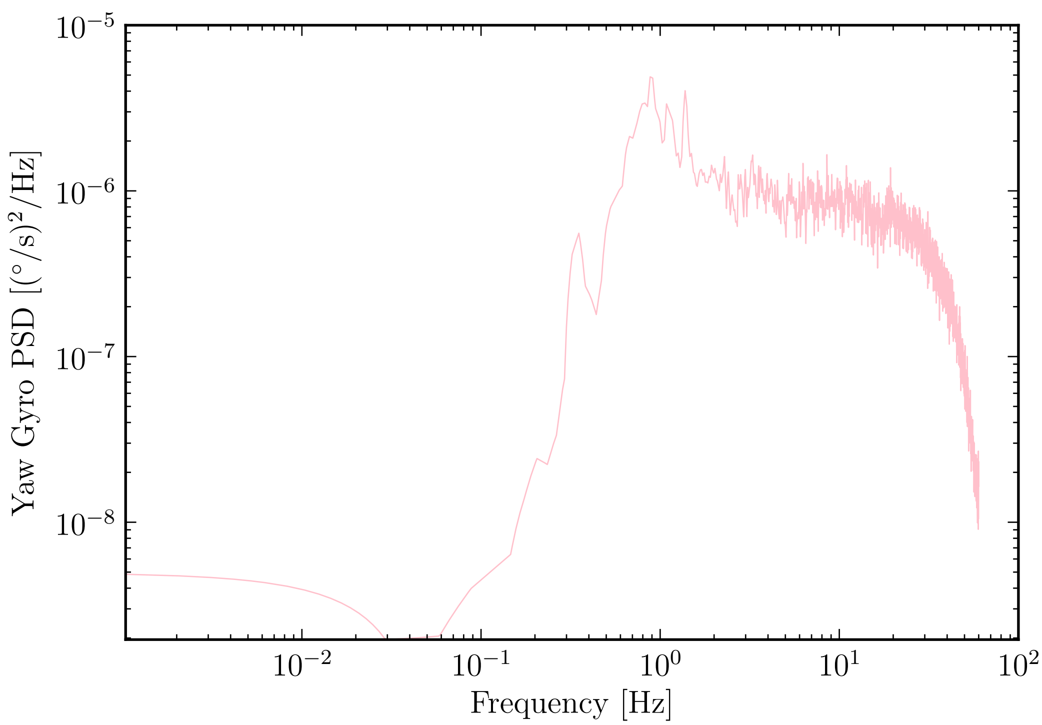


Figure 2.40: The PSD of the yaw gyro timestream from Figure 2.39

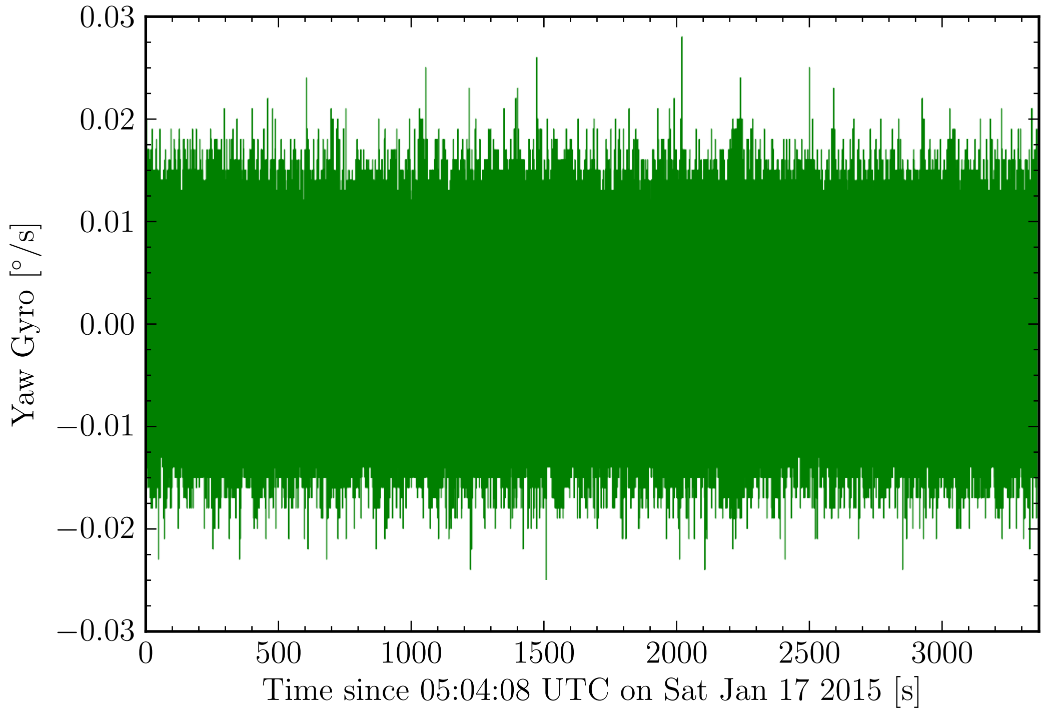


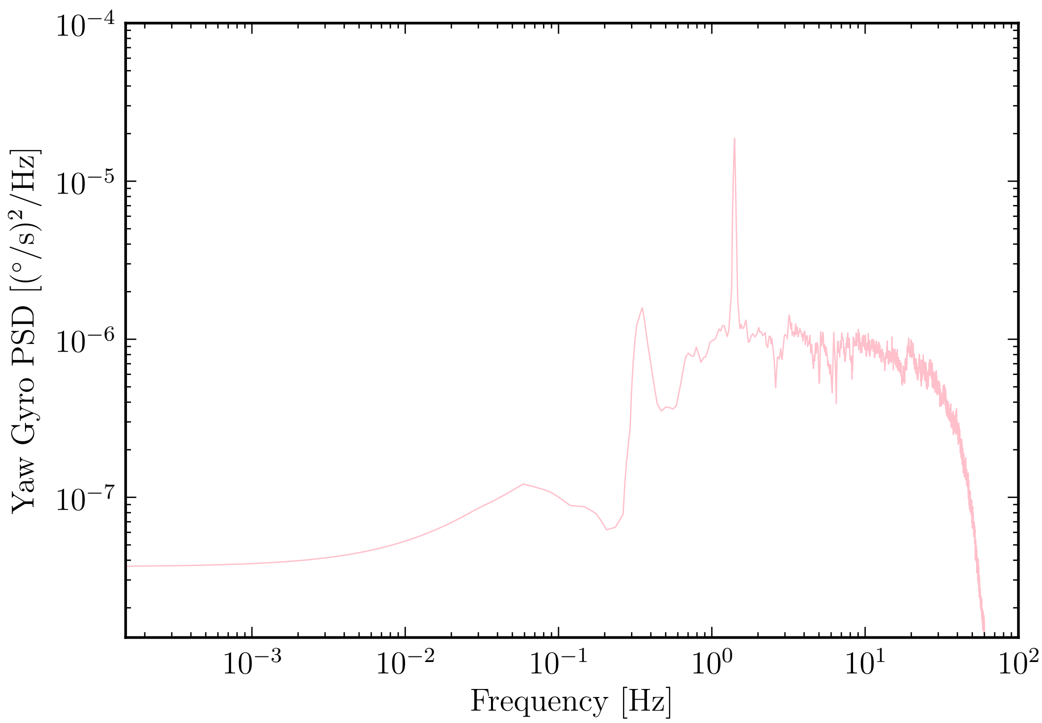


Figure 2.41: *Top* – The yaw gyro timestream during Stop mode, with the pivot in velocity mode. *Bottom* – The PSD of the above timestream.

### 2.10.3 Pointing Stability: Stop Mode

Figure 2.41 shows just under an hour of yaw gyro data taken in Stop mode with the pivot in velocity mode. The figure also shows the PSD of the yaw gyro timestream. A prominent 1.4 Hz line is visible. The nulling effect of the control system servoing out the gyro noise at lower frequencies is visible, just as it was in Az-El Goto mode.

### 2.10.4 Scan Performance

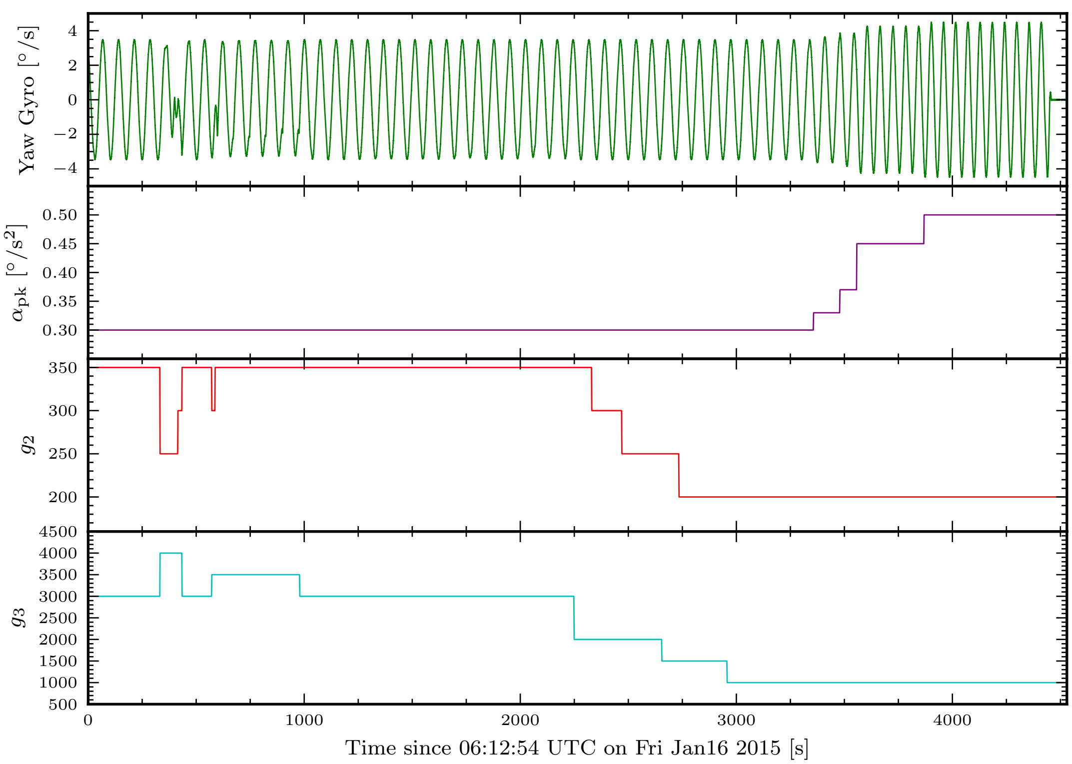


Figure 2.42: End-of-flight scan testing with the pivot in velocity mode. *Top to Bottom* – yaw gyro timestream (green), peak Az scan acceleration (purple), pivot gain $g_2$ (red), and pivot gain $g_3$ (cyan).

During LOS, it was speculated that the payload moment of inertia around the yaw axis was larger than it had been during testing in the high bay prior to launch, perhaps due to the addition of ballast and the CSBF solar panels. A reduced value of $\alpha_{\rm pk}$ =

0.3 deg·s$^{-2}$ was settled upon, with the pivot in torque mode. However, the possiblity that performance could be improved with further tuning was considered. Near the end of flight, motor gain tuning was revisited, this time with the pivot in velocity mode. The scan initially became unstable at higher peak accelerations. However, as shown in Figure 2.42, suitable adjustment of the pivot gain proportional to reaction wheel torque, and the pivot gain proportional to the Az velocity request ($g_2$ and $g_3$ from Section 2.6.2.2, Eq. 2.14), allowed $\alpha_{\mathrm{pk}}$ to be increased to 0.5 deg·s$^{-2}$. This value resulted in a peak scan speed of 4.6 deg/s.

## 2.11 The 2015 LDB Flight

SPIDER launched on January 1, 2015 at 03:59 UTC from the LDB facility. During the 2014–2015 austral summer, the facility was located 3 miles from Williams Field at coordinates[11] of (77°51′ S, 167°12.06′ E). The balloon altitude was stable over the flight, showing only diurnal variations (Figure 2.43). The flight lasted approximately 17.5 days. The payload landed on the West Antarctic Ice Sheet at (76°21.90′ S, 87°27.14′ W) at a (geodesic) distance of 2275 km from the launch site. Figure 2.44 shows the flight path around the continent. The helium main tank lasted approximately 16 days, which was better than the expected hold time.

The decision was made to terminate the flight soon after the end of science operations, in West Antarctica, over concerns that the balloon's flight path would eventually take it over ocean. The remote location of the landing site made a full recovery of the payload in the remainder of the season logistically difficult. Fortunately, the site was close to the Sky Blu field camp, operated by the British Antarctic Survey (BAS). The camp is supported by a blue ice runway and Twin Otter aircraft. On February 4, 2015, a BAS team led by Sam Burrell was able to reach the payload and recover the 5 pressure vessels, the

[11] These coordinates are as reported by the GPS in the LDB high bay

FLCs, the 6 MCCs, the 3 star cameras, the 3 GoPros, and the SIP. These items together contained all of the recorded data on the experiment. The items were shipped out of Rothera Station to North America via Punta Arenas, Chile. The full recovery of the data drives marked the complete success of the flight campaign. As of this writing, the remainder of Spider lies on the ice, awaiting recovery in the 2015–2016 season.

Figures 2.45 through 2.47 show some events from launch day. The last image was taken by one of three GoPro cameras mounted on the experiment before launch. The downward-facing limb camera's battery lasted some 30 hours after activation, capturing footage of launch, ascent, and scan testing at float altitude.

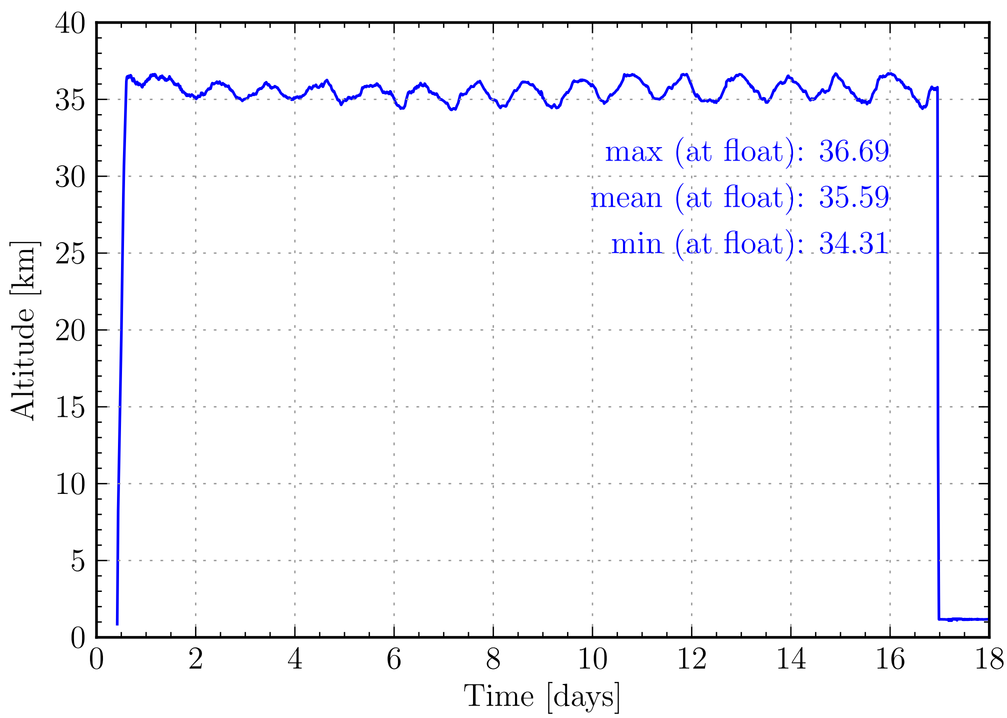


Figure 2.43: The altitude of Spider vs. time during the 2015 LDB flight, according to data from the SIP

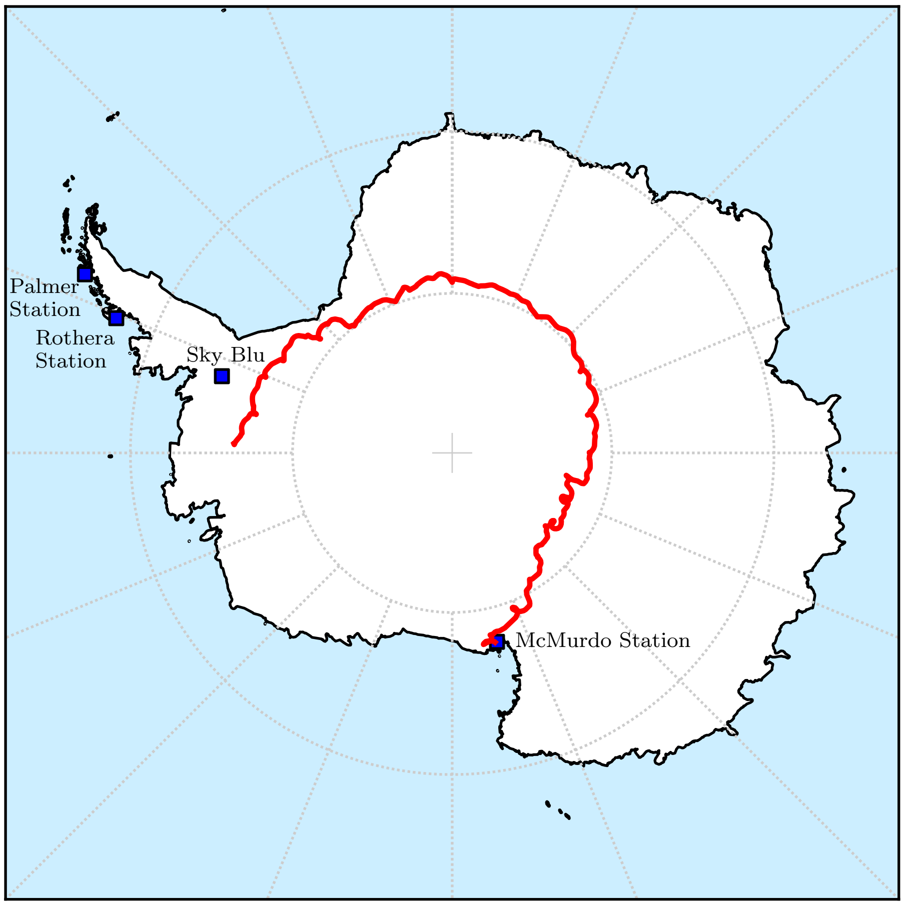


Figure 2.44: The path of the SPIDER 2015 LDB flight. The locations of Rothera Station and the Sky Blu field camp, operated by BAS, are shown.

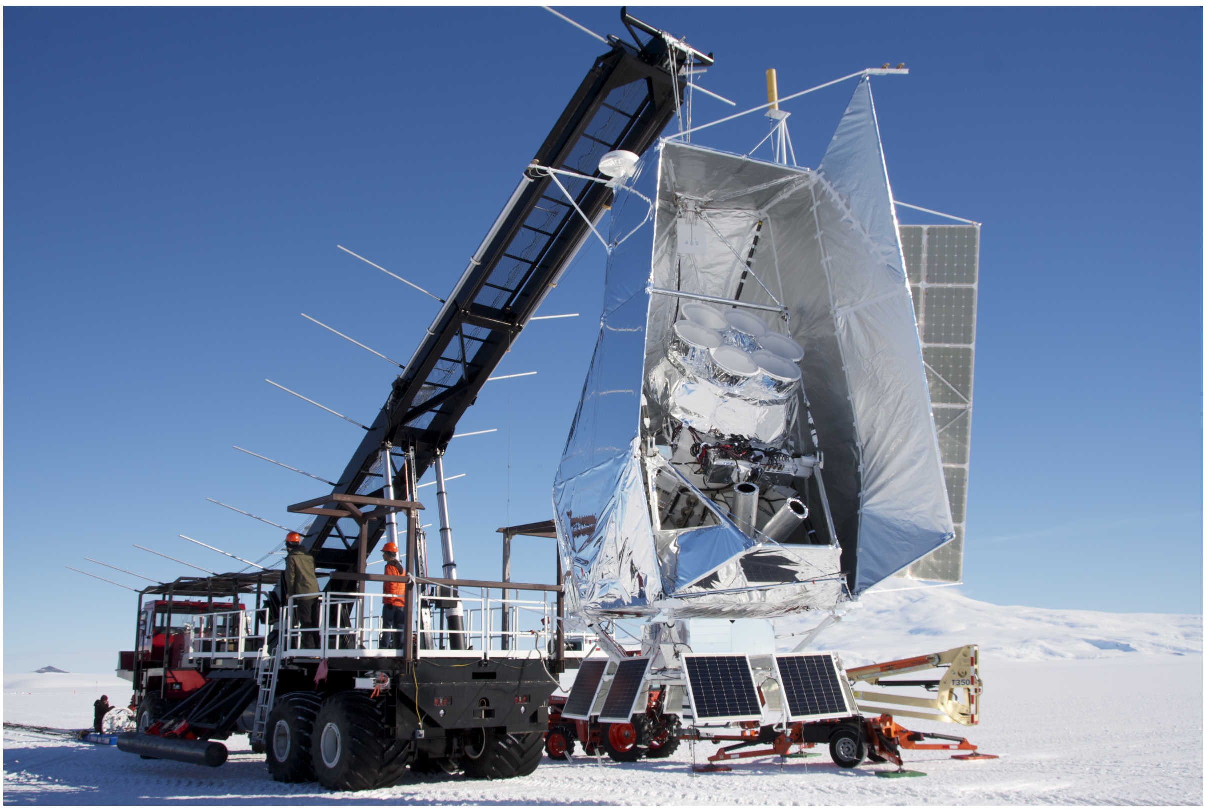

Figure 2.45: The Spider experiment hangs from the launch vehicle, known as *The Boss*, on launch day

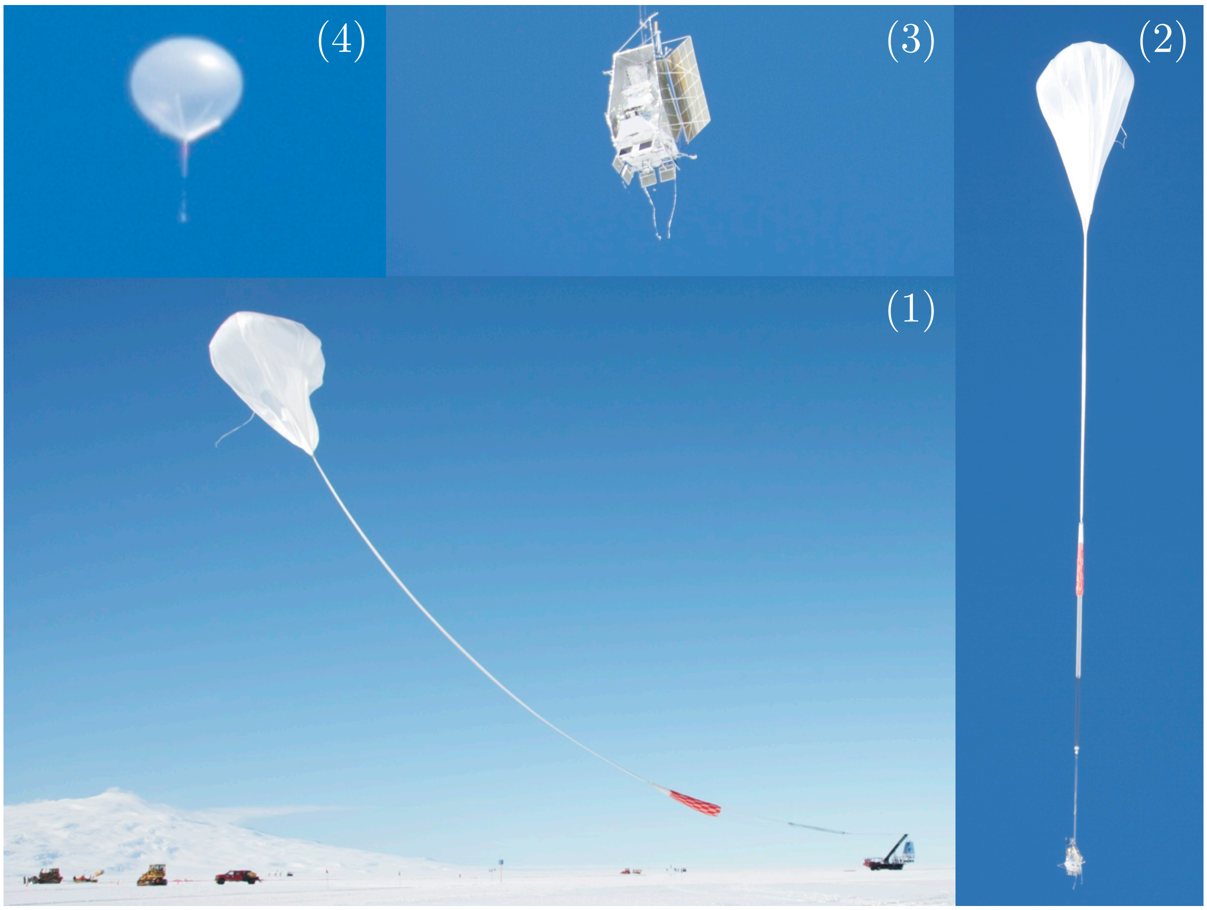


Figure 2.46: Photographs from the SPIDER launch. 1) The balloon and flightline shortly after release. 2) The balloon ascending, just after launch. 3) SPIDER, now aloft. 4) The fully-inflated balloon at 35 km altitude, as seen from the ground.


Figure 2.47: The Earth limb as imaged by a GoPro camera mounted on Spider. The image has been colour-corrected and corrected for lens aberration. Geographic features on or near Ross Island are labelled. The wisp is due to a scratch on the GoPro casing. The object in the foreground is the magnetometer boom and its support wire.

# Chapter 3

# The BLASTPol Experiment

A portion of the work conducted for this thesis was devoted to a second experiment: the Balloon-borne Large-Aperture Submillimetre Telescope for Polarimetry, or BLASTPol. BLASTPol is a modification of the original BLAST: a balloon-borne 1.8 m telescope for submillimetre-wavelength observations [72]. BLAST had three focal plane arrays of bolometric detectors with bands centred at 250 µm, 350 µm, and 500 µm. These arrays were prototypes of the ones used for the SPIRE instrument on the *Herschel* Space Observatory [73]. BLAST fulfilled at least three purposes:

1. The observation of thermal dust emission in Galactic star-forming regions [74].

2. The observation of resolved nearby galaxies in the submillimetre [75].

3. The detection of submillimetre emission from numerous extragalactic sources—high redshift galaxies—revealing information about the Cosmic Infrared Background (CIB) [76, 77], the cosmic history of star formation [78], and galaxy clustering [79].

BLASTPol was BLAST converted into a polarimeter by the placement of a polarizing grid in front of the feed horn array of each focal plane. In addition, an achromatic half-wave plate (HWP) was installed to modulate the polarization of the sky signal, allowing each detector to make an independent measurement of the $Q$ and $U$ Stokes parameters for linear polarization (Section 1.6). BLASTPol's main purpose was to measure *polarized* thermal emission from dust in Galactic molecular clouds, in order to learn more about the role played by magnetic fields in the early stages of star formation (Section 3.1.2).

A BLAST test flight was launched from Fort Sumner, New Mexico in 2003 (hereafter referred to as *BLAST03*). BLAST also had two science flights: the first was launched from Kiruna, Sweden in 2005 (hereafter *BLAST05*). The second was an LDB flight launched from McMurdo Station, Antarctica in 2006 (hereafter *BLAST06*). In *BLAST06*, after a successful flight and landing, a failure of the parachute separation system caused the payload to be dragged by winds for approximately 200 km across the Antarctic ice, until it eventually came to rest in a crevasse field [80]. Fortunately, the data hard drives were recovered, as were the optics, receiver, and detectors, enabling the experiment to be rebuilt as BLASTPol. BLASTPol had two LDB science flights launched from McMurdo in 2010 and 2012 (hereafter *BLASTPol10* and *BLASTPol12* respectively). This chapter gives an overview of the BLASTPol experiment, in order to set the context for the data analysis work presented in Chapters 4 and 5. More details of the experiment and its flights are in Refs. [81–83], and in an upcoming paper [84].

## 3.1 Scientific Objectives

### 3.1.1 Alignment of Dust Grains in Molecular Clouds

For some time, it has been known that interstellar dust grains can be polarized, achieving partial alignment with the direction of the local magnetic field. This effect leads to a linear polarization of light that is partially absorbed by the dust grains (Figure 3.1a), as has been observed in optical polarimetry of background starlight [85, 86]. This technique can be used to infer the projected plane-of-sky component of the magnetic field aligning the dust. However, the technique can only be used in areas of lower density, such as in the diffuse ISM or on the periphery of molecular clouds. In denser regions, dust extinction becomes too large to see background sources in the visible and near-IR. A promising method for probing denser regions is submillimetre polarimetry [87]. Submillimetre wavelengths are sensitive to the thermal emission from the 10–20 K dust in molecular clouds, where star formation takes place. Due to the alignment of the dust grains, this emission is also linearly polarized (Figure 3.1b), with the degree of polarization dependent on the fractional degree of alignment.

The exact mechanism for dust grain alignment is still an area of active research. One proposed mechanism is known as the Radiative Alignment Torque (RAT) model [88]. In this theory, due to a net helicity, dust grains can be spun up by photons in the interstellar radiation field. As a result, they acquire a net magnetic dipole moment. The directions of the dipole moment, of the rotation axis, and of the axis of maximal moment of inertia, all become coincident. This axis precesses around the magnetic field vector. The RAT model shows that the angle of precession eventually becomes small, meaning that the rotational axes of the grains become aligned with the magnetic field. It also makes the prediction that denser, more heavily-shielded regions within clouds, or areas with a weaker ambient radiation field, should have dust with a smaller degree of fractional alignment. Observations to map the strength and direction of polarized submillimetre emission in molecular clouds can therefore probe dust physics as a function of environment, and can test the RAT and other grain alignment models.

### 3.1.2 The Early Stages of Star Formation

Another application of submillimetre polarimetry is understanding the physical processes relevant during the early stages of star formation. It is generally understood that some physical processes must regulate the gravitational collapse of cores and other dense substructures within molecular clouds, since these structures are observed to have lifetimes greater than a free-fall time [74]. It also now generally-accepted that both magnetic fields and supersonic turbulence play some role in this regulation, but the relative importance

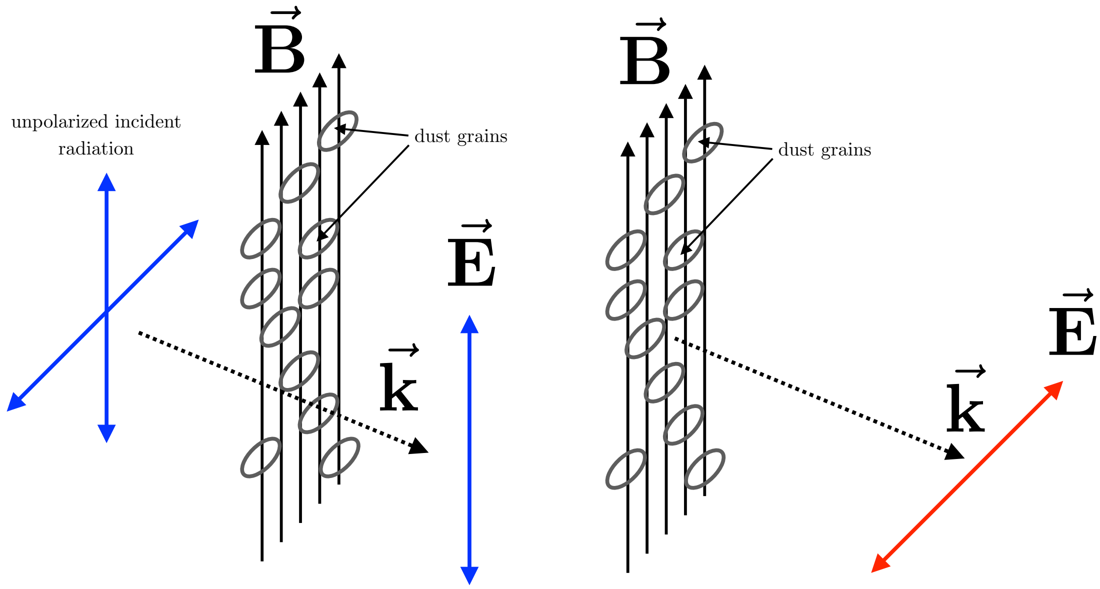


(a) Polarization in absorption

(b) Polarization in emission

Figure 3.1: Aspherical spinning dust grains (grey) preferentially align with their rotational (i.e. short) axes parallel to the direction of the local magnetic field (**B**, in black). *Left* – unpolarized visible or near-IR light (in blue) from a background source, such as a star, is incident on the dust grains. The light polarization component parallel to the long axes of the grains is preferentially absorbed. Light therefore emerges linearly polarized in the direction perpendicular to the grains' long axes, and parallel to the **B**-field. *Right* – dust grains emit thermal radiation preferentially in the direction of their long axes. Emission in the far-IR and submillimetre (red) from aligned dust grains is therefore linearly polarized in the direction parallel to the grains' long axes, and perpendicular to the **B**-field.

of each is not fully-understood [89]. Theoretical models with strong magnetic fields [90] predict smooth, ordered field lines on large scales, while models in which turbulence dominates [91] predict more chaotic, disordered field lines. In addition, magnetic-dominated models predict a different relationship between the magnetic field direction and the morphology of cloud substructures compared to turbulence-dominated models. For example, strong field models predict that dense molecular cloud cores should be oblate, with magnetic field lines aligned with the minor axes of the cores. Collapse occurs slowly through ambipolar diffusion (the breaking of flux-freezing) in which neutral material is able to diffuse slowly perpendicular to the field lines, which are pinched into a mild hourglass shape [89]. In contrast, in turbulent models with weaker magnetic fields, star formation is delayed primarily by turbulence dissipating dense substructures before they can collapse. However, this turbulence eventually decays, allowing collapse to take hold in cores that are already self-gravitating. In this scenario, the field in a core becomes increasingly ordered with time, and takes on a pronounced hourglass shape, since the weaker magnetic fields cannot slow collapse perpendicular to the field direction as effectively as in the strong field case [89].

Until the advent of submillimetre polarimetry to map out the plane-of-sky component of the **B**-field directions in star-forming regions, there was a dearth of observational data to constrain these models. BLASTPol was conceived as an experiment to probe the importance of magnetic fields in star formation. It occupies a niche between experiments such as *Planck*, which has provided a coarse resolution ($\sim 5'$ FWHM) submillimetre survey of the entire sky, and ALMA, which can provide sub-arcsecond polarimetry that can resolve magnetic fields within cores and disks, but is not sensitive to fields within the surrounding cloud. With its nominally sub-arcminute resolution, BLASTPol was designed to be able to observe cloud fields down to the scale of cores and other substructures, while still having the sensitivity and mapping speed (afforded by a balloon-borne observing platform) to trace those fields across entire clouds. Using this capability, BLASTPol was intended to answer the following questions, which are elaborated upon in [81]:

1. Is molecular core morphology determined by large-scale magnetic fields?

2. Do filamentary structures within clouds have magnetic origins?

3. How strong are **B**-fields in clouds, and how does the strength vary among clouds?

## 3.2 Optical Design

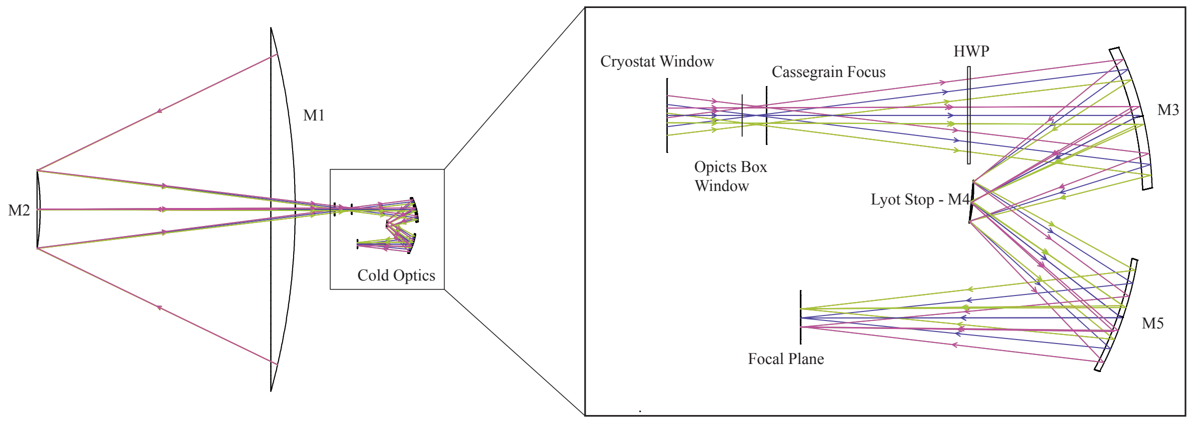


Figure 3.2: A ray-tracing diagram of the BLASTPol optics.

Figure 3.2 shows a schematic of the BLASTPol optical system, which is an on-axis Cassegrain design. The primary mirror (M1) is a 1.8 m aluminum hyperboloid[1]. The aluminum hyperbolic secondary (M2) is 40 cm in diameter. From the secondary, light passes through the window of a liquid nitrogen ($LN_2$) and liquid helium (LHe) cryostat into the cooled optics box, where it is re-imaged onto the detector arrays by a series of ∼1.5 K spherical reflective optical elements (M3, M4, and M5) in an Offner relay configuration. Element M4 acts as the Lyot stop, determining the illumination of the primary for each feed. A series of two dichroic beam splitters, which are not shown, reflect wavelengths shorter than a cutoff, and transmit longer wavelengths, thus directing light to each of the 250 µm, 350 µm, and 500 µm arrays. The bands are further defined at the long wavelength end by the cutoff of the waveguides coupling each feed horn to the detector arrays. In general, the band edges are also defined by additional metal-mesh filters [54]. Each band has an approximate width of 30%, and their spectral responses were measured before flight using a Fourier Transform Spectrometer (FTS).

A cryogenic (4 K) achromatic HWP [92] is placed in the optics box, 19 cm beyond

[1] Although M1 is referred to as a paraboloid in previous literature, its conic constant is actually -1.029, making it slightly hyperbolic, and making the telescope technically a Ritchey-Chrétien telescope.

the Cassegrain focus. This placement reduces the effect of localized defects in the HWP structure. The plate is 10 cm in diameter and consists of five layers of sapphire each 500 μm thick, glued together by 6 μm layers of polyethylene. In flight the HWP was rotated in discrete steps of 22.5°, once for each complete up and down scan of a target (Section 3.4.2.2). The HWP rotator consists of steel ball bearings housed in a stainless steel structure driven by a gear train and a G-10 shaft that connects to a stepper motor external to the cryostat. A potentiometer is used for position sensing at 4 K.

## 3.3 Cryostat and Receiver

### 3.3.1 Cryogenic System

To operate, the detectors must be cooled to ∼300 mK. This is achieved in several stages. The upper half of the cryostat houses cylindrical $LN_2$ and LHe tanks with capacities of 43 L and 32 L respectively. They are maintained at pressures slightly greater than 1 atmosphere during flight. The tanks provide 77 K and 4 K cooling stages in the form of shields that extend around the optics box in the lower half of the cryostat. In between these shields is a vapour-cooled shield (VCS) that is cooled by the passage of boil-off from the liquid cryogens through a heat exchanger. The VCS reaches ∼35 K, and reduces the thermal loading on the LHe stage.

A small $^4$He reservoir is connected to the main helium tank through a capillary. In flight, this "pumped pot" is vented to the external atmosphere (∼3 mbar). It is connected to a vacuum pump during operations in the laboratory. Liquid is forced across the capillary by the pressure difference between the main tank and the pot. This pumped helium bath provides a ∼1.5 K stage with 20 mW of cooling power. This stage cools the optics box, and also maintains the temperature of the condensation point of a closed-cycle $^3$He refrigerator. This fridge cools the detectors from 1.5 K down to approximately 300 mK. It contains a charcoal pump and heater, and operates in essentially the same manner as the SPIDER fridges described in Section 2.2.3. The fridge can provide 30 µW of cooling power for four days, and takes two hours to be cycled. The overall hold time of the cryostat is approximately 13 days.

### 3.3.2 Detectors

The BLASTPol focal plane arrays consist of 139, 88, and 43 bolometers at 250 µm, 350 µm, and 500 µm respectively. Table 3.1 summarizes the yield of working bolometers based on analysis of *BLASTPol12* data. Table 3.2 lists the bolometers that were identified to be non-functional, due to elevated noise or other criteria. This result is also shown visually[2] in Figures 4.6 through 4.8. The devices are "spider-web" bolometers made up of a silicon nitride micromesh absorber to which a Neutron Transmutation Doped Germanium (NTD-Ge) thermistor has been glued [93]. Radiation is coupled to the bolometers by an array of smooth-walled conical feed horns. The feeds have $2F\lambda$ spacing where $F$ is the focal ratio of the entire optical system.

In BLASTPol, the photo-lithographed polarizing grids mounted in front of each feed horn array were patterned to alternate the polarization angle by 90° from feed-to-feed along the scan direction (Figure 3.3). Therefore the time to measure one Stokes parameter was given by the bolometer spacing divided by the typical scan speed of 0.1 deg/s, and is much shorter than the timescale characteristic of the $1/f$ knee in the bolometer noise spectra at $\sim$0.05 Hz.

### 3.3.3 Readout

In the detector bias circuit, two load resistors, having resistance $R_L \gg R_{\mathrm{bolo}}$, are placed in series with the bolometer (see Figure 2 of [62]). An AC bias voltage is applied to

[2]However, some of the bolometers that were identified to be bad still appear to have viable beam maps in these figures. It is likely that they were bad intermittently during the flight, and not during the period of time in which the beam calibration source was observed.

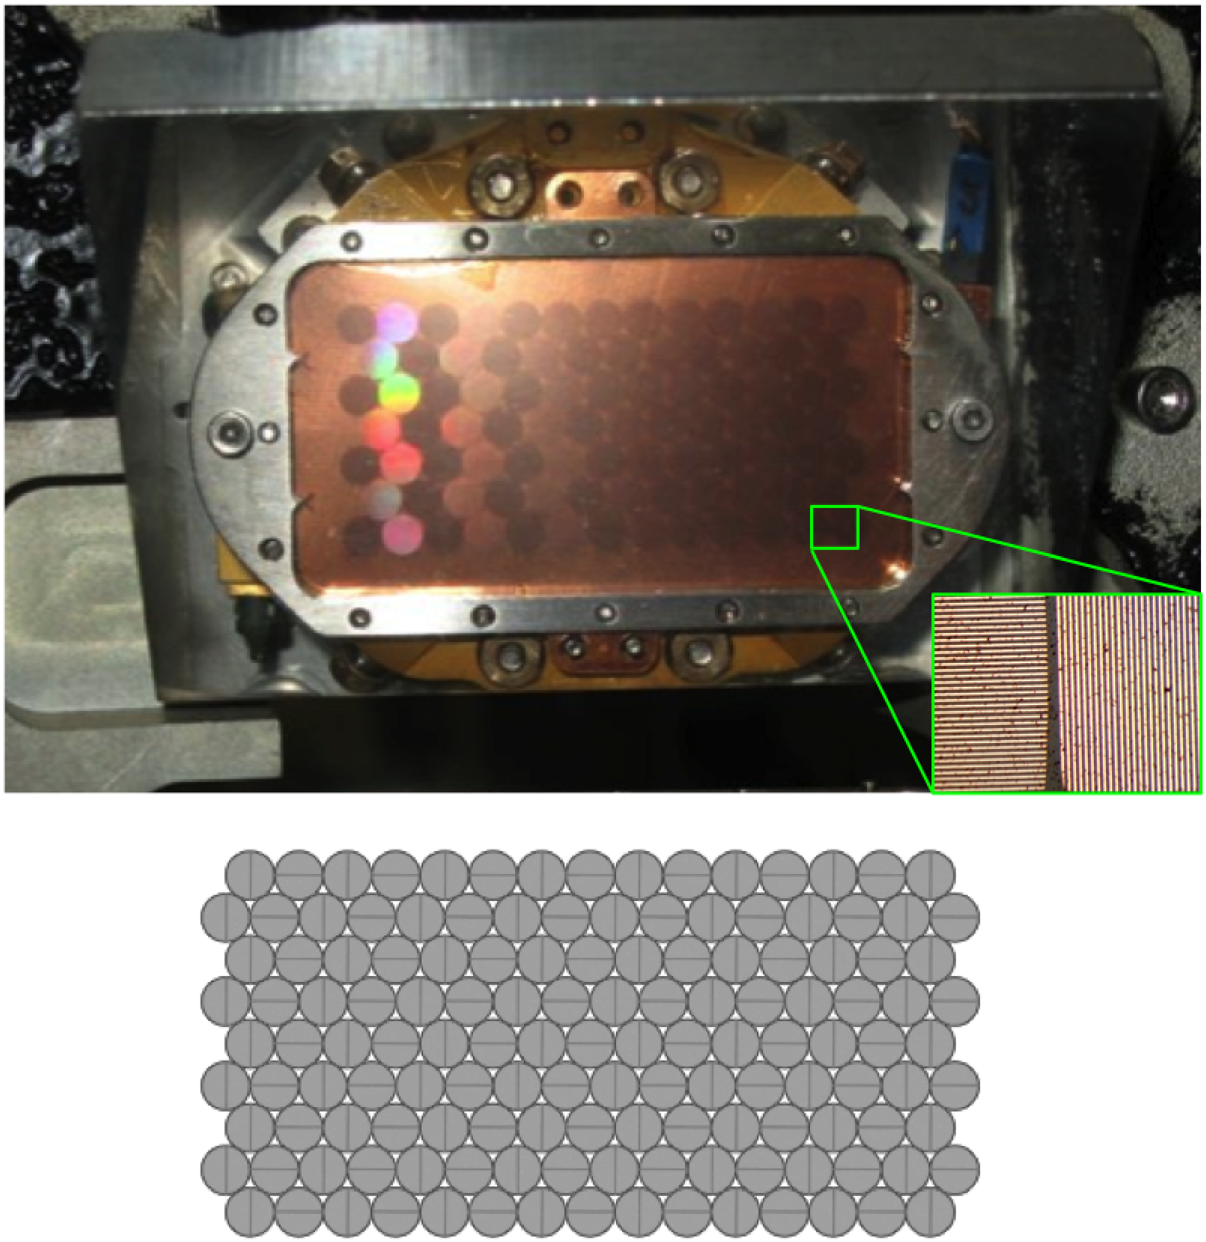

Figure 3.3: *Top* – a photograph of the BLASTPol 350 µm array showing the photo-lithographed polarizing grid placed in front of it. *Top Inset* – a close-up of the grid between two adjacent pixels, showing the 90° change in polarization angle. *Bottom* – A diagram of the 250 µm array, showing the alternating pattern of polarization sensitivity for each bolometer.

Table 3.1: The yield of working bolometers in the BLASTPol detector arrays in 2012

| Array | Total | Working | Yield |
|---|---|---|---|
| 250 µm | 139 | 129 | 92.8% |
| 350 µm | 88 | 77 | 87.5% |
| 500 µm | 43 | 40 | 93.0% |
| **Total** | **270** | **246** | **91.1%** |

Table 3.2: A list of the bolometers determined to be non-functional in *BLASTPol12*. Bolometers are listed by name. For instance B2A02H indicates the 250 µm focal plane, row A, column 02, and horizontal (vs. vertical) polarization orientation.

| 250 µm | 350 µm | 500 µm |
|---|---|---|
| B2A02H | B3A01V | B5D02H |
| B2A15V | B3B05H | B5E01V |
| B2B01V | B3B07H | B5E08H |
| B2B04H | B3C01V | |
| B2B13V | B3C11V | |
| B2C12H | B3D03H | |
| B2D03V | B3D10V | |
| B2F02H | B3E11V | |
| B2G02H | B3G01V | |
| B2H12H | B3G02H | |
| | B3G04H | |

this series combination, which results in an approximately constant bias current being applied to the bolometer. The voltage across the bolometer is therefore a measure of its resistance, which has a known temperature dependence. This voltage is read out differentially by a pair of JFET amplifiers. The amplifiers are located in the JFET cavity of the cryostat, in order to be as close as possible to the bolometers, reducing capacitive pickup on the readout signal lines. The bolometer signal then passes to the pre-amplifier crate, where it is further amplified and bandpass filtered. From here the signal propagates to the Data Acquisition System (DAS), a rack containing BLASTbus electronics boards with 24-bit $\Sigma\Delta$ ADCs to sample and digitize the data. These boards also generate the $\sim$200 Hz bias signal and implement a digital lock-in amplifier. The DSPs on these boards carry out further digital filtering. More details of the BLASTbus electronics, including their use in the biasing and readout of NTD-Ge bolometers, can be found in [62] and [63].

## 3.4 Gondola and Pointing System

### 3.4.1 Structural Overview

The BLASTPol gondola was a rebuilt version of the aluminum gondola developed for BLAST. The main features of this design are shown in Figure 3.4. The gondola consists of an outer frame, capable of pointing in azimuth, and an inner frame, which can point in elevation relative to the outer frame. The outer frame has a main rectangular frame whose top surface is an aluminum honeycomb deck to which the flight electronics are mounted. These include the Attitude Control System (ACS) and the two redundant onboard flight computers. The rectangular frame also has pyramidal structures at either end that provide mount points for the inner frame. The reaction wheel is mounted below the main deck, and a cage for housing the SIP extends below it. The inner frame houses the telescope primary and secondary mirrors, the cryostat, the receiver electronics, and

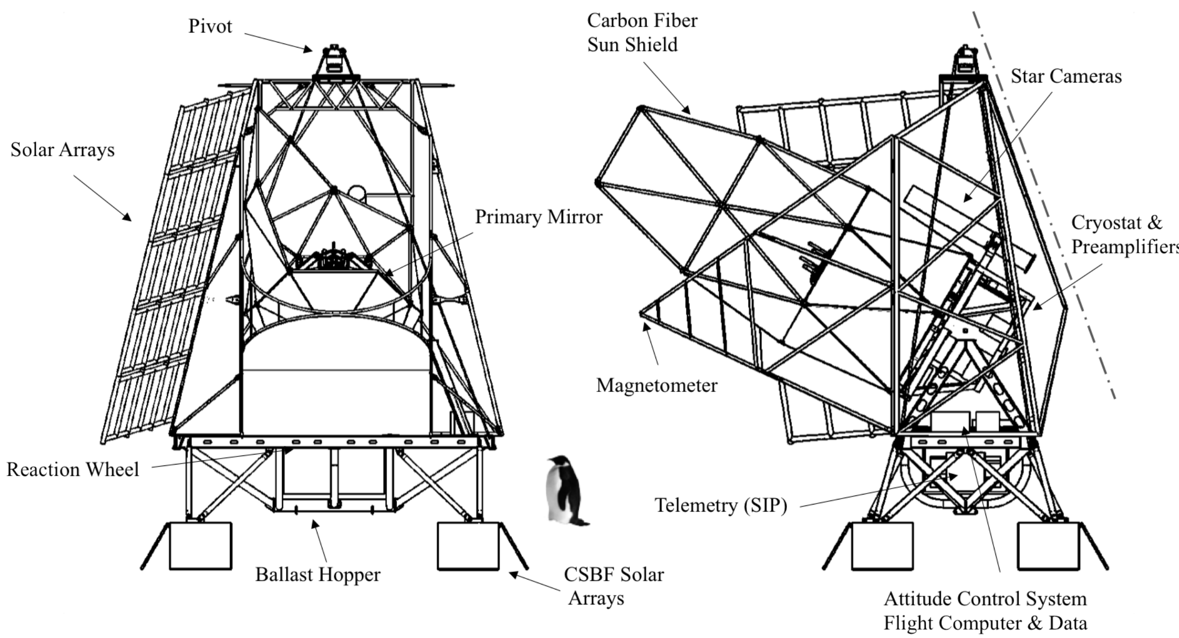


Figure 3.4: The main components of the BLASTPol payload. A 1 m Emperor penguin is shown for scale.

two star cameras. BLASTPol saw the addition of a 4 m baffle to the inner frame, composed of a carbon fibre frame covered in aluminized Mylar®. This baffle allowed the gondola to point to within 45° of the sun in azimuth on the starboard side, increasing the number of available Galactic targets for observation. Surrounding the gondola outer and inner frames is another aluminum frame for the sunshield. This is also covered in aluminized Mylar®. The sunshield protects the telescope from direct solar radiation, and also significantly reduces temperature variations in the optics during flight. The entire payload is suspended from the pivot by four steel cables that attach to the corners of the outer frame.

## 3.4.2 Pointing System

### 3.4.2.1 Pointing Control

Many elements of the Spider pointing control system are of BLASTPol heritage. Therefore, the BLASTPol system operated in a manner similar to that described in Sections 2.6 through 2.10. Pointing in azimuth was achieved with a reaction wheel (RW) and a motorized pivot. The RW motor torque control loop is that given by Figure 2.21. The RW brushless DC motor was a different model than Spider's, with a lower maximum torque rating. A servo drive from Copley Controls was used to control RW motor current, rather than the Advanced Motion Controls (AMC) unit. Unlike Spider, the RW speed setpoint ($\omega_{SP}$, Section 2.6.2.2) is non-zero, in order to avoid static friction. BLASTPol scans are significantly narrower and require less angular acceleration than Spider scans, making a large swing in RW velocity around zero unnecessary. The physical design of the BLASTPol RW also differed from Spider's, consisting of a 1.52 m diameter aluminum honeycomb disk of 7.6 cm thickness. Forty-eight 0.9 kg brass disks were embedded around the circumference of the wheel. The BLASTPol pivot operated in torque mode according to Equation 2.11, plus an additional control term to overcome static friction (see [68] Eqs. 2.29 and 2.30). The pivot motor was identical to the Spider pivot motor, and the DPRALTR-020B080 model of AMC servo drive (from the same series as Spider's servo drive), was used to drive the BLASTPol pivot (see Section 2.6.2.2 and Table 2.3).

The BLASTPol elevation drive design differed significantly from Spider's as well. The two inner-to-outer frame attachment points form a horizontal elevation axis. On one side, the attachment point consists of a free bearing, while on the other side, a direct-drive brushless DC motor is mounted, enabling the inner frame to be rotated around this axis. The motor torque was controlled by an algorithm similar to that of the BLASTPol RW motor torque control loop. The system could servo continuously in elevation, thus removing the effect of pendulations on the inner frame. More details of the BLASTPol pointing control system can be found in [68].

#### 3.4.2.2 Scan Strategy

The most common[3] BLASTPol scan strategy was to scan back and forth in Az while drifting continuously in El. The El drift speed was set such that after each Az scan, the telescope had drifted in El by a specified increment $\Delta\theta$, which could be commanded in flight. Az scans were carried out at a constant speed, with constant acceleration at the turn-arounds. The typical Az scan speed was $\lesssim 0.1$ deg/s, and the typical map width was 1 degree. After reaching a specified El limit, the El drift direction would reverse. The telescope would raster up and down over a target in this manner in approximately 15 minutes, after which the HWP was rotated by 22.5°. The up-going and down-going maps were repeated for each of the HWP positions, which were nominally 0°, 22.5°, 45°, and 67.5°. Full maps of the source at every HWP position could therefore be obtained within an hour.

A few different scan modes were implemented that differed in the way the coverage area was mapped out on the sky. In a *box* scan, the telescope scanned over a rectangular region in Az and El that was centred on particular coordinates in RA and Dec. Therefore, the orientation of the box on the sky changed with time. In a *quad* scan, the telescope scanned within a quadrilateral on the sky defined by its corner points in RA and Dec. In a *cap* scan, the region mapped on the sky was circular, rather than quadrangular. None of these scan modes for mapping science targets was carried over to SPIDER. An entirely new scan strategy was devised to meet the particular needs of that experiment (Sections 2.8.4 and 2.9). However, other pointing modes not used for mapping science targets *were* carried over from BLASTPol to SPIDER with very few changes, including Az-El Goto, Anti-Sun, RA-Dec Goto, and Drift (Sections 2.8.1 through 2.8.3).

[3] El scans were attempted, which would have provided additional cross-linking. However, the bolometer timestreams obtained during such scans contain large drifts, possibly due to shifting cryogens, or changing air mass with elevation.

#### 3.4.2.3 Pointing Reconstruction

In flight, the telescope pointing was calculated in real-time to better than $5'$ RMS. This solution was computed using a suite of fine and coarse pointing sensors. These included fibre optic rate gyroscopes, a dGPS, an elevation encoder, inclinometers, a magnetometer, sun sensors, and a pair of CCD-based daytime star cameras [94]. For historical reasons, these were known as the Integrating Star Camera (ISC) and the Other Star Camera (OSC). Each star camera had an independent computer capable of identifying stars down to $8^{\text{th}}$ magnitude. It could then compare observed star patterns to a database in order to determine position on the celestial sphere at $\sim$1 Hz. Post-flight pointing reconstruction used only gyroscope-integrated star camera pointing solutions. The post-flight reconstruction method is an extended Kalman filter algorithm [95] similar to that used by WMAP [96]. Post-flight pointing reconstruction accuracy of $< 5''$ RMS has been achieved for BLASTPol. Many more details of pointing reconstruction for BLASTPol and Spider can be found in [69] and [97].

## 3.5 The Power System

The power system described in Section 2.5 was first developed for BLASTPol. Therefore, both *BLASTPol10* and *BLASTPol12* flew with the batteries, solar cells, charge controller, and Power Breakout Board (PBOB) described there. These were wired in a configuration similar to that shown in Figure 2.6. Having a lower power budget than Spider, BLASTPol required only two $3\times3$ arrays of solar panels, as opposed to two $3\times4$ arrays. In *BLASTPol10*, these two arrays were wired in parallel; the power system was not split into two separate systems for the inner frame and outer frame. Therefore, there was only one charge controller, and the two 24 V battery packs[4] were wired in parallel

[4]Once again, each 24 V battery pack consisted of the series combination of two ODYSSEY® PC1200 12 V lead-acid battery units.

to provide enough capacity for the whole payload. In *BLASTPol12*, the power systems for the inner frame and outer frame were split for the purpose of isolation, necessitating a second charge controller.

## 3.6 The 2012 BLASTPol Flight and Data Analysis

### 3.6.1 The LDB Flight

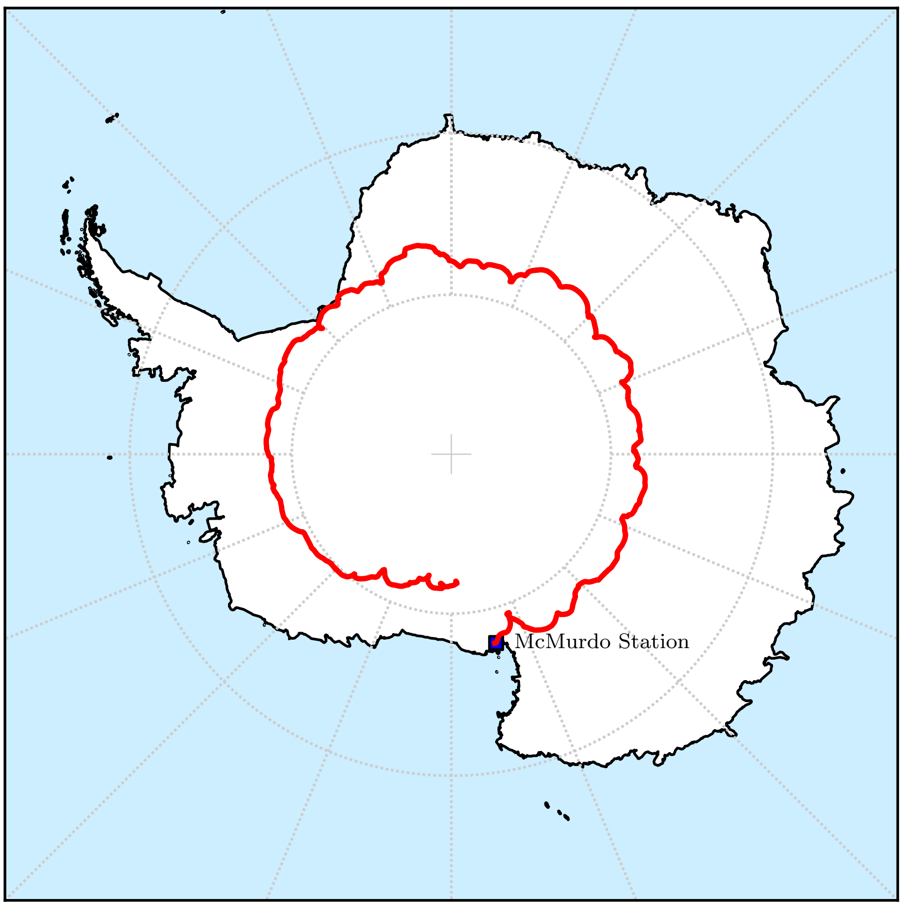


Figure 3.5: The *BLASTPol12* flight path around Antarctica

This section summarizes the 2012 BLASTPol flight. Details of *BLASTPol10* can be found in [82]. BLASTPol launched from the Antarctic Long Duration Balloon (LDB) Facility, which is near Williams Field (77°52.75′ S, 167°3.63′ E), at 18:57 UTC on December

25, 2012. The flight lasted 16 days, 3 hours, and 17 minutes, ending on January 10, 2013 at 22:14 UTC. The payload landed near the Ross Ice Shelf at $(81°58.07' \mathrm{S}, 177°51.32' \mathrm{E})$, approximately 500 km from the launch site. Figure 3.5 shows the flight path around the continent. The instrument was fully recovered in the first week of February, having suffered only minor damage. The helium tank lasted 12.5 days, which was consistent with the expected hold time. The nitrogen tank was depleted shortly before termination.

Table 3.3: Targets observed in *BLASTPol12*. The last two entries were wide area maps made after the failure of the second of the two star cameras.

| Target | Type[a] | Distance[b] [pc] | Map Area[c] [deg$^2$] | Obs. Time [h] |
|---|---|---|---|---|
| Vela C/Vela C Ref. | GMC | 700 | 3.9/14.0 | 43.91/10.80 |
| Carina Nebula | GMC, C | 2300 | 2.7 | 4.19 |
| CG 12 | DC | 550 | 0.1 | 1.72 |
| G331.5-0.1 | GMC, C | 7500 | 3.6 | 4.35 |
| Lupus I | DC | 155 | 1.7 | 15.38 |
| IRAS 08470-4243 | C | 700 | 1.0 | 4.94 |
| Puppis | MC | 1000 | 1.4 | 13.18 |
| VY CMa | C (Star) | 1200 | 0.4 | 4.97 |
| Puppis Wide | — | 1000 | $\sim$113 | 38.64 |
| Vela Wide | GMC | 700 | 23.1 | 89.49 |

[a]GMC = Giant Molecular Cloud, C = Calibrator, DC = Dark Cloud, MC = Molecular Cloud

[b]all distances are approximate

[c]map areas are from [63], with the exception of CG 12, whose map area is taken from [84]

During the flight, most of the instrument subsystems performed very well. The one major system failure occurred to the star camera solid-state drives (SSDs). The OSC drive failed near the beginning of the flight, leaving only the ISC for pointing. The ISC failed six days into the flight, leaving the instrument to rely on the coarse sensors, and requiring a change in the observing strategy to wider maps. These failures are suspected to be due to a known firmware bug in Intel 320 SSDs that may have caused the star camera computers to fail to boot properly from the drives after a cosmic ray upset.

In spite of this problem, degree-scale polarization maps of Galactic molecular clouds were obtained. Table 3.3 lists the majority of BLASTPol targets observed in 2012. It excludes some targets observed for calibration, such as Saturn and IRAS 15100-5613. The first entry in the table lists areas and observing times for both the main Vela C science map, and for the Vela C reference map. The Vela C reference scan was carried out to improve cross-linking and to map the extended region of lower dust column density surrounding the target. The number of hours spent observing each target is estimated from the total number of data frames obtained for each source that were not flagged out due to poor quality. One flag in particular occurred when the Tracking and Data Relay Satellite System (TDRSS) transmitters were turned on for the telemetry downlink and for commanding. These were found to produce significant noise in the receiver. In total, some 300 hours of data were recorded.

### 3.6.2 Data Analysis

The experiment produces a timestream of signal values, or Time-Ordered Data (TOD) for each bolometer. These must then be combined together along with pointing information to produce maps of the sky. This operation, known as map making, is described extensively in [63]. Some low-level data processing steps (sometimes referred to collectively as *data reduction*) must be applied to the TOD before they can be used to produce maps of sufficient quality for scientific analysis. In the BLASTPol data analysis pipeline, some of these steps include:

1. *Despiking and Deconvolution* – the removal of spikes in the bolometer signals from cosmic rays, and the characterization and removal of the bolometer impulse response from the timestreams.

2. *Timestream pre-processing* – the flagging of bad data in the timestreams associated with HWP moves, TDRSS transmission, fridge cycles and other events. This step also involves fitting of polynomial or exponential functions to remove long-timescale thermal or other drifts from the TOD. Elevation-dependent features are specifically identified and de-correlated [68].

3. *Calibration pulse fitting* – the use of pulses from a lamp that regularly illuminated the bolometer arrays during flight (Section 4.5) to remove the effects of detector gain drift.

4. *Characterization of instrumental effects* – a very broad category that includes the determination of pointing offsets, beam non-idealities (Chapter 4), in-flight instrumental polarization [98], and other effects.

In addition to these low-level data processing tasks, the major task of post-flight pointing reconstruction must also precede map making. The result of this effort is the generation of the telescope RA, Dec, and $\phi$ (parallactic angle) for as many samples of the TOD as possible.

# Chapter 4

# BLASTPol Beam Analysis

Characterizing the in-flight beam, or point spread function (PSF), of the telescope optics is a crucial step in the data analysis pipeline. Both flights of BLASTPol were affected by non-ideal structure in the beam that led to degraded image quality. In *BLASTPol10*, the beam shape was affected by an IR blocking filter at the cryostat window that was melted by the sun during ascent. This problem did not occur during *BLASTPol12*. However, the PSFs from the 2012 flight still exhibit unusual structure for instrumental reasons that are unknown at present. The beams are elongated, and most of their power is split between a primary and a secondary lobe that are separated from each other along the elongation axis. (Figure 4.1). A tertiary peak of much lower power is also visible to the right of the main lobes. This chapter describes the effort that was undertaken to characterize the *BLASTPol12* beam shape in each band. The goal of this work was primarily to develop an effective beam template for use elsewhere in the data analysis pipeline. Other uses for the beam modelling included determining the time-variable offset between the star camera and the telescope bore-sight (Section 4.4.2), determining the relative bolometer pointing offsets for each array (Section 4.4.3), and determining a flat-fielding coefficient for each bolometer (Section 4.5). In most cases, each peak of the beam was well-parameterized by a two-dimensional elliptical Gaussian function. This model was fitted using a least squares algorithm (Section 4.1). The beam analysis used maps of the bright compact source IRAS 08470-4243. These maps were made using the naive map makers naivemap and naivepol [92], which were developed for past BLAST and BLASTPol data analysis.

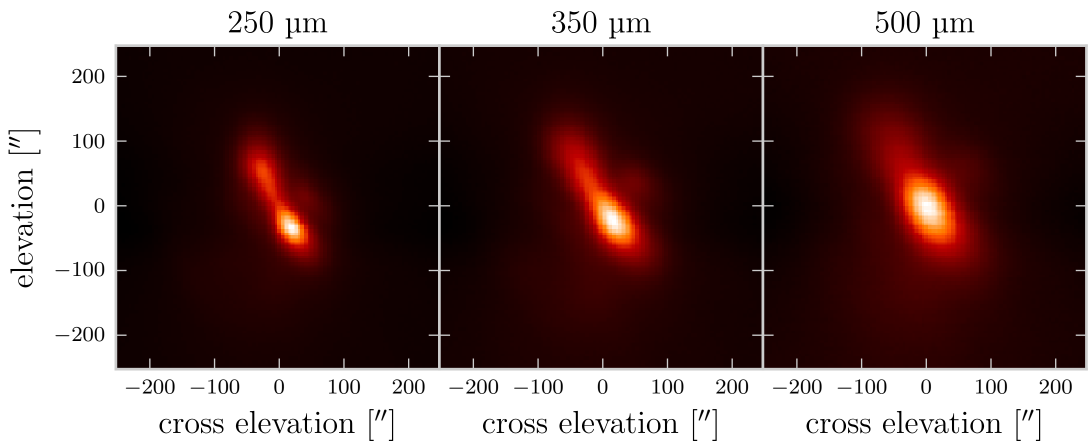


Figure 4.1: General features of the *BLASTPol12* beam in each band. Each image is a map of IRAS 08470-4243 in telescope coordinates. The 5″ pixel scale maps were made by combining all bolometers in each band, and all observations of this source throughout the flight.

As shown in Figure 4.2, maps of IRAS 08470-4243 consistently have the intensity of secondary peak at $\gtrsim 60\%$ of the primary peak intensity, while the tertiary intensity is between 20% and 30% of the primary. However, a corresponding map of Saturn also shown in Figure 4.2 has slightly lower relative intensity values for the secondary and tertiary peaks. This could be due to the known variation of beam properties with time (Section 4.3), or pointing effects. Diffuse structure is also visible at the 10% level in maps of the IRAS source. However, it does not appear in the Saturn maps, suggesting that it is extended structure associated with IRAS 08470-4243, and is not intrinsic to the beam. Although Saturn is closer to being a true point source for BLASTPol, observations of the IRAS source were used for all of the analysis work presented in this chapter, because the 2012 Saturn observations did not include complete bolometer array coverage.

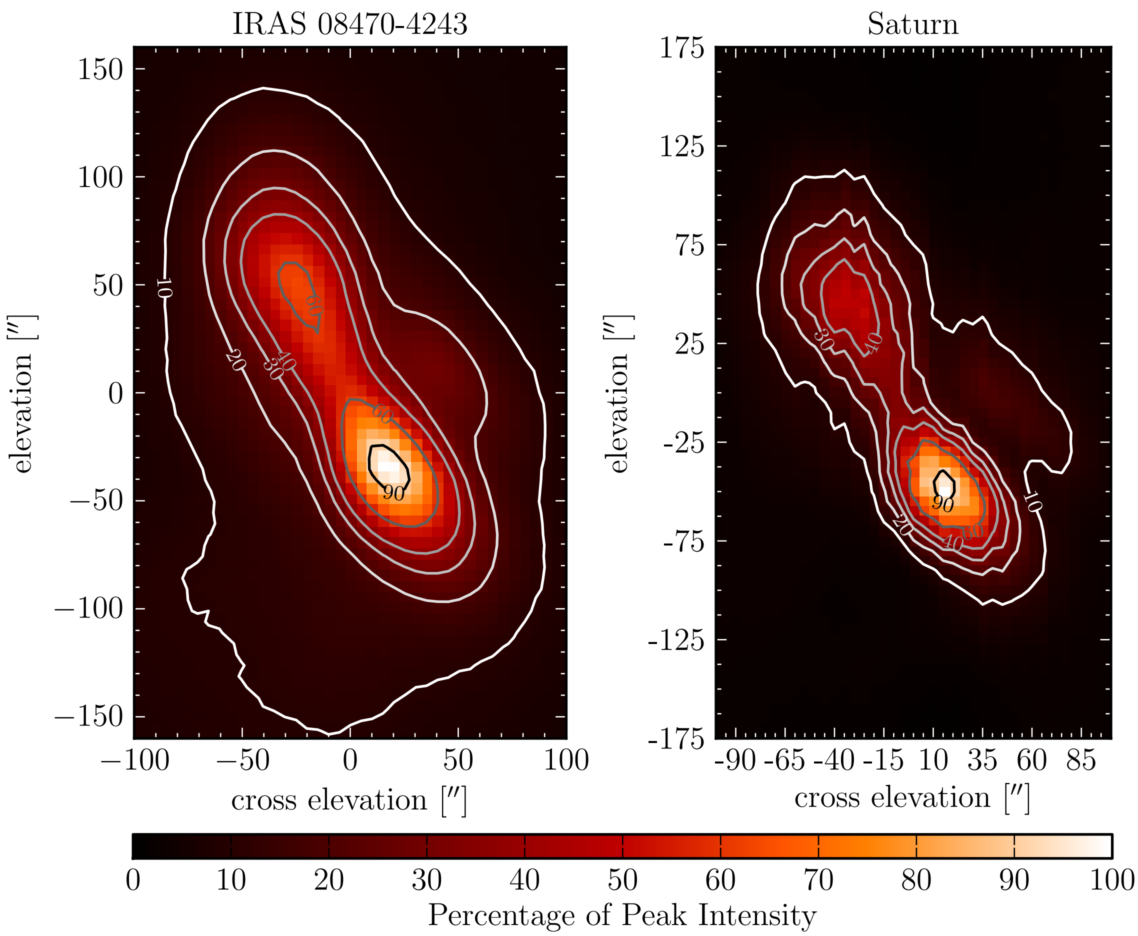


Figure 4.2: *Left* – A more detailed view of the 250 µm map of IRAS 08470-4243 that is shown in the leftmost panel of Figure 4.1. The intensity scale and the overlaid contours show the beam intensity value relative to the brightest pixel. *Right* – A similar relative intensity map of Saturn at 250 µm. The Saturn map includes only 4000 BLASTbus slow frames of data, while the map of the IRAS source combines all of the frames listed in Table 4.2.

## 4.1 Fitting of Elliptical Gaussians

The beam was modelled by fitting a sum of two-dimensional elliptical Gaussian functions to beam maps. In a $(u, v)$ coordinate system aligned with the major and minor axes of the ellipse, the $i^{\text{th}}$ Gaussian in the fit model is of the form

$$f_i(u, v) = A_i \exp\left[-\left(\frac{(u - u_i)^2}{2(\sigma_u^2)_i} + \frac{(v - v_i)^2}{2(\sigma_v^2)_i}\right)\right]. \tag{4.1}$$

If the $(u, v)$ coordinate system is rotated by an angle $\theta_i$ from an $(x, y)$ coordinate system aligned with the horizontal and the vertical (in telescope coordinates), then the expression for $f_i(x, y)$ can be obtained simply by applying the rotation matrix

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \tag{4.2}$$

which results in the expression

$$f_i(x, y) = A_i \exp\left[-\left(a_i(x - x_i)^2 + 2b_i(x - x_i)(y - y_i) + c_i(y - y_i)^2\right)\right], \tag{4.3}$$

where

$$\begin{aligned} a_i &\equiv \frac{\cos^2\theta_i}{2(\sigma_u^2)_i} + \frac{\sin^2\theta_i}{2(\sigma_v^2)_i} \\ b_i &\equiv -\frac{\sin(2\theta_i)}{4(\sigma_u^2)_i} + \frac{\sin(2\theta_i)}{4(\sigma_v^2)_i} \\ c_i &\equiv \frac{\sin^2\theta_i}{2(\sigma_u^2)_i} + \frac{\cos^2\theta_i}{2(\sigma_v^2)_i}. \end{aligned}$$

The fit model $F(x, y)$ is then a sum of $N$ independent Gaussians:

$$F(x, y) = \sum_{i=1}^{N} f_i(x, y) \tag{4.4}$$

where $N = 2$ or 3 depending on the application. To create a template of the time-averaged beam in each band, the model was fitted to 100 pixel $\times$ 100 pixel maps of IRAS 08470-4243 at $5''$ pixel scale. These maps combined data from all bolometers within a given array, and from all observations of this source throughout the flight. Examples of such maps are shown in Figure 4.1. In this case, the fitting included all three peaks. In contrast, for applications where single-bolometer maps were fitted, such as for beam centroiding, the tertiary peak was excluded from the fit. These single-bolometer maps were typically 50 pixel $\times$ 50 pixel maps at $10''$ pixel scale. Some combination of the coarser resolution and variation in beam shape across the arrays caused the three-peak model not to converge for all bolometers. The two-peak model was more stable, and deemed sufficient for these applications. The three-peak model was also not a good fit to any 500 μm map.

Fitting was carried out using the Python `scipy.optimize.leastsq` routine to minimize the residual between the beam map and $F(x, y)$. Typically $A_i$, $x_i$, $y_i$, $(\sigma_u)_i$, $(\sigma_v)_i$, and $\theta_i$ were all free parameters in this optimization. The exception occured at 500 μm where a better fit was achieved with the two-peak model by constraining $\theta_1 = \theta_2$ and $(\sigma_v)_2 = [(\sigma_v)_1/(\sigma_u)_1](\sigma_u)_2$. A pixel brightness threshold was applied; for the two-peak and three-peak fitting, only pixels brighter than 27.5% or 20% (respectively) of the peak map value were included in the fit. These cuts encompassed only the two or three peaks, preventing the diffuse structure of IRAS 08470-4243 from affecting the fit results.

The results of applying this fitting procedure to all-bolometer maps in each band are presented in Figures 4.3, 4.4, and 4.5. The beam fit parameters for these maps are listed in Table 4.1, giving the quantitative properties of the time- and array-averaged beam for each band.

Table 4.1: Beam fit parameters for the fit results shown in Figures 4.3 through 4.5. For each band, these numbers describe the properties of a BLASTPol beam that is averaged over time (all IRAS 08470-4243 observations) and across each array (all bolometers). The full width at half maximum (FWHM) of each Gaussian along each axis is equal to $\sigma(2\sqrt{2\ln 2})$. The $u$ direction is along the major axis of each ellipse, and the $v$ direction along the minor axis. The $\theta$ parameter is the rotation angle of each ellipse, measured clockwise from the horizontal. The coordinates $(r_{\mathrm{sep}}, \theta_{\mathrm{sep}})$ give the positions of the secondary and tertiary peaks relative to the primary peak position, with $\theta_{\mathrm{sep}}$ also measured clockwise from the horizontal.

| | | $A_i/A_1$ | $(\mathrm{FWHM})_u$ [″] | $(\mathrm{FWHM})_v$ [″] | $\theta$ [°] | $r_{\mathrm{sep}}$ [″] | $\theta_{\mathrm{sep}}$ [°] |
|---|---|---|---|---|---|---|---|
| 250 μm | peak 1 | 1 | 73.4 | 41.8 | 57.10 | — | — |
| | peak 2 | 0.669 | 87.7 | 51.4 | 66.75 | 97.4 | 62.78 |
| | peak 3 | 0.260 | 54.1 | 23.7 | 56.99 | 54.3 | −63.88 |
| 350 μm | peak 1 | 1 | 102.3 | 57.3 | 54.51 | — | — |
| | peak 2 | 0.517 | 101.4 | 60.8 | 60.54 | 110.6 | 59.48 |
| | peak 3 | 0.247 | 55.4 | 30.9 | 51.00 | 70.2 | −62.44 |
| 500 μm | peak 1 | 1 | 129.8 | 73.7 | 59.40 | — | — |
| | peak 2 | 0.257 | 74.4 | 42.3 | 59.40 | 121.8 | 61.23 |

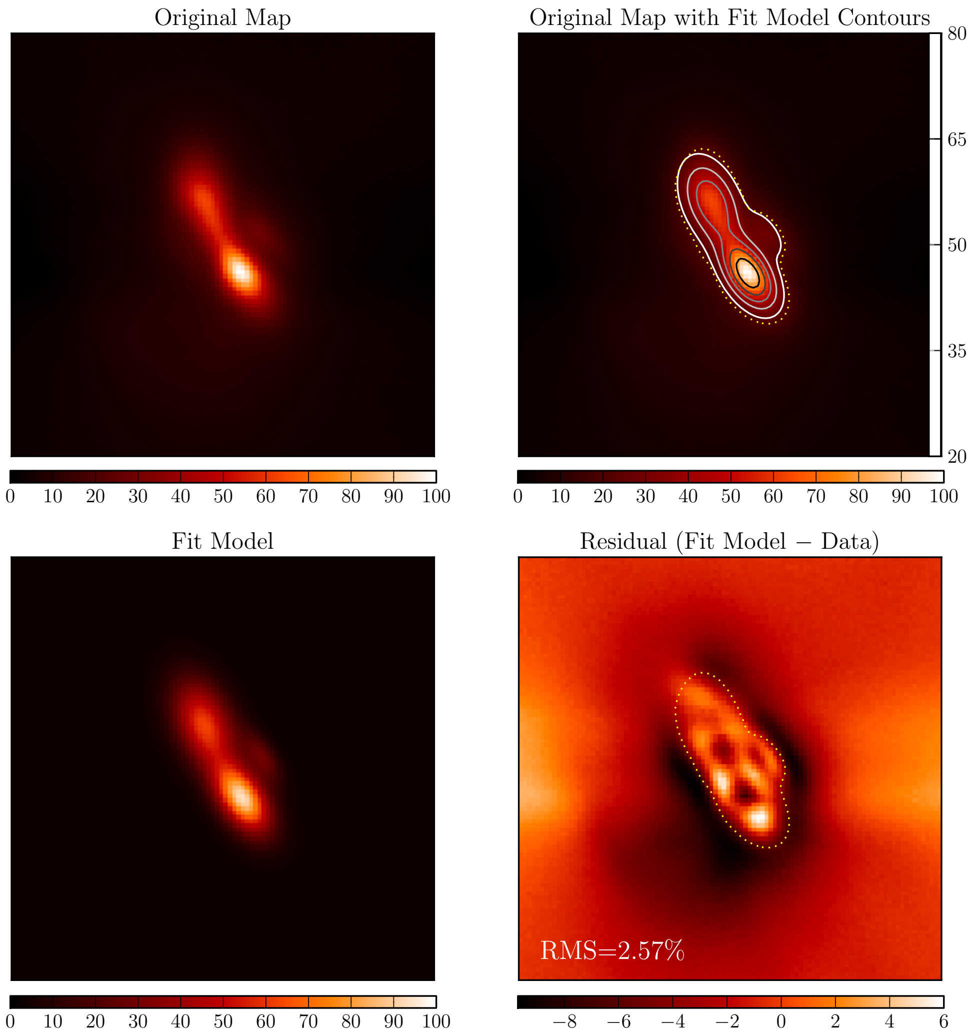


Figure 4.3: Results of applying the fitting procedure to an all-bolometer map at 250 µm. *Top Left* – The original beam map used for fitting. *Top Right* – The original map with solid contours of the fit model superimposed. The yellow dotted contour is the 20% level of the data, enclosing the map pixels that were included in the fit. *Bottom Left* – An intensity map of the fit model. *Bottom Right* – The residual (model − data), expressed as a percentage of the full dynamic range of the original map. The RMS value is indicated. Only the pixels enclosed by the yellow dotted 20% map contour were used in calculating the RMS residual.

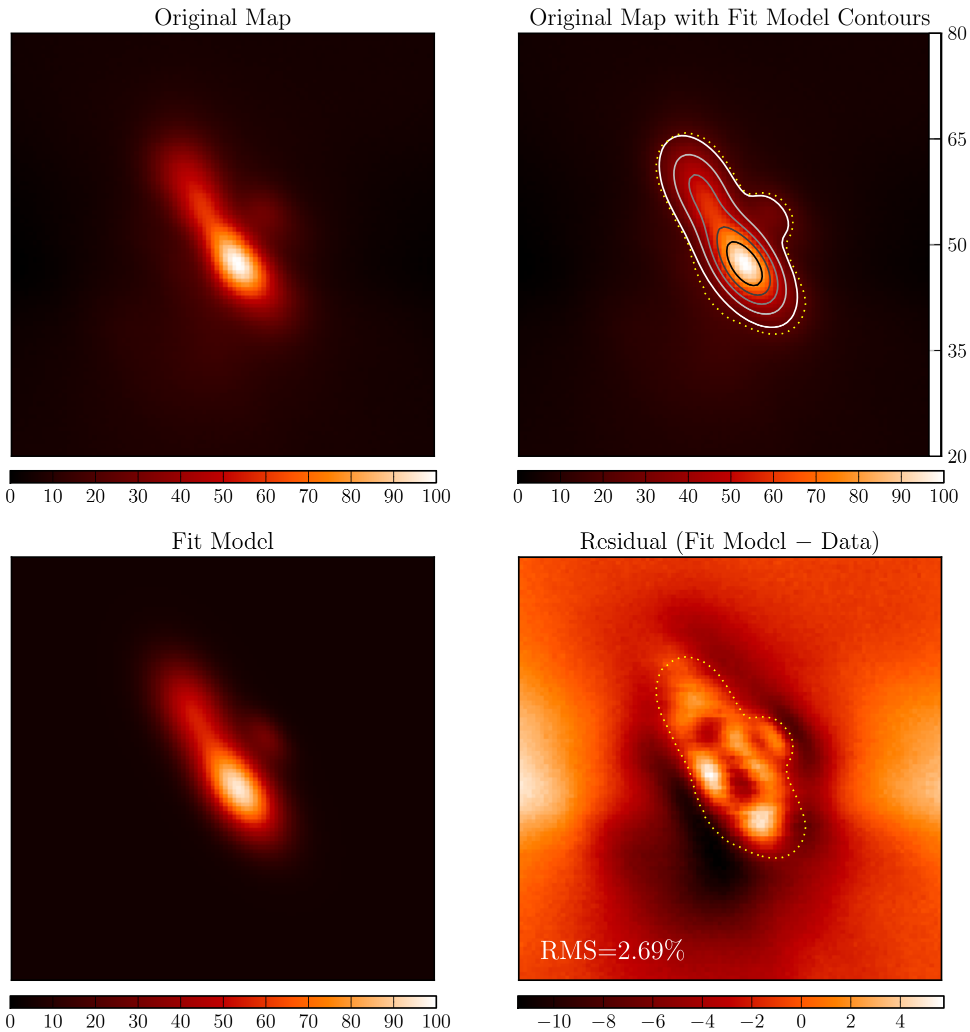


Figure 4.4: Results of applying the fitting procedure to an all-bolometer map at 350 µm. The panels in this figure are analogous to the corresponding ones in Figure 4.3.

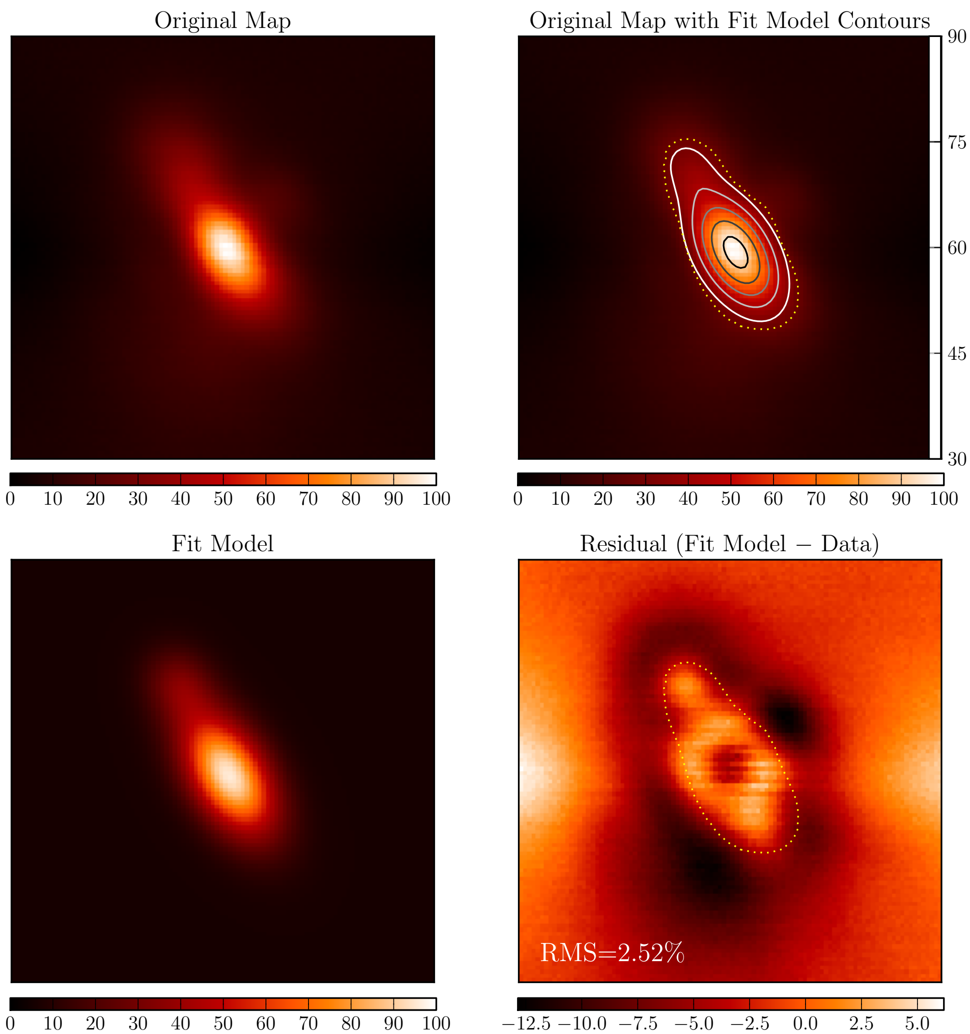


Figure 4.5: Results of applying the fitting procedure to an all-bolometer map at 500 µm. The panels in this figure are analogous to the corresponding ones in Figures 4.3 and 4.4, except that the yellow dotted contour is the 27.5% level of the beam map rather than 20%, since the fit was restricted to the primary and secondary lobes in this case.

## 4.2 Variation of the Beam Across the Arrays

Section 4.1 examined the properties of a time-averaged and array-averaged beam obtained by fitting to naivepol all-bolometer maps in each band. As described in that section, naivepol single-bolometer maps were also made at a two times coarser resolution. By applying the same fitting procedure to each of these maps, it was possible to investigate variations in the beam properties that occur spatially across each focal plane array. The single-bolometer maps are displayed in Figures 4.6, 4.7, and 4.8 below. Each map in these figures is a naivepol Stokes $I$ map of IRAS 08470-4243. The spatial arrangement of the maps in these figures matches that of the bolometers in the corresponding focal plane. Bolometer outlines with no data or lack of a beam indicate bad bolometers.

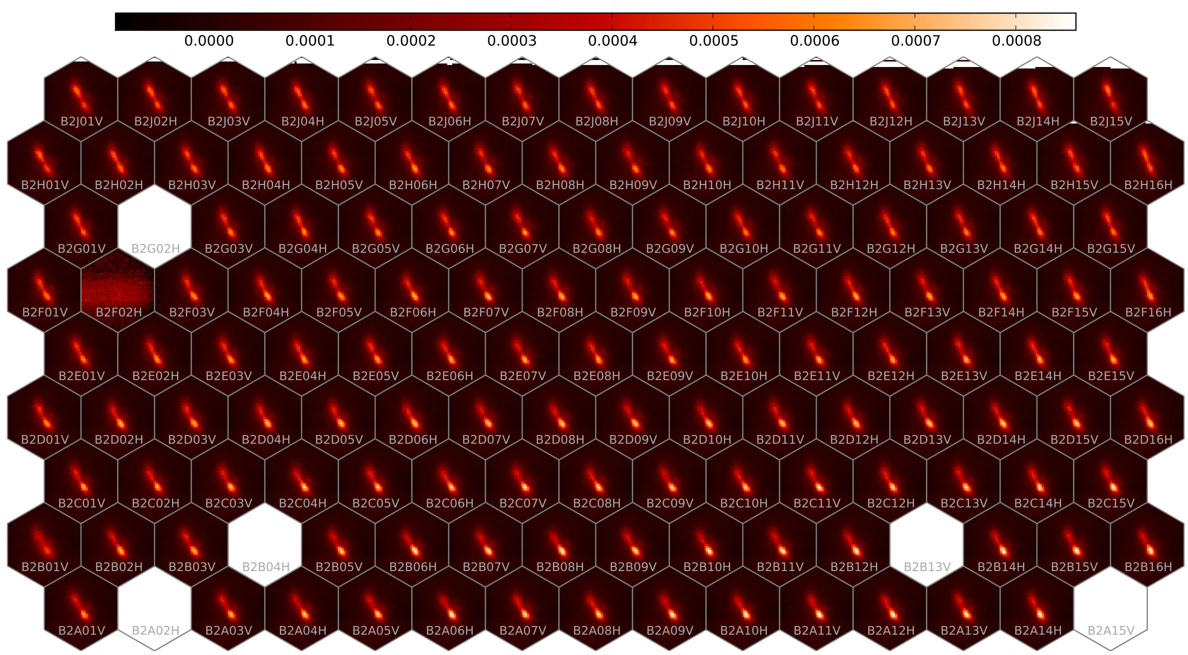


Figure 4.6: Beam maps for each bolometer in the 250 μm array

Certain qualitative trends are apparent from the array maps. To quantify them, each of the fit parameters of the two-peak model is plotted below as a function of bolometer. The fit parameter values are plotted from left to right within a row, and the rows are

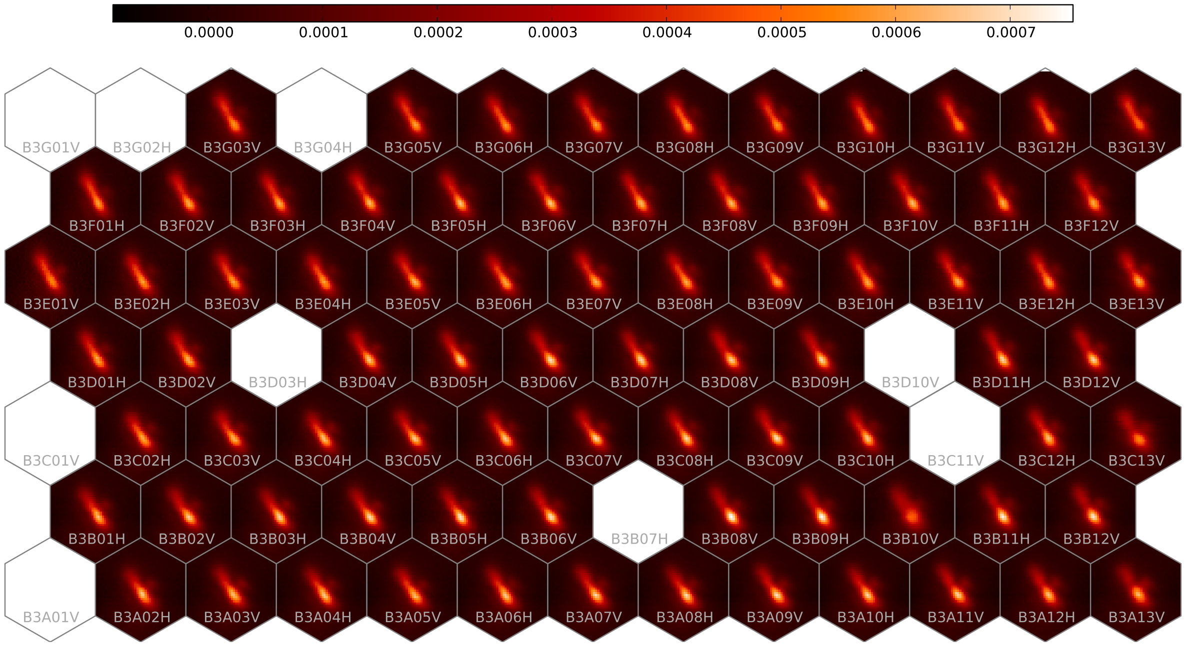


Figure 4.7: Beam maps for each bolometer in the 350 µm array

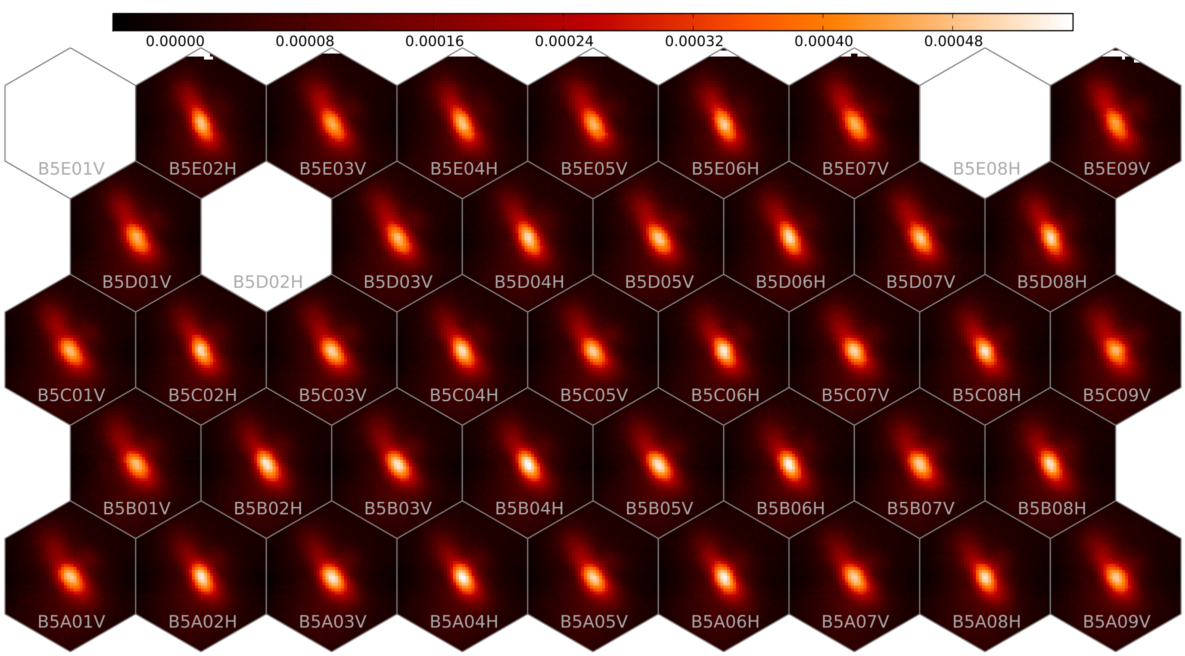


Figure 4.8: Beam maps for each bolometer in the 500 µm array

plotted sequentially from left to right, starting with the bottom row and moving upward, making trends across the array in this direction apparent. Histograms of each fit parameter are also shown. One visible trend in the array maps is a change in the relative intensities of the primary and secondary peaks, particularly at 250 µm and 350 µm. There appears to be significantly more beam power in the primary than the secondary for detectors in the bottom row (row A), but this difference diminishes moving upwards to higher rows, until the two peaks seem comparable in the top rows of the arrays. This trend is confirmed by Figure 4.9, which plots $A_2/A_1$ for each bolometer, for each of the arrays. In the 250 µm array (Fig. 4.9a), this ratio even exceeds unity in the topmost two rows; the primary and secondary beam lobes are reversed.

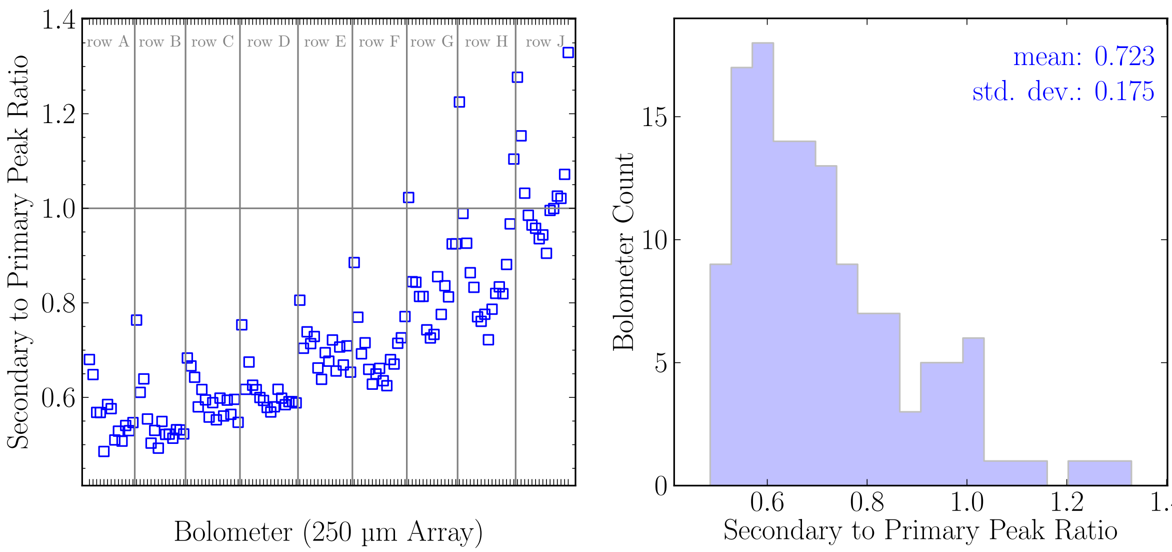


(a) Secondary to primary peak ratio at 250 µm

Figure 4.9: *Left* – Ratio of the secondary to the primary Gaussian peak ($A_2/A_1$) as a function of bolometer in each array. *Right* – Histograms of the peak ratio data.

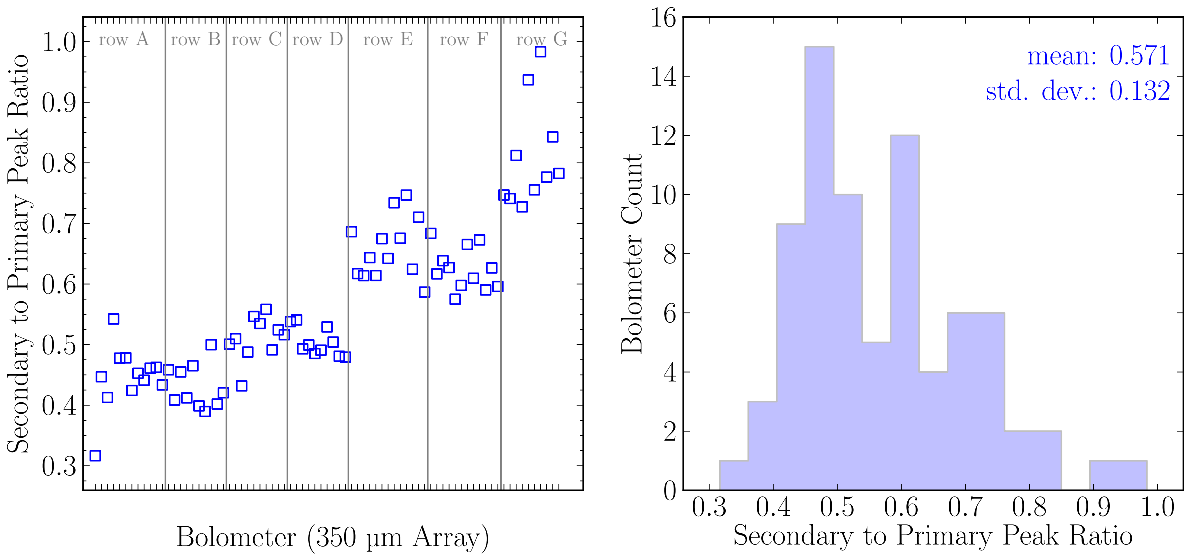


(b) Secondary to primary peak ratio at 350 μm

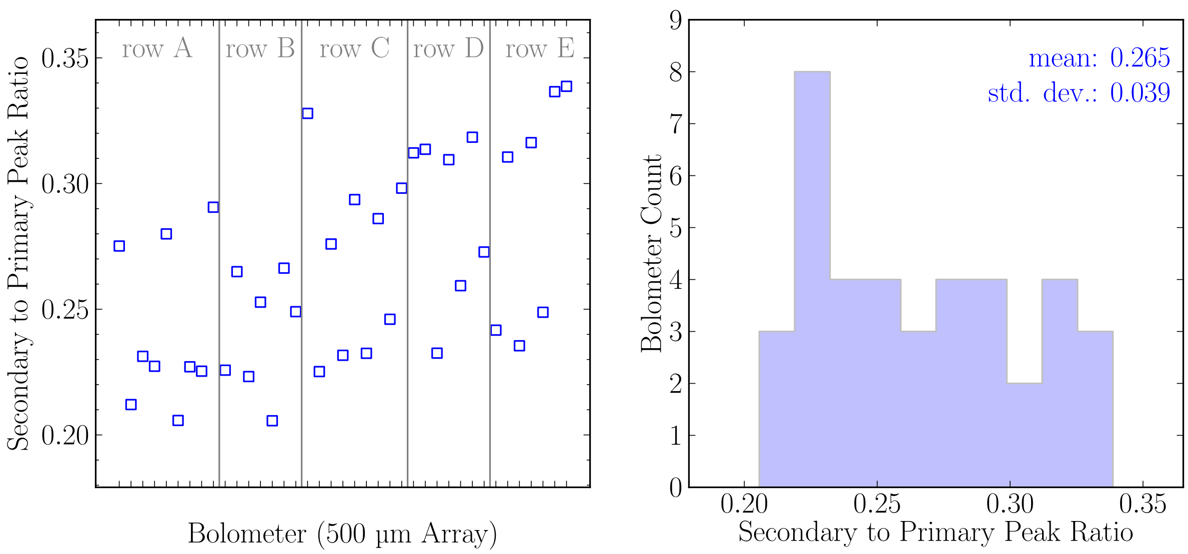


(c) Secondary to primary peak ratio at 500 μm

Figure 4.9: (*Continued*) Secondary to primary Gaussian peak ratio vs. bolometer.

Figure 4.10 examines the overall rotation angle $\theta_{\rm sep}$ of the beams, which is defined as the rotation angle of a line connecting the centres of the primary and secondary peaks, measured clockwise from the horizontal.

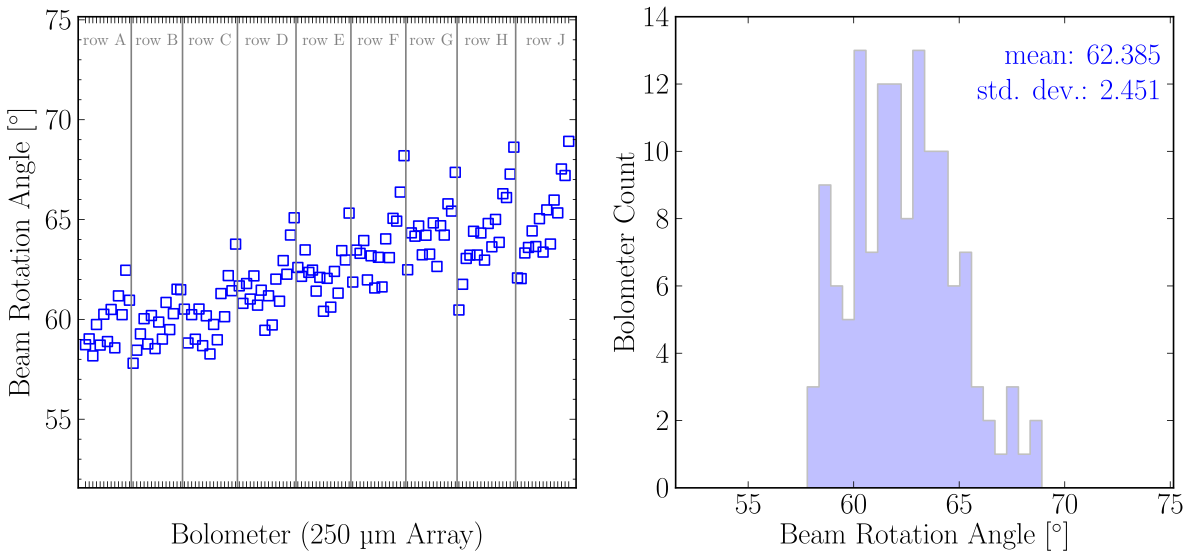


(a) Overall beam rotation angle at 250 µm

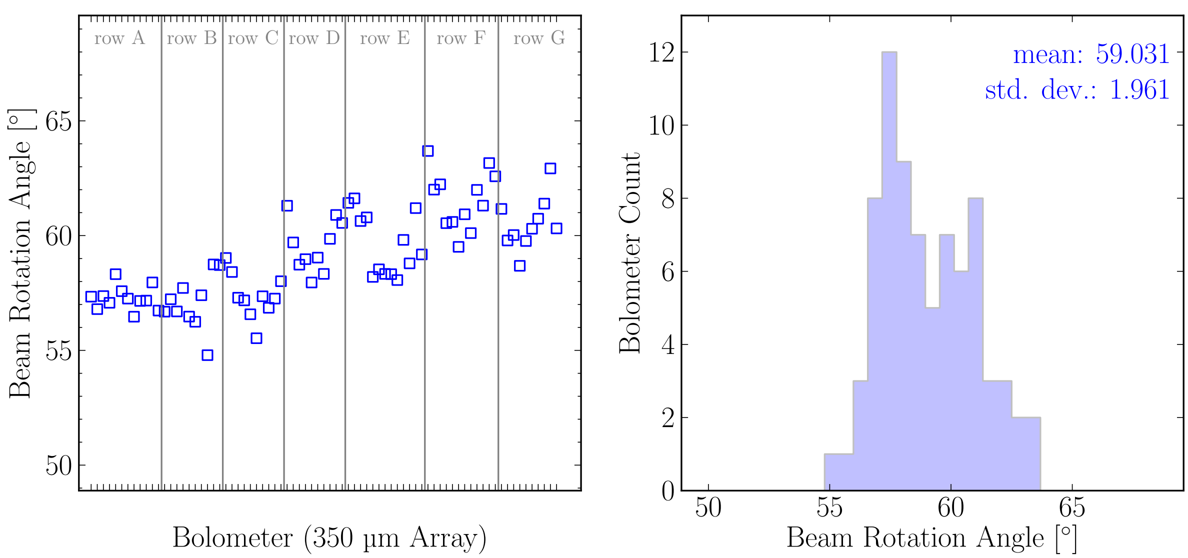


(b) Overall beam rotation angle at 350 µm

Figure 4.10: *Left* – Overall rotation angle ($\theta_{\text{sep}}$) of the beam as a function of bolometer in each array. *Right* – Histograms of the rotation angle data.

This figure shows variations in beam rotation angle of several degrees, as well as evidence that the overall rotation angle gets steeper, meaning the elongation axis becomes closer to vertical, moving upwards in the rows. This measure is also sensitive to changes in the relative locations of the two peaks.

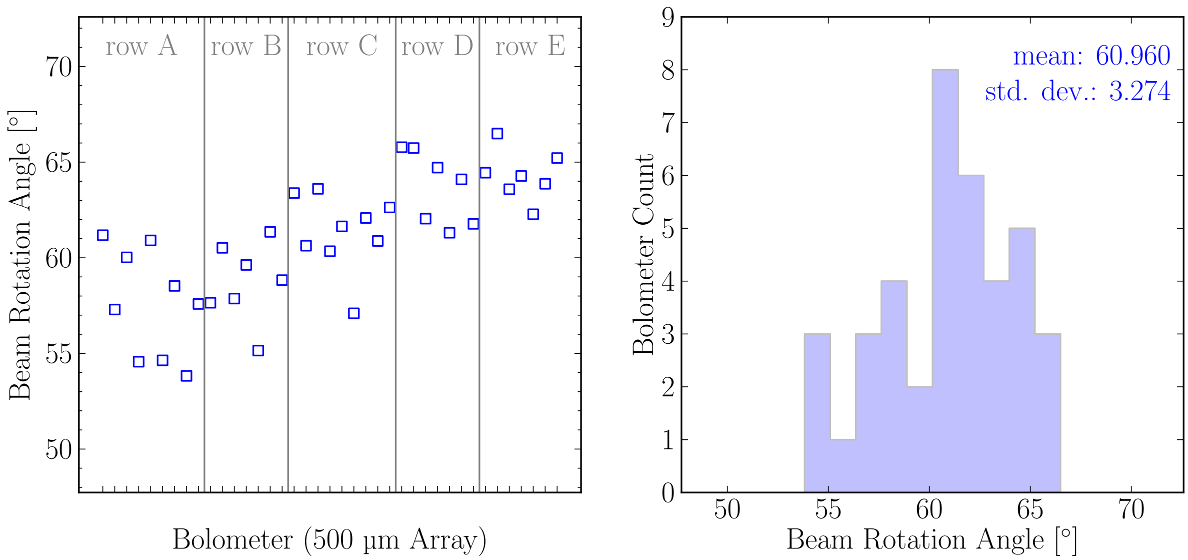


(c) Overall beam rotation angle at 500 μm

Figure 4.10: (*Continued*) Overall beam rotation angle vs. bolometer.

The separation $r_{\text{sep}}$ between the centres of the primary and secondary peaks, in arcseconds, is shown in Figure 4.11. The RMS deviation from the mean separation is at the 5 arcsecond level in the case of all three bands. The 250 μm and 500 μm arrays (Figs. 4.11a and 4.11c) appear to have a slight upward trend in beam elongation with increasing row.

Full widths at half maximum along the major and minor axes of each elliptical Gaussian are plotted in Figures 4.12 and 4.13 respectively. At 250 μm and 350 μm, the FWHM in both directions is comparable between the primary and the secondary peak, with the secondary peak tending to be slightly wider (Figs. 4.12a, 4.12b, 4.13a, and 4.13b). At 500 μm the secondary peak is distinctly smaller (Figs. 4.12c and 4.13c). The lower angular resolution has caused the two distinct beam lobes to blend together, making this asymmetric two-Gaussian model a better fit.

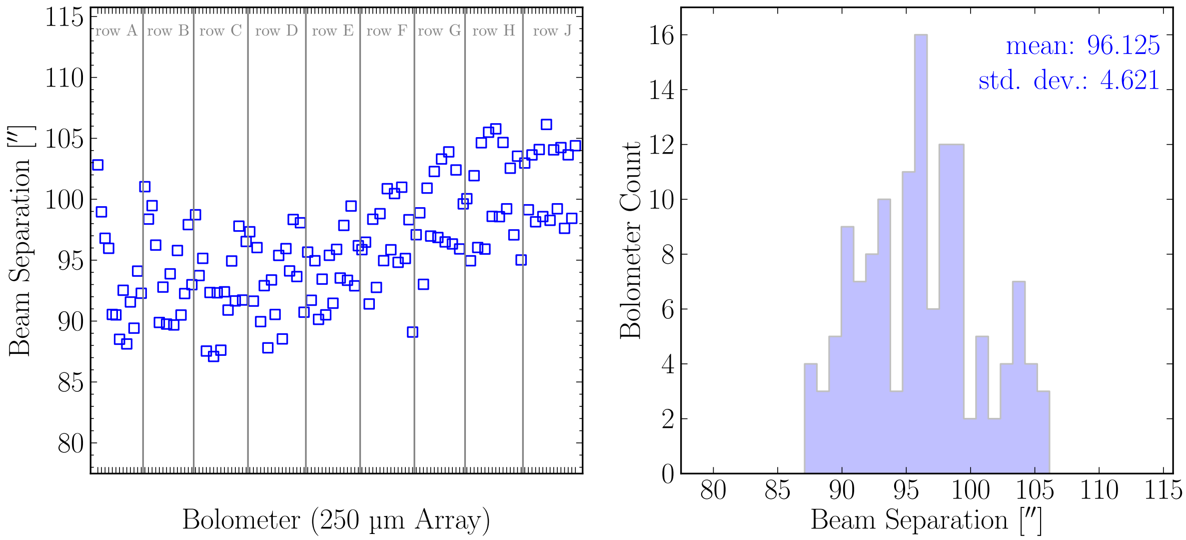


(a) Secondary to primary peak separation at 250 µm

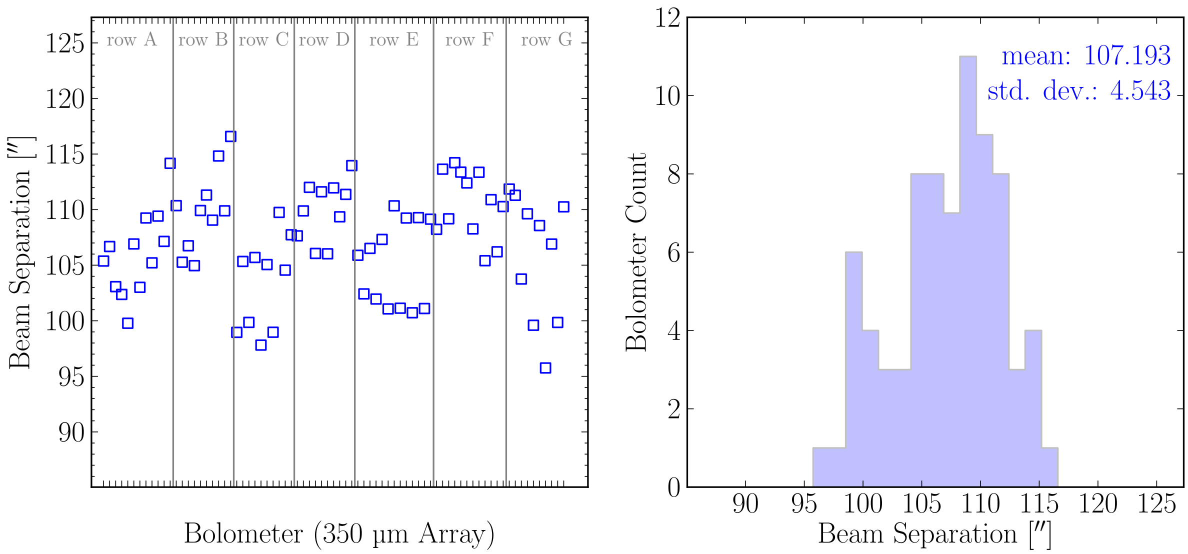


(b) Secondary to primary peak separation at 350 µm

Figure 4.11: *Left* – Separation between the secondary and primary Gaussian peak ($r_{\rm sep}$) as a function of bolometer in each array. *Right* – Histograms of the peak separation data.

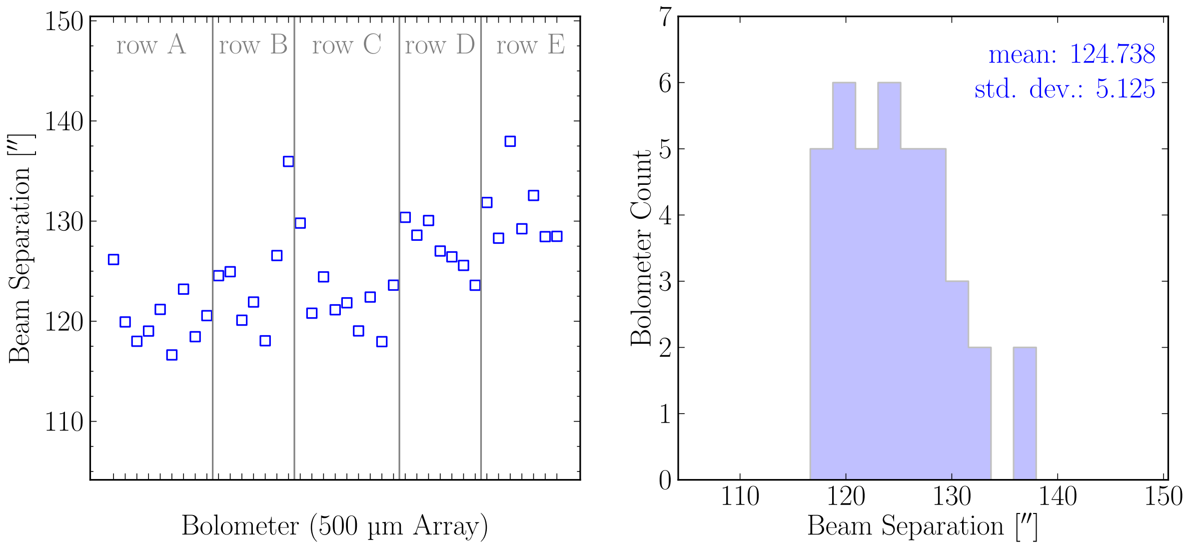


(c) Secondary to primary peak separation at 500 µm

Figure 4.11: (*Continued*) Secondary to primary Gaussian peak separation vs. bolometer.

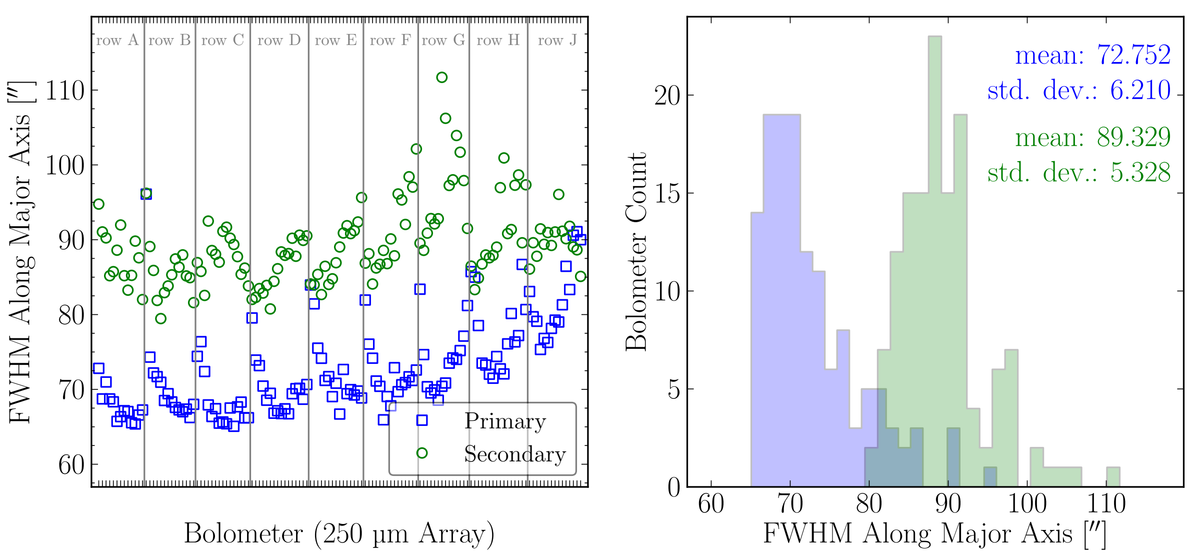


(a) FWHM along the major axes at 250 µm

Figure 4.12: *Left* – FWHMs along the major axes of the primary and secondary Gaussian peaks as a function of bolometer in each array. *Right* – Histograms of the FWHM data.

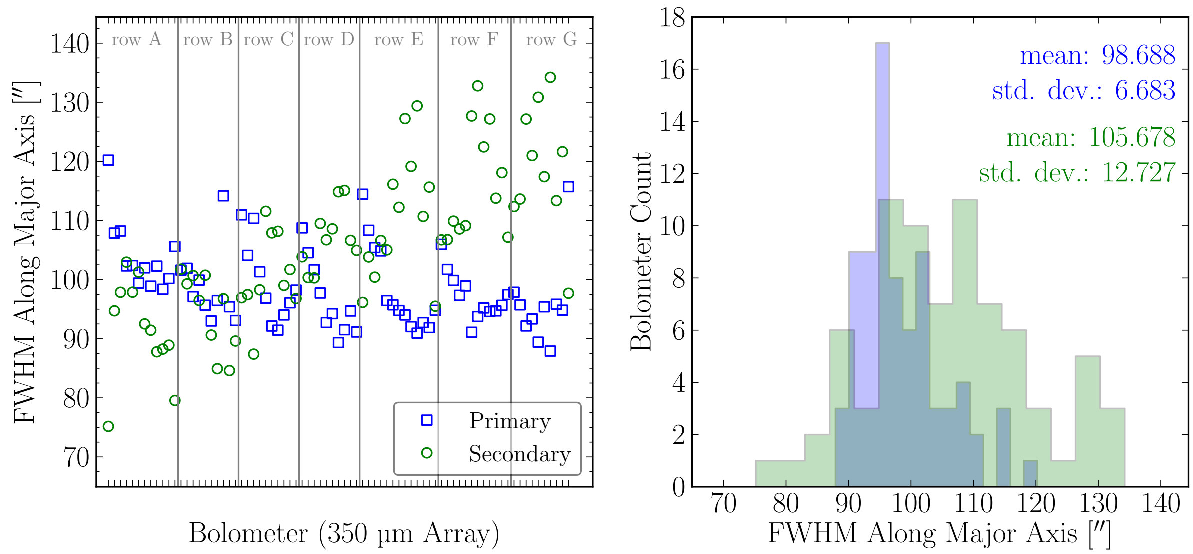


(b) FWHM along the major axes at 350 µm

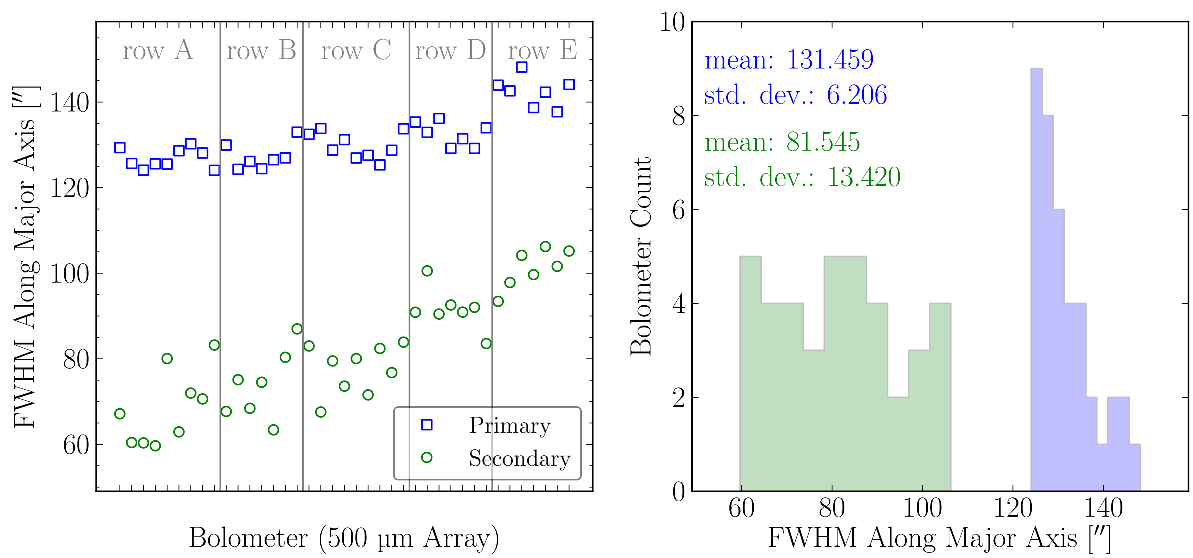


(c) FWHM along the major axes at 500 µm

Figure 4.12: (*Continued*) FWHM along the major axes of the two Gaussian peaks vs. bolometer

Figure 4.14 shows the individual rotation angles $\theta_1$ and $\theta_2$ of the major axes of the primary and secondary Gaussian peaks, measured clockwise from the horizontal. At 250 and 350 µm, the secondary peak rotation angles are systematically higher, meaning that the major axis of the secondary ellipse is closer to vertical. This difference gives the

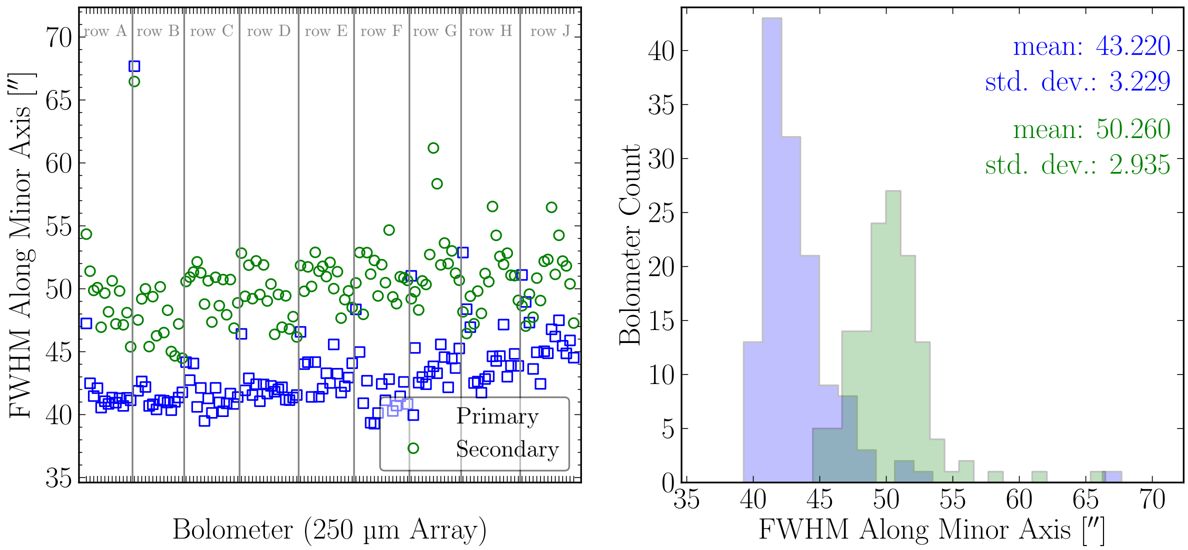


(a) FWHM along the minor axes at 250 μm

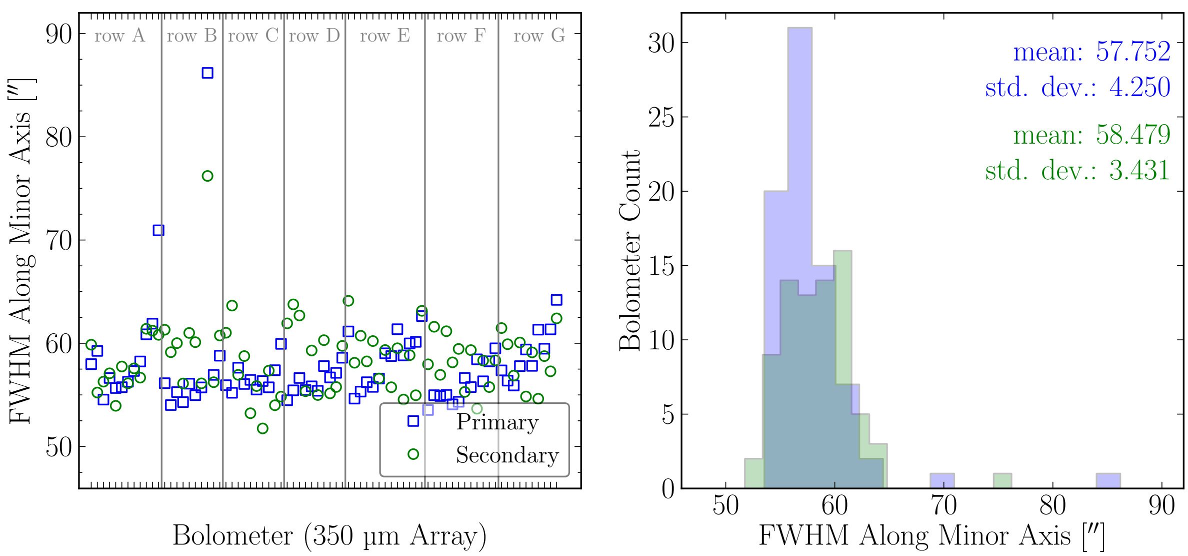


(b) FWHM along the minor axes at 350 μm

Figure 4.13: *Left* – FWHMs along the minor axes of the primary and secondary Gaussian peaks as a function of bolometer in each array. *Right* – Histograms of the FWHM data.

beam a bent appearance. For the 500 μm fitting, the primary and secondary peaks were constrained to have the same rotation angle, so only one set of angles is shown.

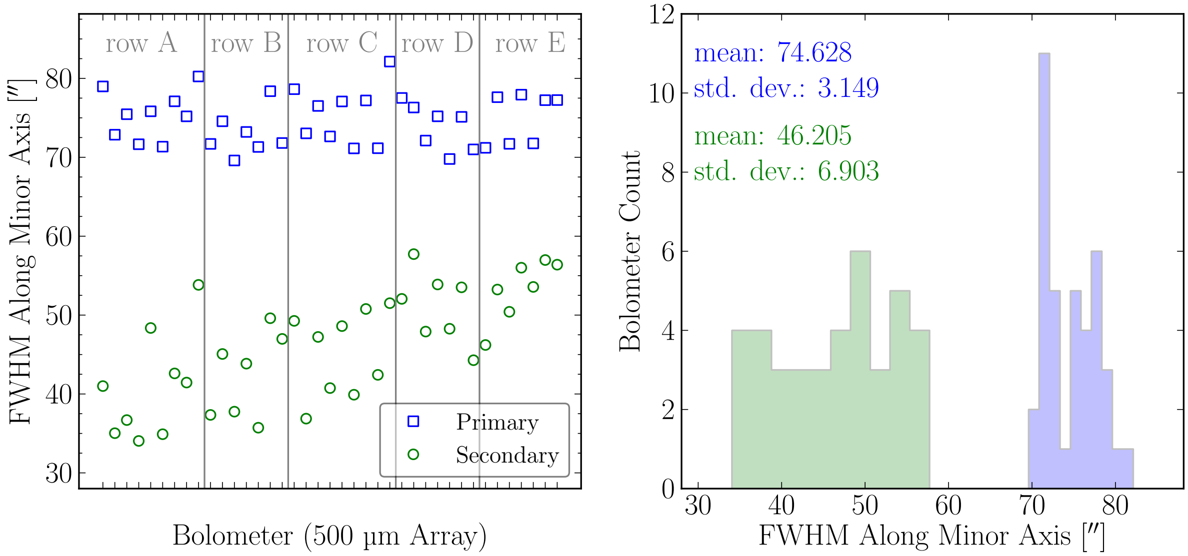


(c) FWHM along the minor axes at 500 µm

Figure 4.13: (*Continued*) FWHM along the minor axes of the two Gaussian peaks.

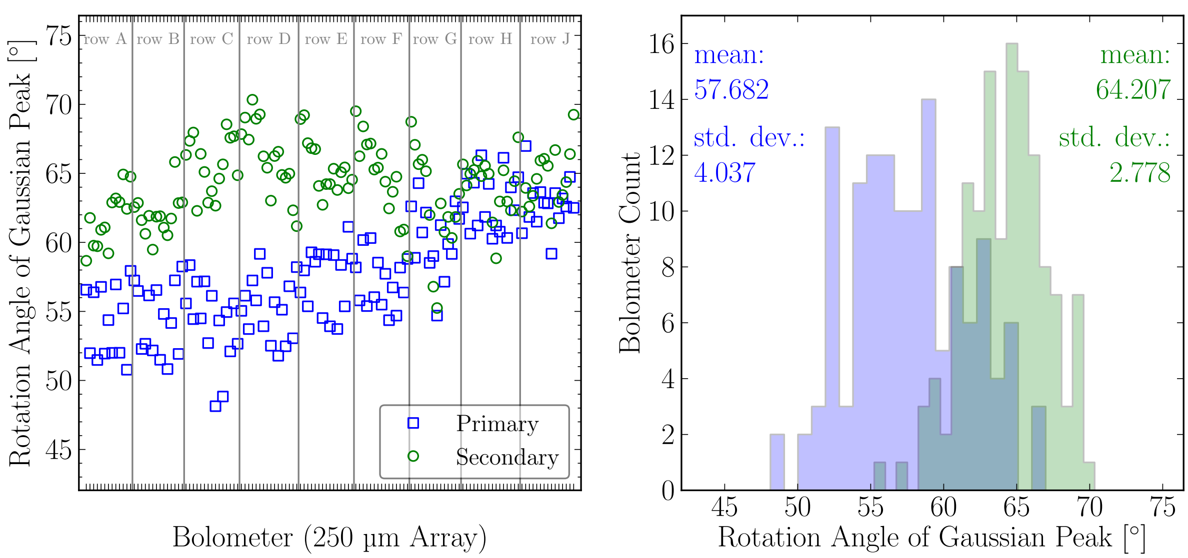


(a) Peak rotation angles at 250 µm

Figure 4.14: *Left* – Rotation angles of the two Gaussian peaks ($\theta_1$ and $\theta_2$) as a function of bolometer in each array. *Right* – Histograms of the peak rotation angle data.

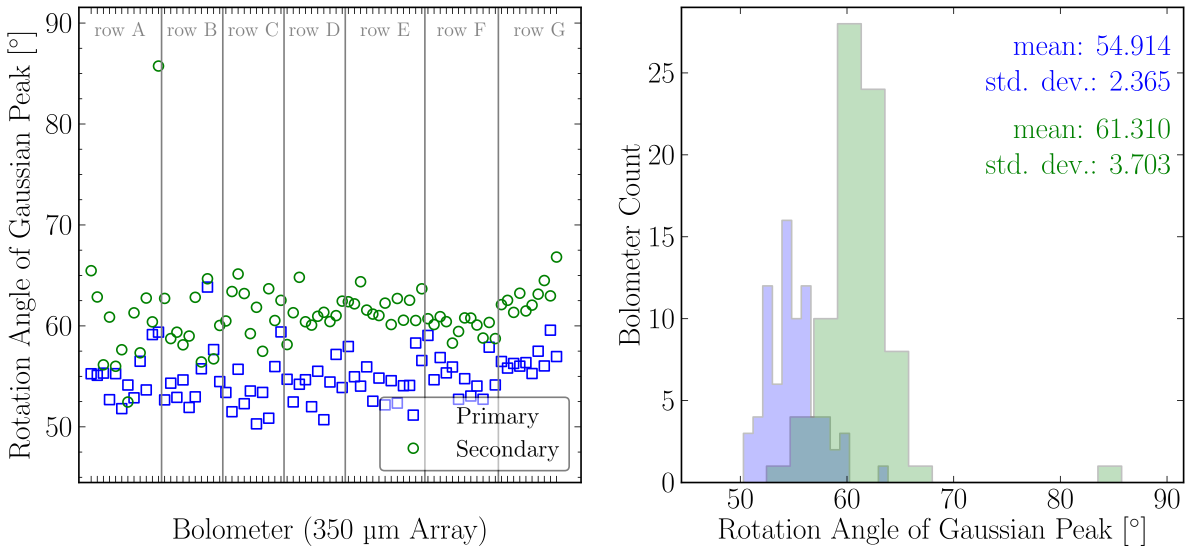


(b) Peak rotation angles at 350 µm

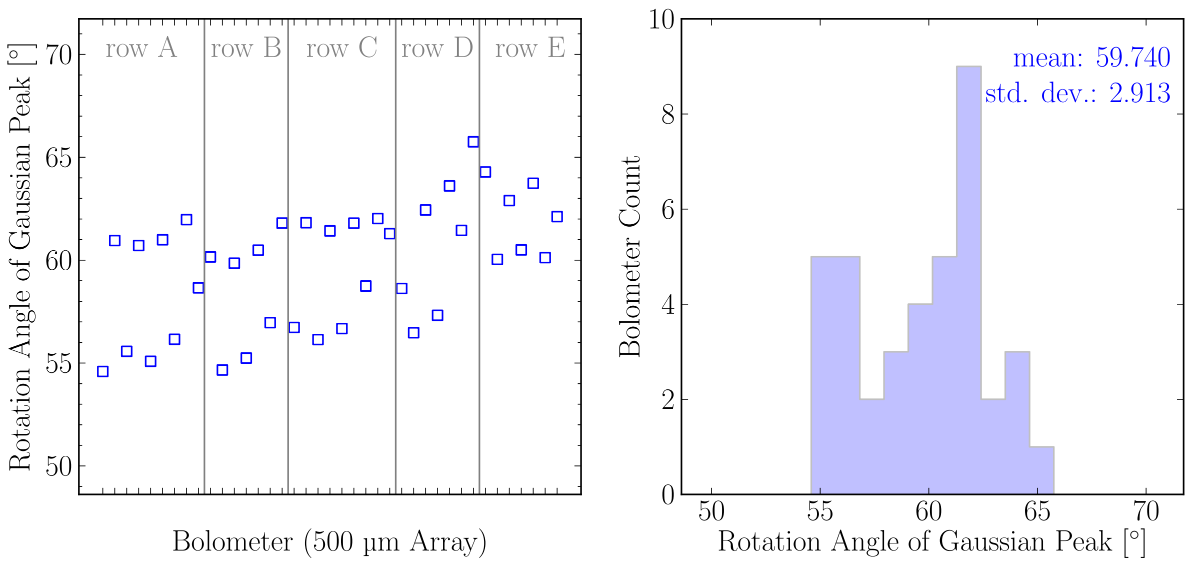


(c) Peak rotation angles at 500 µm

Figure 4.14: (*Continued*) Rotation angles of the two Gaussian peaks vs. bolometer.

## 4.3 Variation of the Beam with Time

Variation of the beam properties with time has been observed in qualitative comparisons of maps of IRAS 08470-4243 made from observations taken at different times during the flight (Figure 4.15).

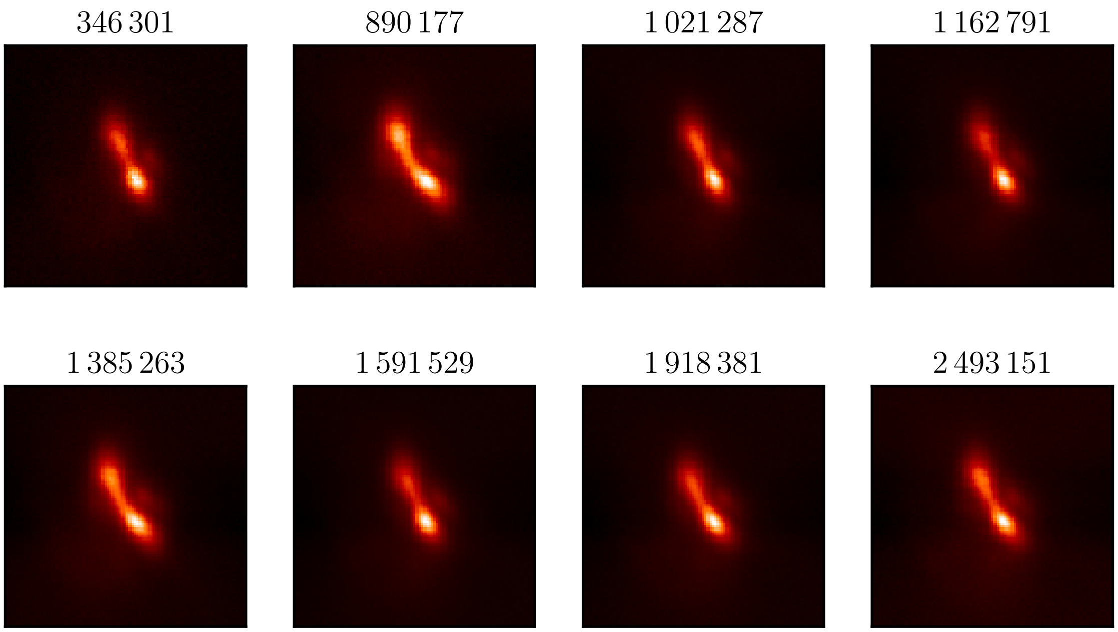


(a) The time series at 250 µm

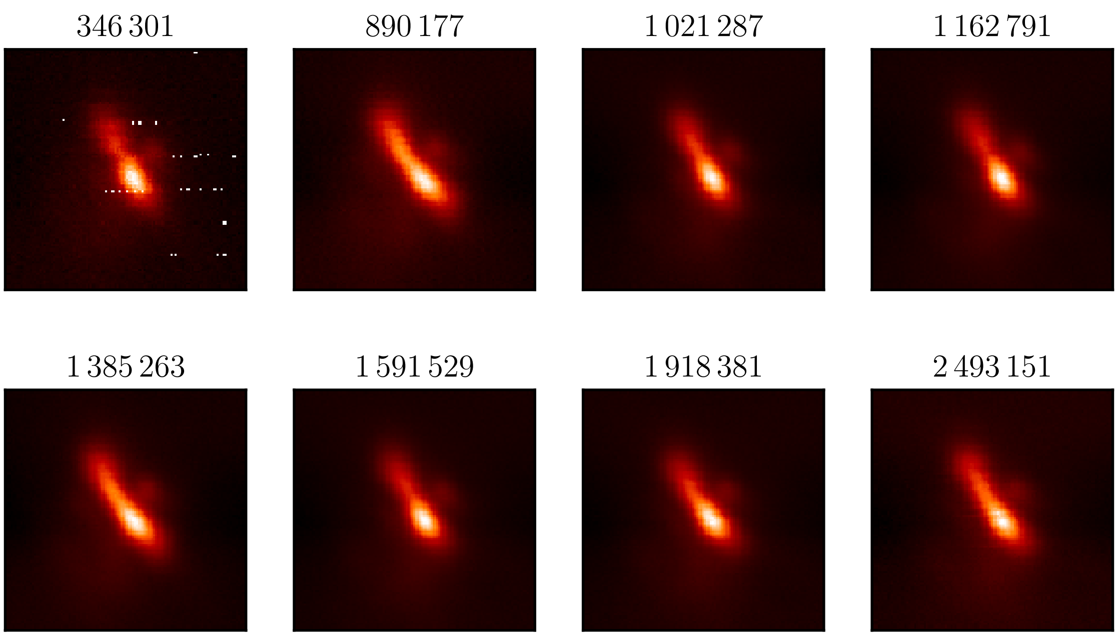


(b) The time series at 350 µm

Figure 4.15: Time series of beam maps using selected data sets from Table 4.2. These images show significant variation in beam rotation angle, peak separation, peak ratio, and FWHM, with time.

To quantify these variations, the two-Gaussian fit model was applied to all-bolometer maps made for each individual data set corresponding to an IRAS 08470-4243 observation. The variation of the fit parameters between maps shows how the array-averaged beam changes between data sets taken at different times. The BLASTbus slow frame numbers and frame ranges of these observations are given in Table 4.2.

Table 4.2: The individual data sets used to produce all-bolometer maps vs. time

| Starting Index | Number of Frames | HWP Position [°] |
|---:|---:|---:|
| 338 197 | 3847 | 38.6 |
| 346 301 | 3915 | 15.5 |
| 350 332 | 3878 | 62.4 |
| 354 299 | 3240 | 83.5 |
| 888 930 | 1245 | 62.0 |
| 890 177 | 2782 | 62.0 |
| 893 048 | 1080 | 83.6 |
| 1 013 569 | 3457 | 83.6 |
| 1 017 160 | 4019 | 38.7 |
| 1 021 287 | 2819 | 16.7 |
| 1 162 791 | 4066 | 38.6 |
| 1 166 967 | 1126 | 15.7 |
| 1 289 215 | 4150 | 61.9 |
| 1 293 454 | 1712 | 83.6 |
| 1 385 263 | 4251 | 83.6 |
| 1 478 084 | 3712 | 16.6 |
| 1 591 529 | 4113 | 17.2 |
| 1 731 380 | 4074 | 15.7 |
| 1 836 689 | 4206 | 83.4 |
| 1 918 381 | 3711 | 61.4 |
| 1 922 181 | 3390 | 83.9 |
| 2 034 851 | 4010 | 16.2 |
| 2 183 497 | 4094 | 38.7 |
| 2 490 201 | 2861 | 61.9 |
| 2 493 151 | 1905 | 83.6 |

The BLASTbus slow data rate is ∼5 Hz, therefore the typical data set size of approximately 4000 slow frames corresponds to just under 15 minutes of data. Each of these observations was carried out at a single HWP position.

The fit model was applied to beam maps for each of the data sets listed, at both 250 µm and 350 µm. Since these are single-HWP-angle data sets, a flux map of the source was made for each one using naivemap, rather than a Stokes $I$ map using naivepol. During the fitting procedure, a pixel brightness threshold of 27.5% of the peak map value was applied at 250 µm. This threshold was increased to 32.5% at 350 µm in order to restrict the fit to the two main beam lobes. Figure 4.16 shows the ratio of the secondary to the primary Gaussian peaks $A_2/A_1$ for each of the beam maps. There is a clear variation back and forth between higher peak ratios of approximately 0.8 and lower peak ratios of 0.5 to 0.6.

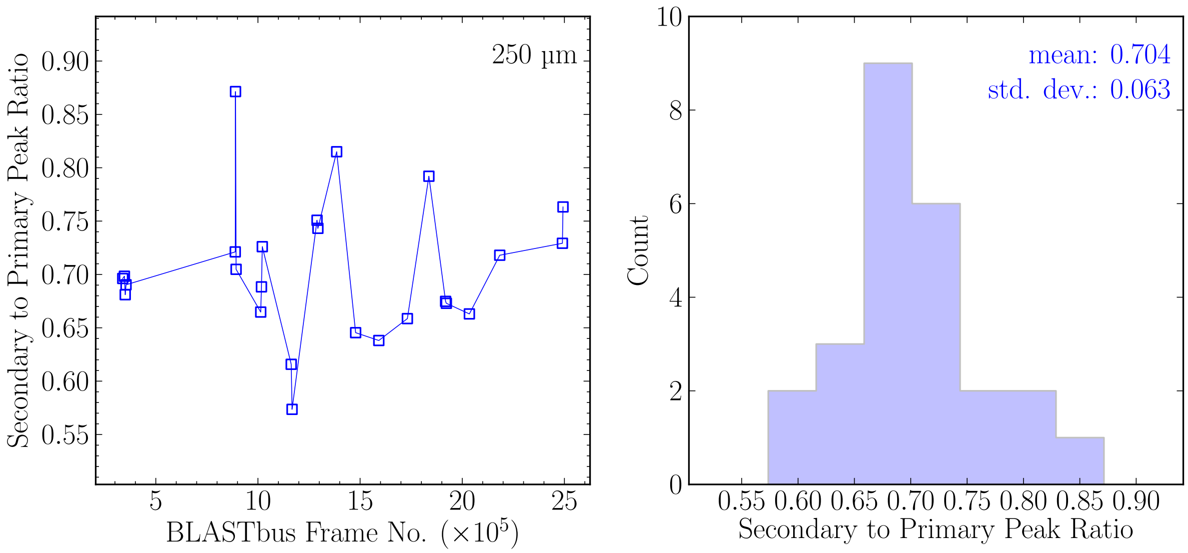


(a) Secondary to primary peak ratio at 250 µm

Figure 4.16: *Left* – Ratio of the secondary to the primary Gaussian peak ($A_2/A_1$) as a function of time. *Right* – Histograms of the peak ratio data.

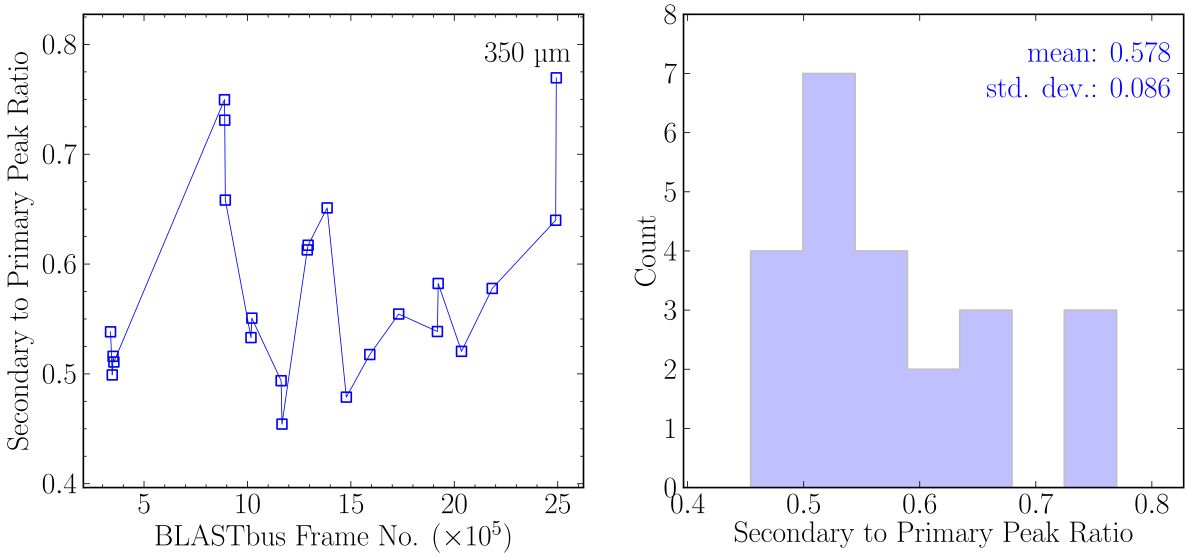


(b) Secondary to primary peak ratio at 350 µm

Figure 4.16: (*Continued*) Secondary to primary Gaussian peak ratio vs. time.

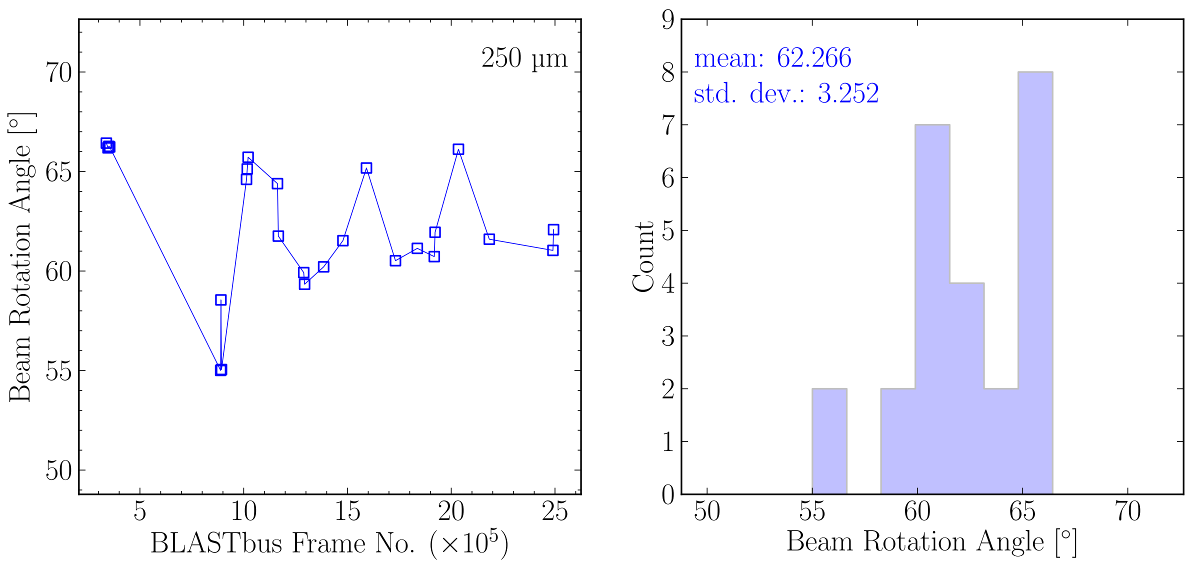


(a) Overall beam rotation angle at 250 µm

Figure 4.17: *Left* – Overall rotation angle ($\theta_{\mathrm{sep}}$) of the beam as a function of time. *Right* – Histograms of the rotation angle data.

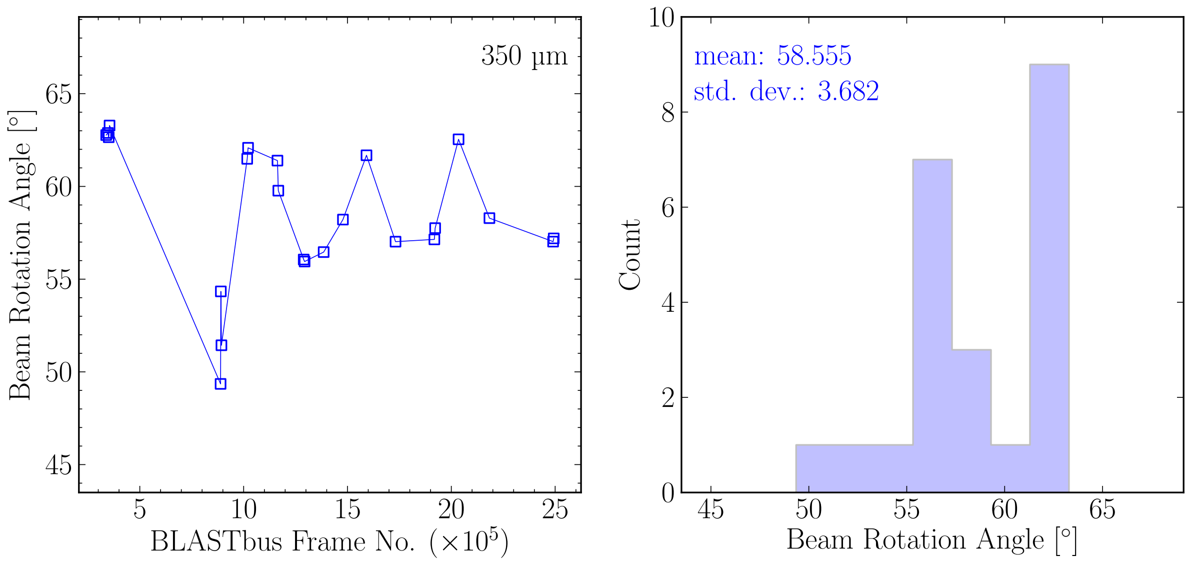


(b) Overall beam rotation angle at 350 µm

Figure 4.17: (*Continued*) Overall beam rotation angle vs. time.

Other fit parameters vary with time in a similar manner. Figure 4.17 shows the overall beam rotation angle $\theta_{\text{sep}}$ versus time. This is the rotation angle of a line connecting the centres of the primary and secondary peaks, measured clockwise from the horizontal. This quantity oscillates between steeper and shallower angles, corresponding to the beam elongation axis moving clockwise towards vertical, and counterclockwise away from vertical, respectively.

Figure 4.18 shows that this oscillation is visible in the peak separation as well. At 250 µm, there is a 10″ variation between the smallest and largest peak separations (Fig. 4.18a). This variation increases to 15″ at 350 µm (Fig. 4.18b). The histograms for this parameter appear bimodal, indicating a clear distinction between the beams with more widely-separated peaks at some times, and the beams with less widely-separated peaks at other times.

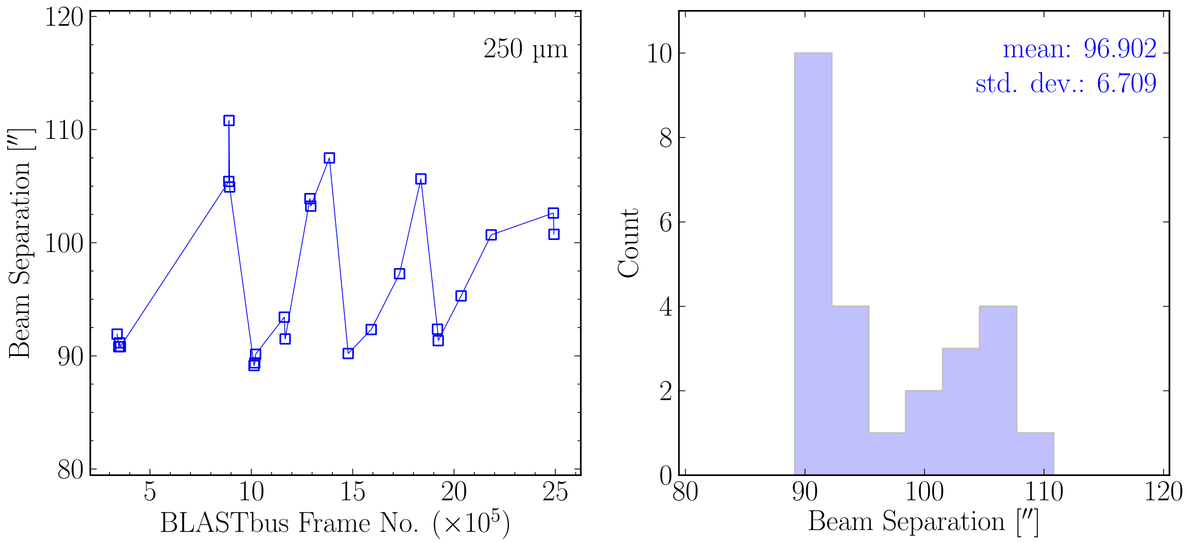


(a) Secondary to primary peak separation at 250 µm

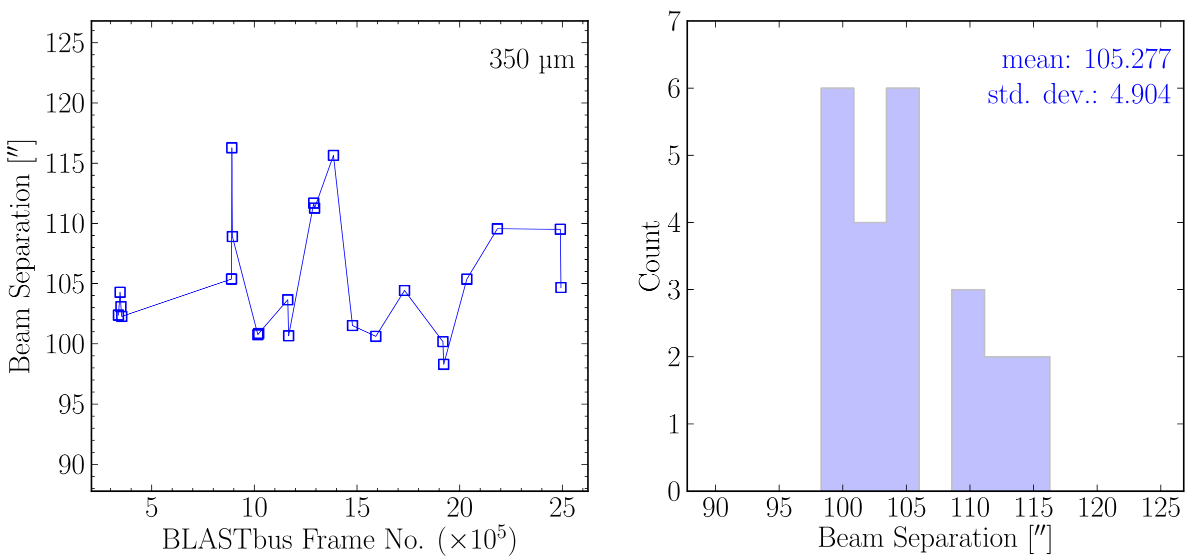


(b) Secondary to primary peak separation at 350 µm

Figure 4.18: *Left* – Separation between the secondary and primary Gaussian peak ($r_{\rm sep}$) as a function of time. *Right* – Histograms of the peak separation data.

The FWHMs of the two peaks along their major and minor axes are shown in Figures 4.19 and 4.20 respectively. At 250 µm, the FWHM along the major axis of the secondary peak remains relatively constant, while that of the primary peak shows 20″ to 30″ variations (Fig. 4.19a). Curiously, the opposite appears to be true at 350 µm; the variations in the FWHM of the secondary are larger than that of the primary (Fig. 4.19b). This discrepancy could simply be a function of how the two-Gaussian fit model accommodates the differing overall beam shape in the two bands. In both cases the FWHM of the secondary peak is systematically larger. This holds true for the widths along the minor axis, although the primary and secondary are much more comparable in size in this direction, and show much smaller variations with time.

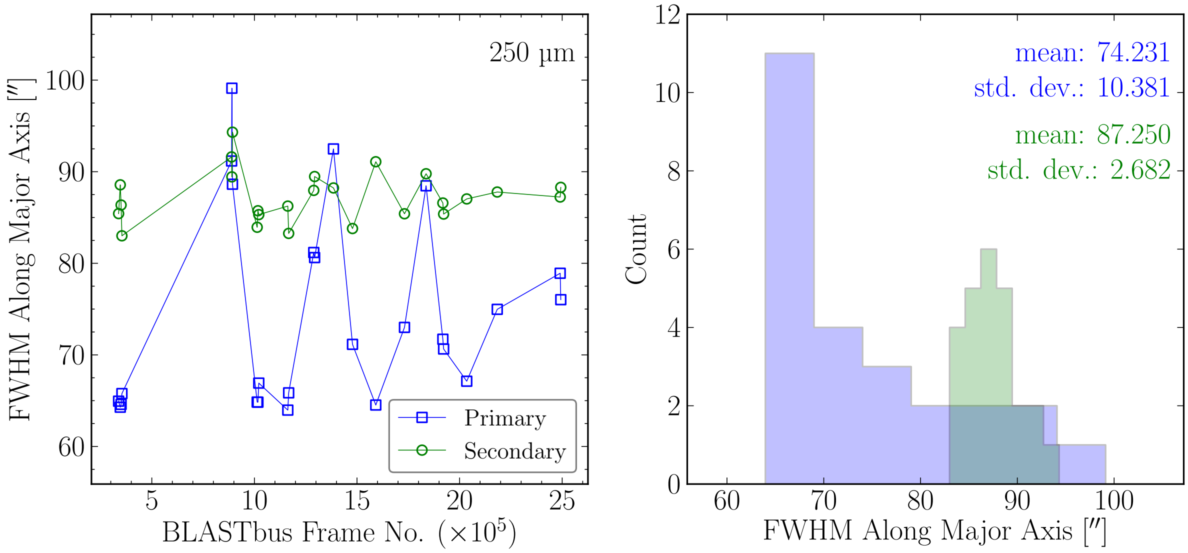


(a) FWHM along the major axes at 250 μm

Figure 4.19: *Left* – FWHMs along the major axes of the primary and secondary Gaussian peaks as a function of time. *Right* – Histograms of the FWHM data.

Figure 4.21 shows the rotation angles of the individual peaks, $\theta_1$ and $\theta_2$, as a function of time. Both angles appear to oscillate, with the changes in the primary being significantly larger. These oscillations are also anti-correlated. As the flight progresses, the beam varies from having peaks comparable in rotation angle, to having a primary peak that is shallower (with its major axis closer to horizontal) and a secondary peak that is steeper, giving the beam a more bent appearance.

The overall picture that emerges from this beam modelling is an oscillation with time between two extremes in beam shape. At one extreme, the beam has more closely-

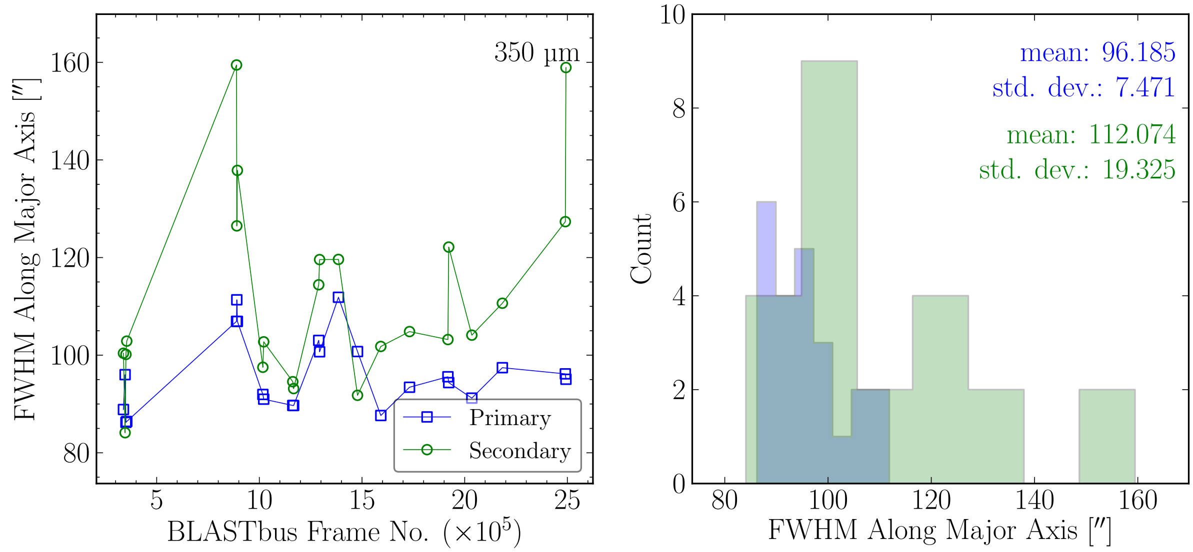


(b) FWHM along the major axes at 350 µm

Figure 4.19: (*Continued*) FWHMs along the major axes of the two Gaussian peaks.

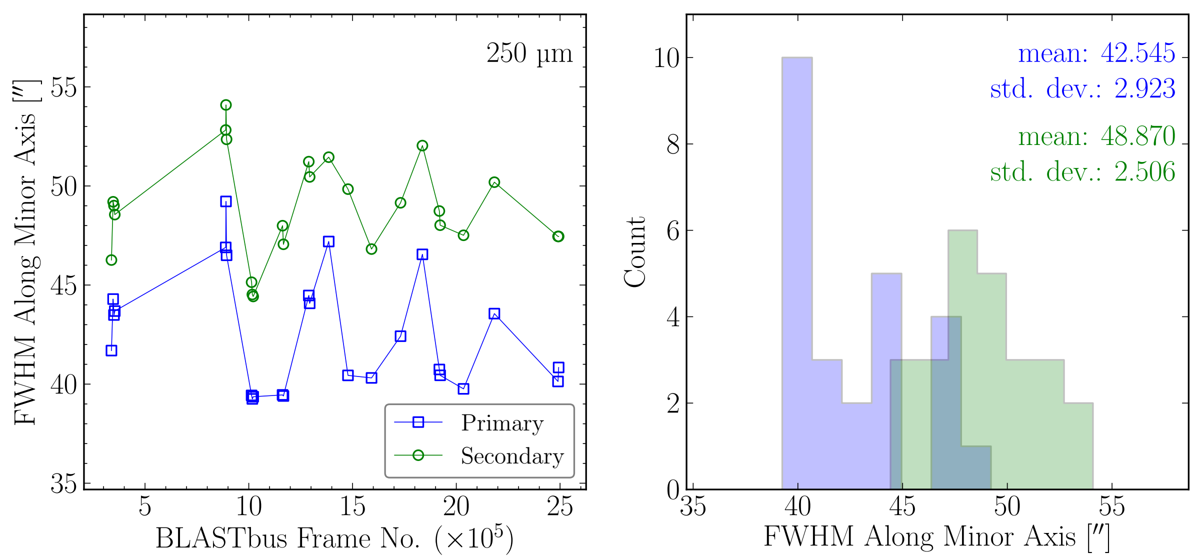


(a) FWHM along the minor axes at 250 µm

Figure 4.20: *Left* – FWHMs along the minor axes of the primary and secondary Gaussian peaks as a function of time. *Right* – Histograms of the FWHM data.

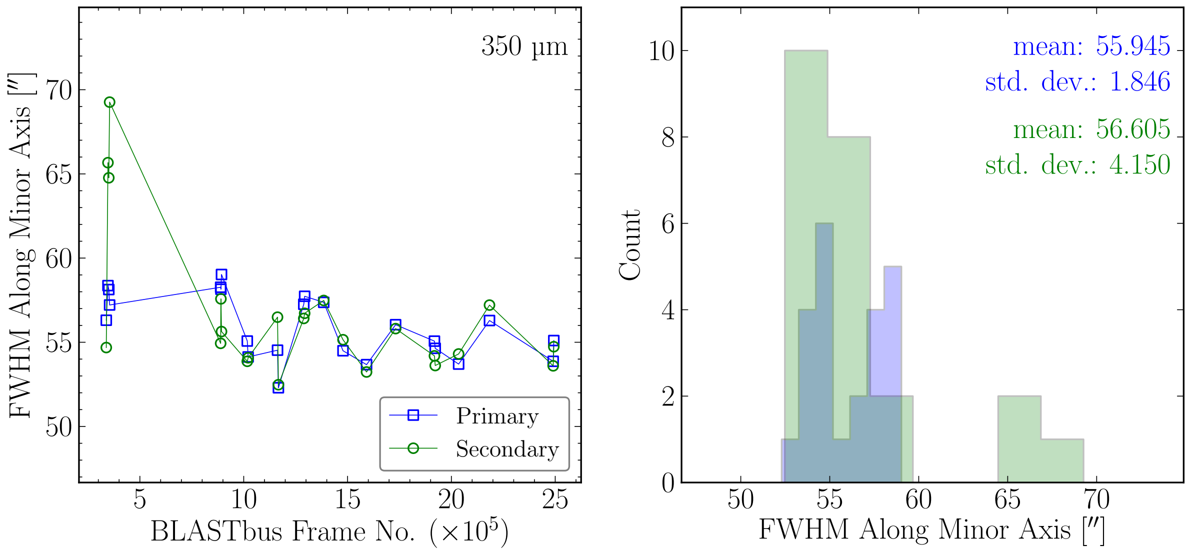


(b) FWHM along the minor axes at 350 µm

Figure 4.20: (*Continued*) FWHMs along the minor axes of the two Gaussian peaks.

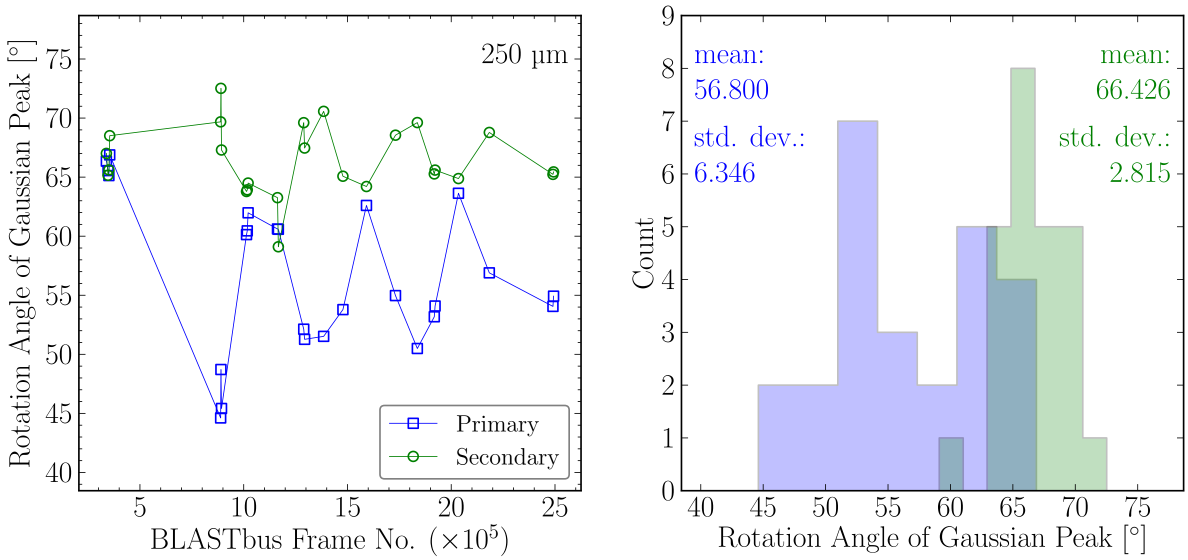


(a) Peak rotation angles at 250 µm

Figure 4.21: *Left* – Rotation angles of the two Gaussian peaks ($\theta_1$ and $\theta_2$) as a function of time. *Right* – Histograms of the peak rotation angle data.

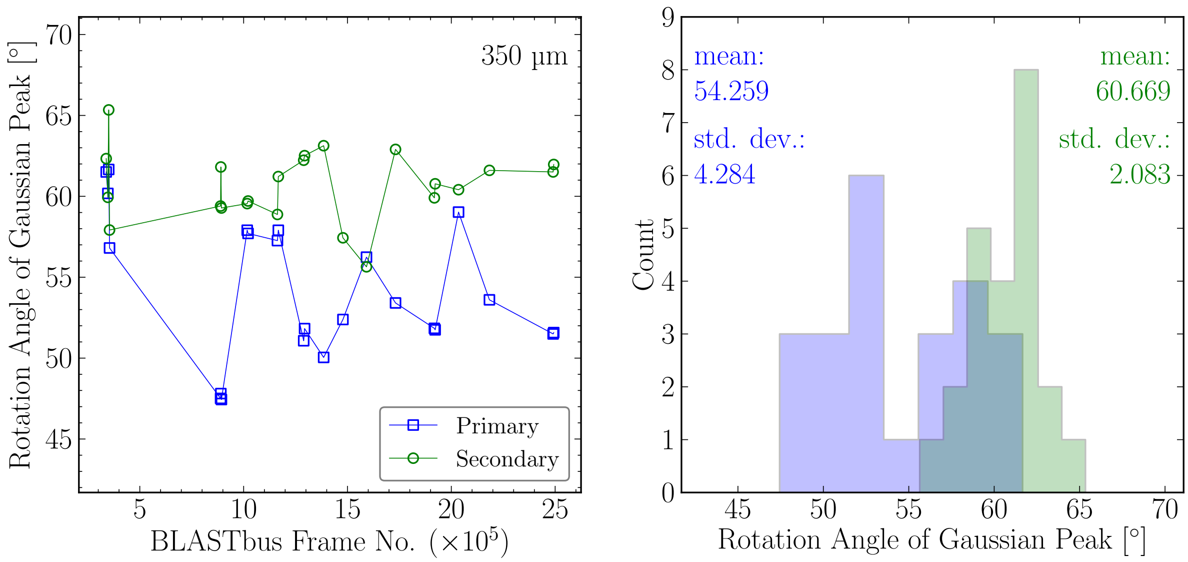


(b) Peak rotation angles at 350 µm

Figure 4.21: (*Continued*) Rotation angles of the two Gaussian peaks vs. time.

separated peaks that are different in amplitude and less elongated. The beam is also steeper overall, and less bent. Examples include the beams mapped at indices 1 021 287 or 1 591 529 shown in Figure 4.15. At the other extreme, the beam has more widely-separated, more elongated peaks with a more comparable amplitude. The beam is also shallower overall, and more bent. Examples include the beams mapped at indices 890 177 or 1 385 263 shown in Figures 4.15. Examining the peak rotation angles, for example (Figure 4.21), the typical period of oscillation is ∼500 000 slow frames, or just over one day. These variations suggest some kind of long-timescale variation in the thermal, mechanical, or other properties of the telescope. However, attempts to correlate the beam fit parameters with other measured quantities such as parallactic angle, telescope elevation angle, primary and secondary mirror temperatures, secondary mirror strut temperatures, and the gradients of these temperatures yield no obvious trends. As of this writing, there is no physical model that explains both the beam shape, and its time-variability.

## 4.4 Beam Centroiding

The Gaussian fitting procedure of the previous sections was also used to determine beam centroids for each bolometer in each focal plane array. This procedure can be divided into three major tasks: defining the beam centroid location (Section 4.4.1), determining the offset between the the star camera pointing solution and the telescope bore-sight (Section 4.4.2), and computing the relative bolometer pointing offsets within each array (Section 4.4.3).

### 4.4.1 Computing the Flux-Weighted Centroid

Due to the asymmetric, dual-lobed structure of the beam, the definition of the beam centroid location was not straightforward. Ultimately, it was decided to use the flux-weighted centroid $(x_c, y_c)$ defined as

$$x_c = \frac{1}{A_{\text{tot}}} \sum_{x=1}^{n} \sum_{y=1}^{m} s(x,y)\,[xA(x,y)] \tag{4.5}$$

$$y_c = \frac{1}{A_{\text{tot}}} \sum_{x=1}^{n} \sum_{y=1}^{m} s(x,y)\,[yA(x,y)]$$

where the map is $m \times n$ pixels, $A(x,y)$ is the map value in pixel $(x,y)$, and $A_{\text{tot}} \equiv \sum_{x=1}^{n}\sum_{y=1}^{m} s(x,y)A(x,y)$ is the total flux in the two beam lobes. The mask $s(x,y)$ excludes pixels below the pixel brightness threshold:

$$s(x,y) = \begin{cases} 1 & \text{if } A(x,y) \geq 0.275 \cdot \max[A(x,y)] \\ 0 & \text{otherwise} \end{cases}$$

Figure 4.22 shows the locations of the beam centroids computed for the all-bolometer, all-data IRAS 08470-4243 maps described in Section 4.1. The selection of the flux-weighted centroid is a compromise given the effects of beam smearing that result from changing parallactic angle. The chosen centroid location is the pivot point around which this beam rotation on the sky will occur. The flux-weighted centroid necessarily lies

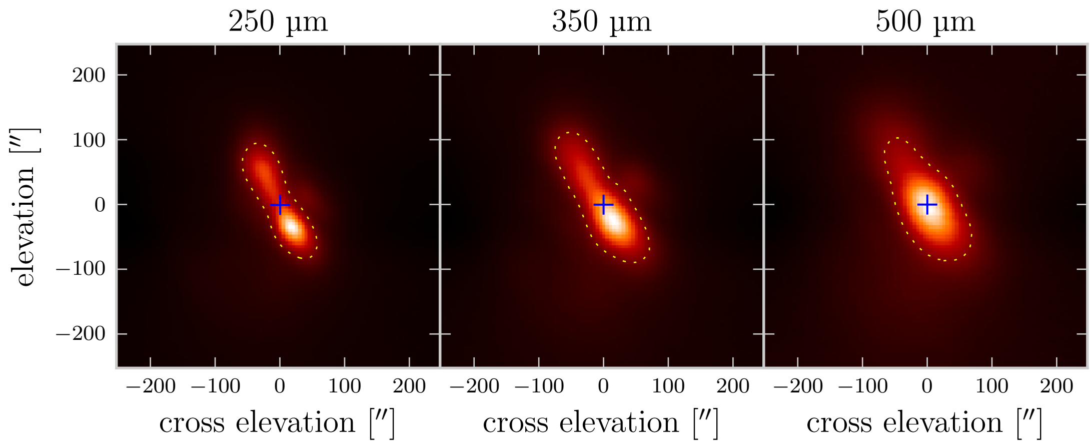


Figure 4.22: Locations of the flux-weighted centroids for the time- and array-averaged beam in each band (blue crosses). The yellow dotted contours show the 27.5% level of the maps, enclosing the pixels that were included in the centroid calculation.

between the two lobes, but closer to the primary, leading to some smearing of both lobes, but not as much as for more extreme cases, such as locating the centroid at one lobe or the other, or at their geometric centre. Table 4.3 gives the offsets between the beam centroids and the brightest pixel locations. If the beam centroid is $(x_c, y_c)$ and the location of the brightest pixel is $(x_{\max}, y_{\max})$, then the second column of this table is $\Delta x \equiv x_c - x_{\max}$ and the third column is $\Delta y \equiv y_c - y_{\max}$.

Table 4.3: Offsets between the location of the brightest map pixel, and the flux-weighted beam centroid in each band

| Band [µm] | Offset in Cross-El ["] | Offset in El ["] |
|---:|---:|---:|
| 250 | −14.51 | +34.23 |
| 350 | −14.97 | +24.68 |
| 500 | +0.16 | −0.16 |

### 4.4.2 The Star Camera to Bore-Sight Pointing Offset

Pointing reconstruction for BLASTPol was provided in part by a pair of imaging star cameras. In *BLASTPol12*, one of the two star cameras experienced solid-state drive (SSD) failure near the beginning of the flight. The remaining camera, known as the integrating star camera, or ISC, lasted for approximately half the flight before experiencing the same type of failure. The next step in the beam centroiding procedure was to determine the pointing offset between the ISC and the telescope bore-sight, in order to produce a telescope bore-sight pointing solution from the ISC one. This star camera offset has two components: a fixed component that is different for each of the three bolometer arrays, and a time-variable component that is common to all three arrays. The time-variable component is necessary due to flexure of the star camera mount and other mechanical changes during the flight.

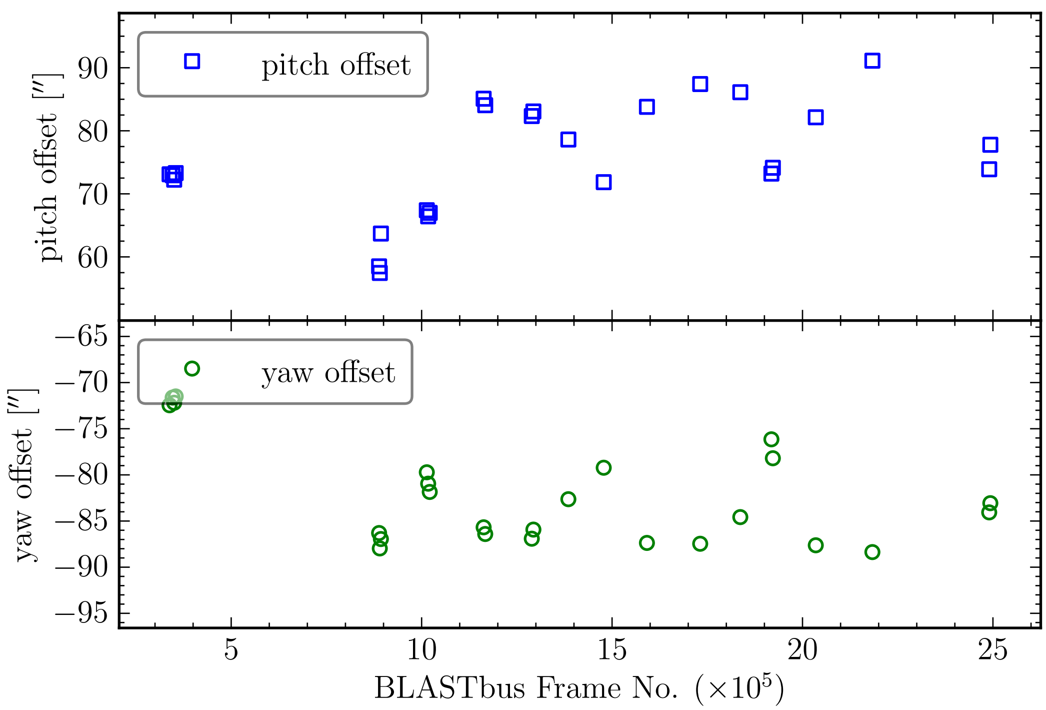


Figure 4.23: Star camera offsets from the telescope bore-sight in pitch and yaw vs. BLASTbus index, for observations of IRAS 08470-4243

Star camera offsets were determined for IRAS 08470-4243 observations by applying the two-Gaussian fitting procedure to the 250 μm all-bolometer maps made for each of the observations listed in Table 4.2. Each of these maps was made to be centred on the known RA and Dec of the source (according to the ISC solution). Once the primary peak location had been identified from the fit, the flux-weighted centroid location was determined using the offsets for 250 μm in Table 4.3. The difference between the coordinates of this centroid and the coordinates of the map centre was taken to be the star camera offset at the time of the observation in question. The fixed, per-array component of the star camera offset was then modified using the offsets for 350 μm and 500 μm in Table 4.3 in order to make the positions of the flux-weighted centroids correct for maps made in these bands as well. Figure 4.23 shows the results for the time-variable component of the star camera offsets at times of IRAS 08470-4243 observations.

### 4.4.3 Relative Bolometer Pointing Offsets

Once the telescope bore-sight pointing solution was known, the final step in the beam centroiding procedure was to determine the relative pointing offsets, from the bore-sight, of individual bolometers within each array. These offsets are considered to be constant in time, and are tabulated along with other detector parameters in a configuration file known as the *bolotable*. BLASTPol map making software such as naivemap, naivepol, and TOAST [99] all make use of the bolotable. Once reliable star camera offsets were available, new bolotable offsets for *BLASTPol12* were computed by applying the two-Gaussian fit model to single-bolometer maps of IRAS 08470-4243. Naive maps were made centred on the known RA and Dec of the source (according to the bore-sight solution), and the difference between the map centre and the flux-weighted beam centroid was taken to be the bolotable pointing offset for the bolometer in question. Figure 4.24 plots the 350 μm map centre to beam centroid offsets in pitch, versus the map centre to beam centroid offsets in yaw, before and after the procedure above. The corrected offsets are also shown in the inset. This result demonstrates that the new bolotable offsets locate the beams for each bolometer correctly, to a high degree of accuracy.

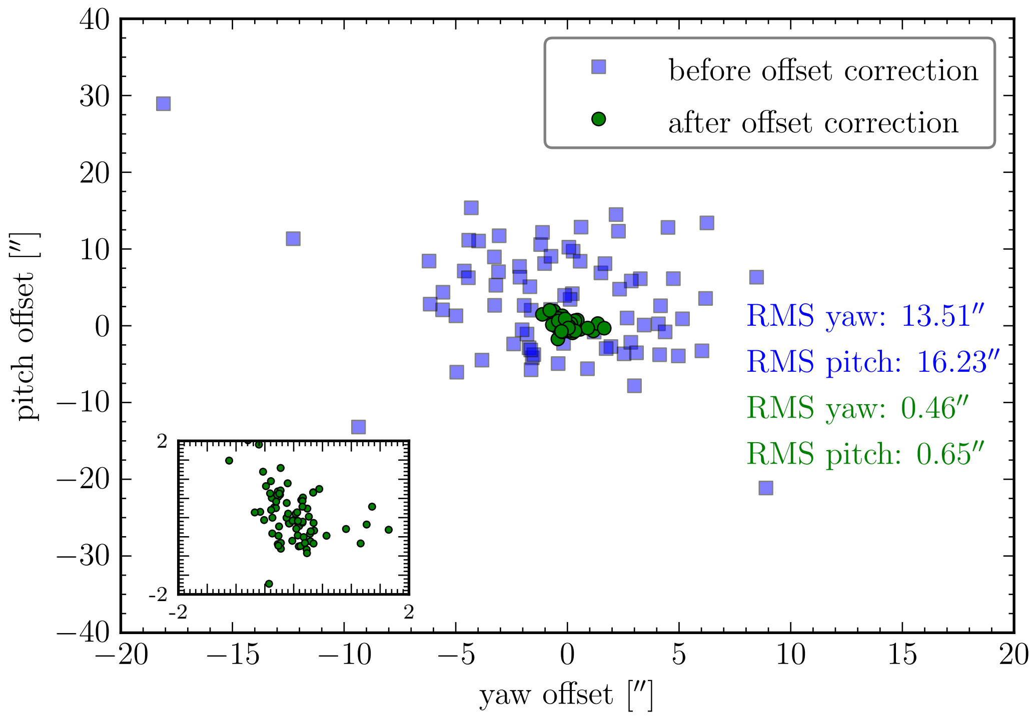


Figure 4.24: Pointing offsets between the map centres and the beam centroids for single-bolometer maps at 350 µm, before and after correction of the bolotable offsets

## 4.5 Bolometer Flat-Fielding

Beam modelling also found use in the application of bolometer flat-fielding. The BLAST-Pol optics box included a calibration lamp placed in a hole at the centre of the Lyot stop. Throughout the flight, to calibrate detector gains, the arrays were illuminated by calibration pulses of constant intensity. Before flat-fielding, the bolometer timestreams had already been corrected for non-linearity (gain variation with DC level) using these cal pulse amplitudes. This ensured a constant gain throughout the flight for a given bolometer. However, different bolometers still had different gains. The goal of flat-fielding is to apply correction coefficients to each bolometer to remove these differences, producing a "flat" response to a given source across the arrays. This correction can be achieved by comparing single-bolometer maps made from observations of the same source, and quantifying the total power in each map in some way. Two methods for doing so were attempted: fit-based flat-fielding and aperture photometry-based flat-fielding.

### 4.5.1 Fit-Based Flat-Fielding

One method of flat-fielding was to apply the two-Gaussian fit model to the single-bolometer $50 \times 50$ pixel Stokes $I$ maps of IRAS 08470-4243 described previously. The integral over the fit model was then computed as a measure of the total beam power:

$$I_{\mathrm{tot}} = \iint F(x,y)\,dx\,dy. \tag{4.6}$$

The integral was computed using the `scipy.integrate.dblquad` function with limits corresponding to the map width and height. The function $F(x, y)$ is the two-Gaussian version of the fit model given by Equations 4.3 and 4.4. In each array, a reference bolometer was chosen near the centre of the focal plane: B2E08H, B3E09V, and B5C05V. The ratio of the integrated fluxes of each bolometer to that of the reference bolometer, $I_{\mathrm{tot}}/I_{\mathrm{ref}}$, was computed to determine the relative sensitivity of each bolometer. During the fitting procedure, a pixel brightness threshold of 30% of the peak value was applied to the maps. This value was found to be suitable for the single-bolometer maps at $10''$ resolution. Background subtraction was also attempted by computing the mean map intensity in two equal-width $100''$ vertical strips at the left and right edges of the map, and then subtracting this value from the maps before fitting. Strips were chosen, rather than a ring, to avoid the diffuse structure of IRAS 08470-4243, which is distributed asymmetrically around the map centre. This subtraction made little difference for most of the maps, since the background level computed using this method was negligibly small. Example results of the fit-based flat-fielding calculations are shown in Figures 4.25a, 4.26a, and 4.27a.

### 4.5.2 Aperture Photometry-Based Flat-Fielding

The second flat-fielding method computed the integrated flux directly from the maps by adding up all pixel values within a circle of radius 150″ centred on the map centre. The map background was estimated from an annulus with radii between 150″ and 180″, and subtracted. These values were taken to be the integrated flux $I_{\mathrm{tot}}$, and they were normalized to $I_{\mathrm{ref}}$ using the reference bolometers listed in the previous section. Example results of the photometry-based flat-fielding calculations are shown in Figures 4.25b, 4.26b, and 4.27b. In these figures, for a given bolometer array, both the fit-based and photometry-based flat-field coefficient visualizations are plotted on the same colour scale, for direct comparison.

In the case of both flat-fielding methods, all of the coefficients shown were generated using single-bolometer maps that combined all of the data ranges in Table 4.2 for which the HWP position was approximately 16°. Corresponding maps for each bolometer were made for data at each of the other (approximate) HWP positions of 38°, 63°, and 84°. The data were binned by HWP angle to see if the normalized coefficient for a given bolometer varied between observations taken at different HWP angles. No variation would be expected for an unpolarized source, however IRAS 08470-4243 is known to be polarized at the level of a few percent. For both flat-fielding methods, variation in the coefficients of individual bolometers was observed that was consistent with this level of polarization.

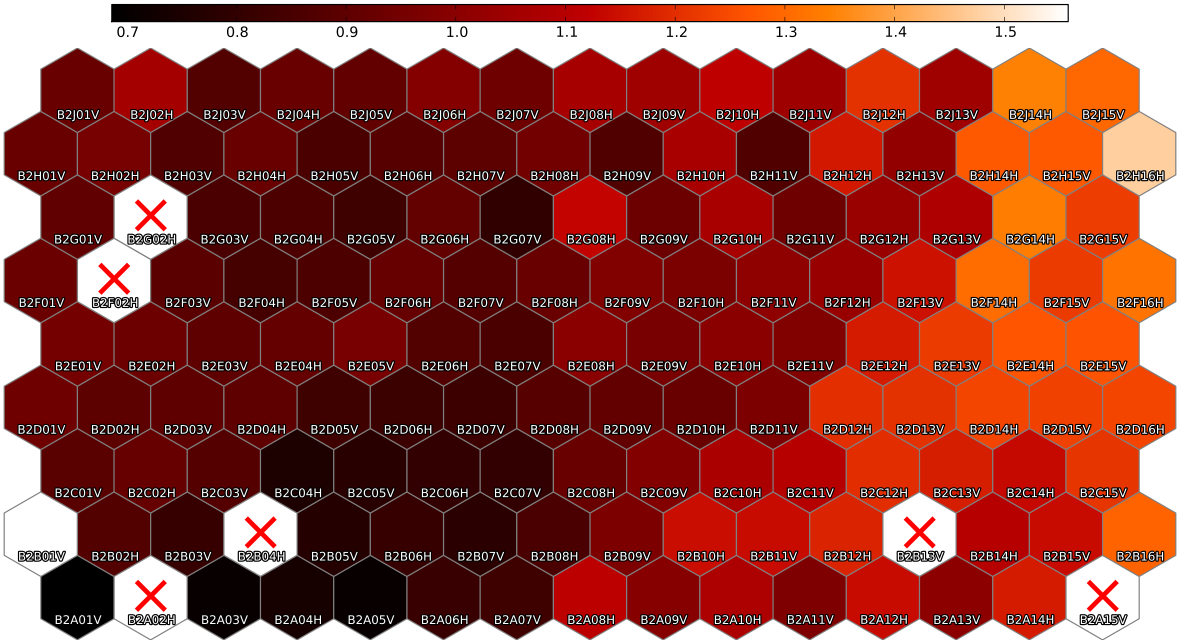


(a) The fit-based flat-fielding coefficients

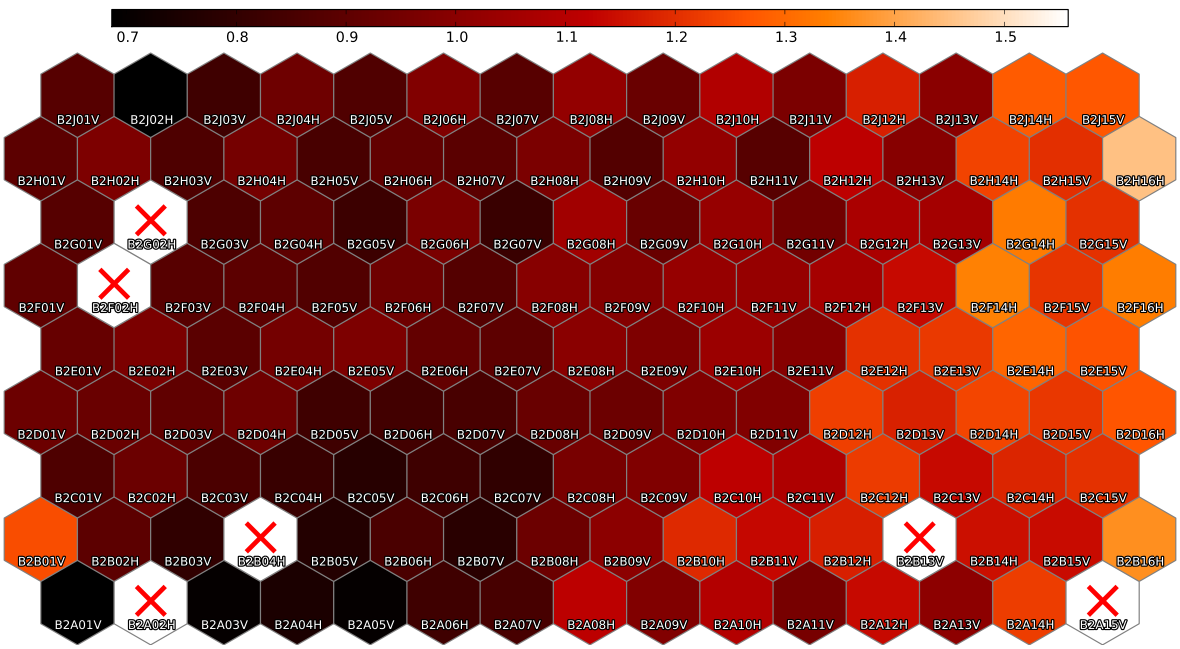


(b) The photometry-based flat-fielding coefficients

Figure 4.25: Flat-fielding coefficients at 250 µm

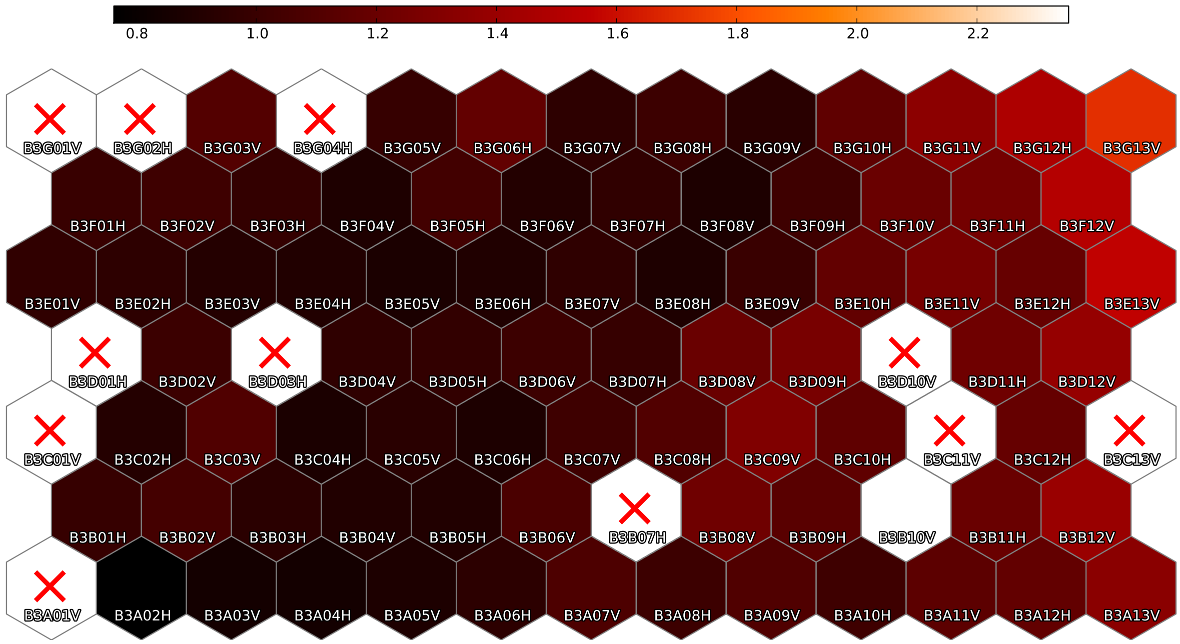


(a) The fit-based flat-fielding coefficients

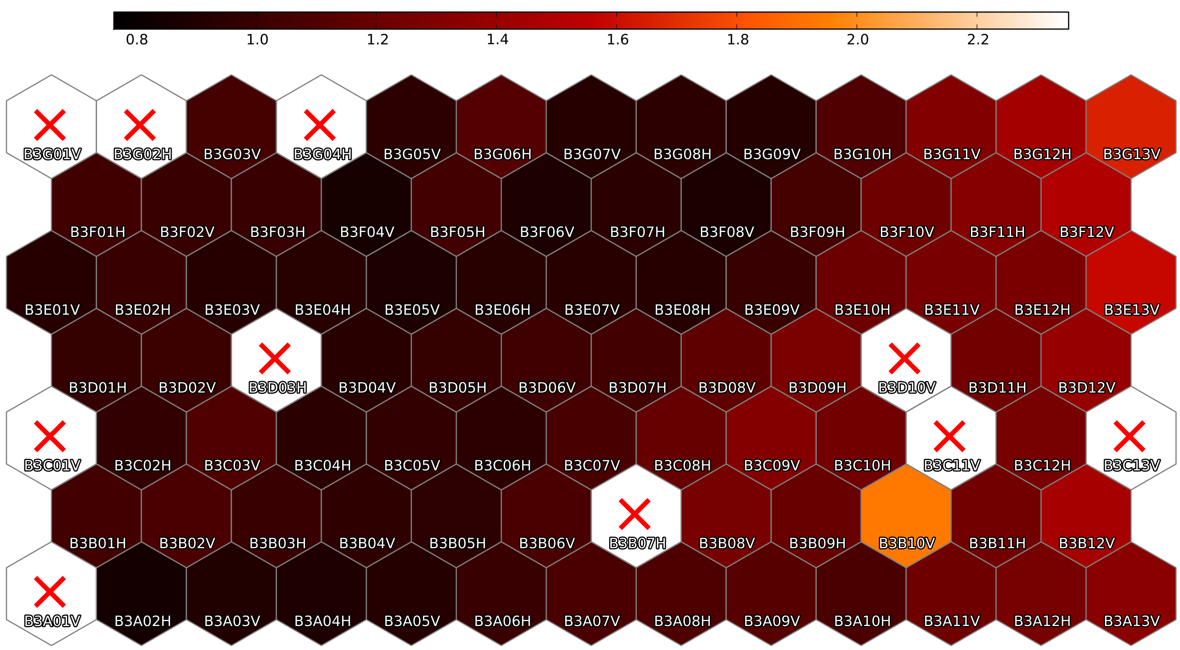


(b) The photometry-based flat-fielding coefficients

Figure 4.26: Flat-fielding coefficients at 350 µm

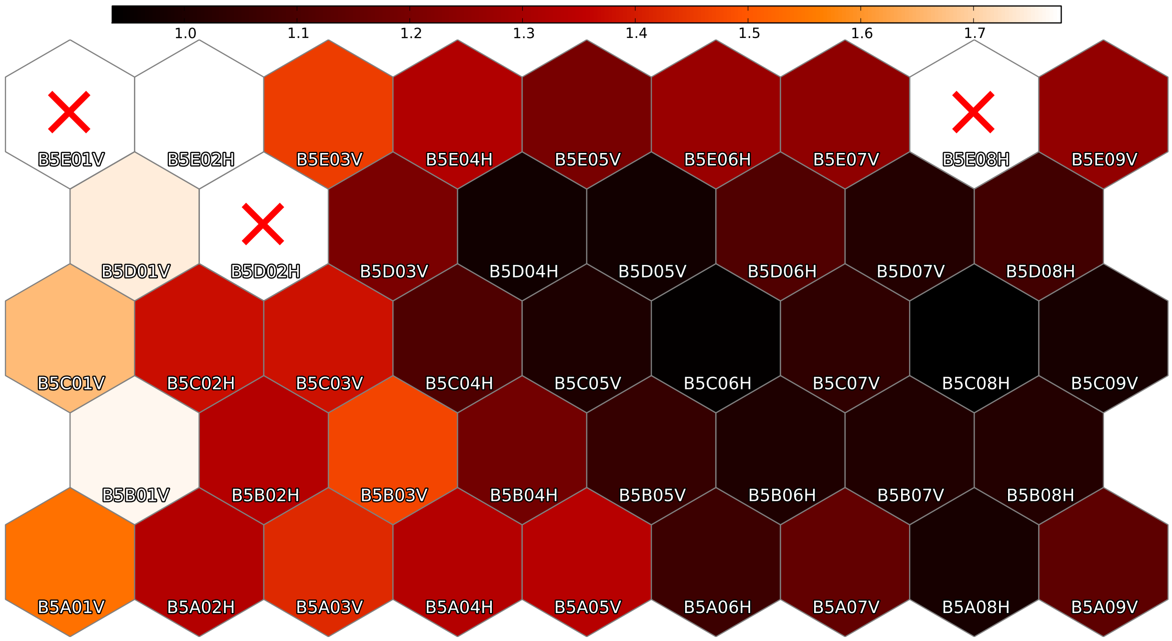


(a) The fit-based flat-fielding coefficients

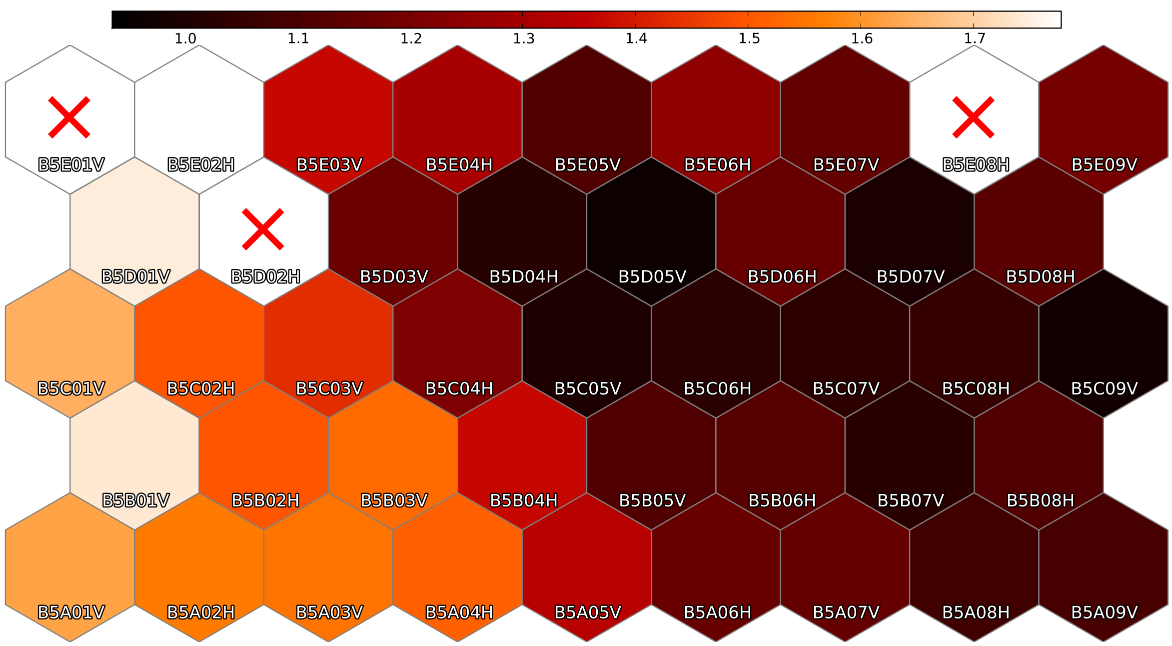


(b) The photometry-based flat-fielding coefficients

Figure 4.27: Flat-fielding coefficients at 500 µm

### 4.5.3 A Comparison of the Two Methods

As shown in Figures 4.25 through 4.27, while the two flat-fielding methods produce different results at the level of individual bolometers, the overall gradients in sensitivity across each of the arrays appear the same for both. Each method had advantages and disadvantages. The aperture photometry-based method, which relied directly on the original map pixels, was more sensitive to coverage gaps in the maps. The fit-based method could be inconsistent: the least squares algorithm would produce a bad fit for some of the bolometers in an array, producing an inaccurate result for the integrated power. The beam maps for the initial analysis were made with subsets of the data binned by HWP angle. However, for the final flat-fielding, it was deemed sufficient to use beam maps combining *all* of the data listed in Table 4.2, which would not have coverage gaps. Therefore, the aperture photometry method was chosen to produce final flat-field coefficients for the analysis. To flat-field the arrays, each bolometer timestream was scaled by the reciprocal of its normalized coefficient: $I_{\rm ref}/I_{\rm tot}$. Figures 4.28, 4.29, and 4.30 show percent differences between the photometry-based and fit-based coefficients for the data set at HWP position 16°. These differences are computed as

$$100 \cdot \frac{(I_{\rm tot}/I_{\rm ref})_{\rm phot} - (I_{\rm tot}/I_{\rm ref})_{\rm fit}}{(I_{\rm tot}/I_{\rm ref})_{\rm fit}}.$$

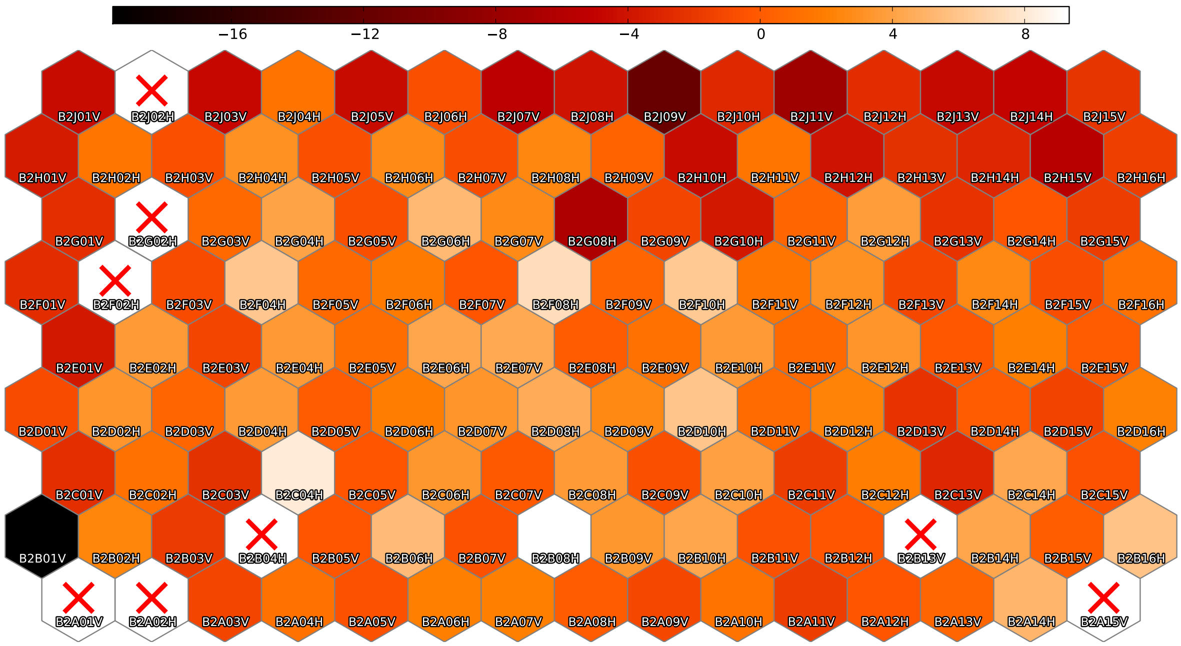


Figure 4.28: Percent differences between photometry-based and fit-based flat-fielding coefficients at 250 µm

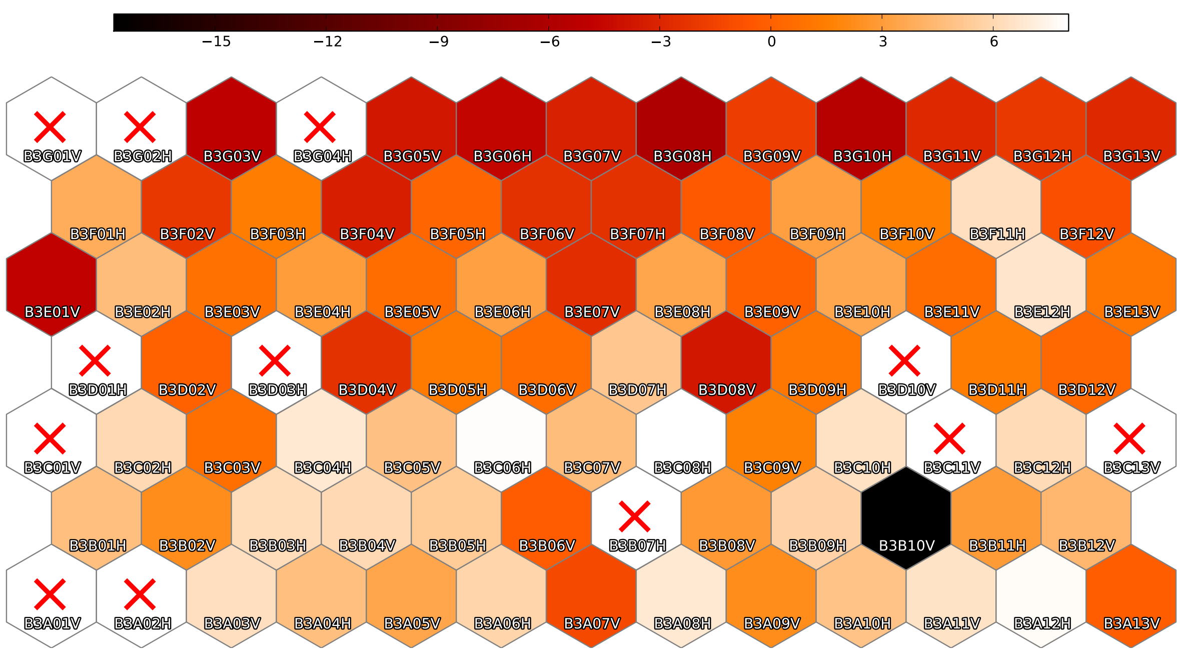


Figure 4.29: Percent differences between photometry-based and fit-based flat-fielding coefficients at 350 µm


Figure 4.30: Percent differences between photometry-based and fit-based flat-fielding coefficients at 500 μm

# Chapter 5

# Polarization Spectrum of the Carina Nebula

## 5.1 The Carina Nebula: Overview

The Carina Nebula (NGC 3372) is the largest and highest surface brightness nebula in the southern sky, appearing in visible light as a giant H II region spanning several square degrees. It is located at an estimated distance of 2.3 kpc [100, 101], within the Sagittarius-Carina spiral arm of the Milky Way. The nebula and surrounding molecular cloud are part of a Giant Molecular Cloud (GMC) complex that spans some 150 pc [102].

Early study of this region focused on the variable star Eta ($\eta$) Carinae, which is among the most massive and luminous stars in the Galaxy. This star is famed for its "Great Eruption", which took place in the mid-19$^{\text{th}}$ century. During this outburst event, the star brightened to an apparent visual magnitude of $-1.0$, before dimming again as it became enclosed in a dusty shroud of ejected material that has become known as the Homunculus Nebula.

In the 20$^{\text{th}}$ century, observations began focusing more on $\eta$ Carinae's surrounding environment: the Carina Nebula itself. Initial observations found no luminous embedded star formation regions in the central part of the nebula, leading to the conclusion that the nebula was an evolved H II region with no active star formation. However, in the last two decades, thanks to observations at many different wavelengths, a new and very different picture has emerged. Active star formation is occurring at the periphery of the nebula, triggered by feedback in the form of intense ultraviolet (UV) radiation and stellar winds from dozens of O-type stars located in the nebula's central clusters. At the same time, this feedback is disrupting the molecular cloud from which these massive stars formed, removing the raw material needed for further star formation in the developing OB association. The observations to date that have led to this overall story are described extensively in [102]. It is noted there that Carina is an ideal laboratory in which to study these effects of feedback from massive stars, in an environment that has not yet been disrupted by supernovae.

In the context of BLASTPol, the Carina Nebula is a much more active and evolved region than the other targets that were observed: relatively quiescent molecular clouds. Therefore, polarization measurements of Carina in the submillimetre have the potential to reveal what the physical properties of the dust are like in a very different radiative environment from that found in these other sources.

## 5.2 Regions in Carina Observed by BLASTPol

The BLASTPol maps of the Carina Nebula encompass two regions, in particular, that have been identified as areas of active star formation. Figure 5.1 shows a filled contour map of Stokes $I$ at 350 µm from the BLASTPol data. On this map, two 8 µm images from the Infrared Array Camera (IRAC) on the *Spitzer Space Telescope* have been overlaid[1]. The leftmost (easternmost) image is the *Spitzer* image of the "South Pillars" region. The rightmost (westernmost) image is of the "West" region. In between the two images, in

[1]These data were obtained from the *Spitzer* Heritage Archive at http://sha.ipac.caltech.edu.

the central part of the nebula, are two clusters of OB stars: Trumpler 14 and Trumpler 16, the latter of which includes $\eta$ Carinae. The dust pillars in the South Pillars region are thought to have been sculpted by radiation and winds from the massive stars in these central clusters. An analysis of the *Spitzer* data has shown that the elongation axes of many (but not all) of these structures point towards the central region, and has identified numerous young stellar objects in both the South Pillars and West regions [103].

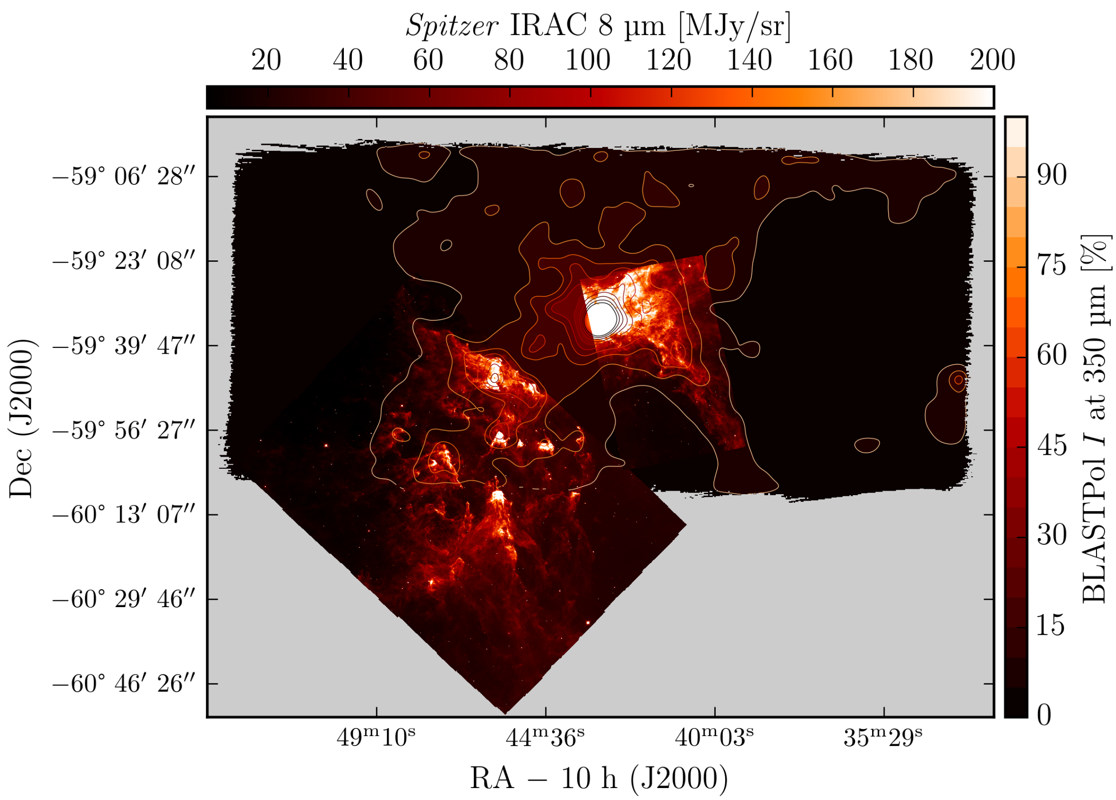


Figure 5.1: BLASTPol 350 μm $I$ map of the Carina Nebula (filled contours), overlaid with *Spitzer* IRAC 8 μm maps made in the same area (images)

## 5.3 Previous Polarimetric Observations of Carina

The Submillimeter Polarimeter for Antarctic Remote Observations (SPARO) measured the polarization at 450 μm in four GMCs, including the Carina Nebula [104]. This survey found that the magnetic field directions in the GMCs were usually parallel to the Galactic plane, and were fairly coherent across the entirety of the clouds. This coherence on relatively large scales has been found by others, and is apparent in the BLASTPol data. The SPARO instrument at the South Pole consisted of a $3 \times 3$ array of detectors that was chopped between a source and two reference positions on the sky. A rotating half-wave plate (HWP) was stepped between six different polarization angles over the course of the observations. Figure 5.2 shows the polarization pseudo-vectors[2] (taken from Table 2 of [104]) measured by SPARO in the Carina Nebula, overlaid on a BLASTPol contour map of the emission at 350 µm, with contours from 5% to 50% in 5% increments (and also a 75% contour). Figure 5.3 shows the corresponding **B**-field directions inferred from these pseudo-vectors.

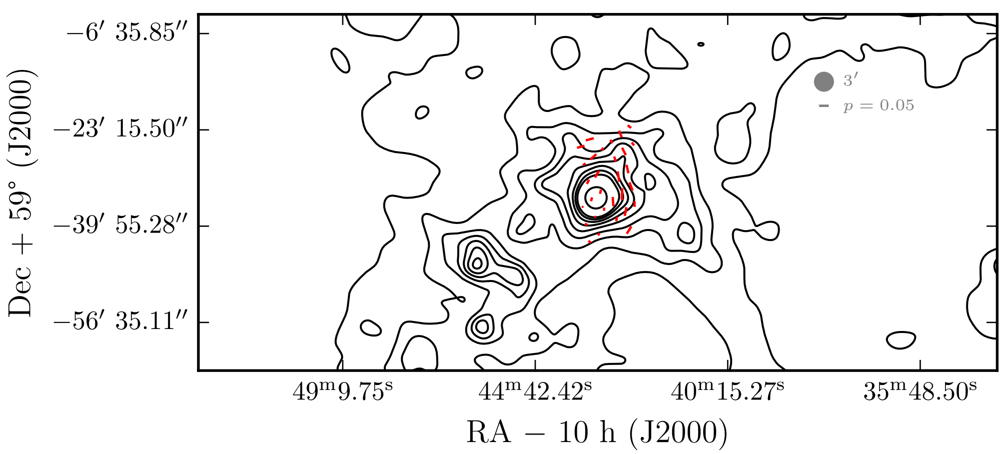


Figure 5.2: Directions and magnitudes of the fractional polarization $p$ in the Carina Nebula, as measured by SPARO at 450 µm (red). The diameter of the grey circle shows the size of the 3′ beam in the deconvolved BLASTPol map (see Section 5.4). The grey rod length is for 5% polarization. The contours show the BLASTPol $I$ map at 350 µm.

[2]The line segments used to depict measured light polarization orientations on the sky are referred to as pseudo-vectors herein. They are not true **E**-field vectors, since the direction of **E** is measured with an ambiguity of 180°.

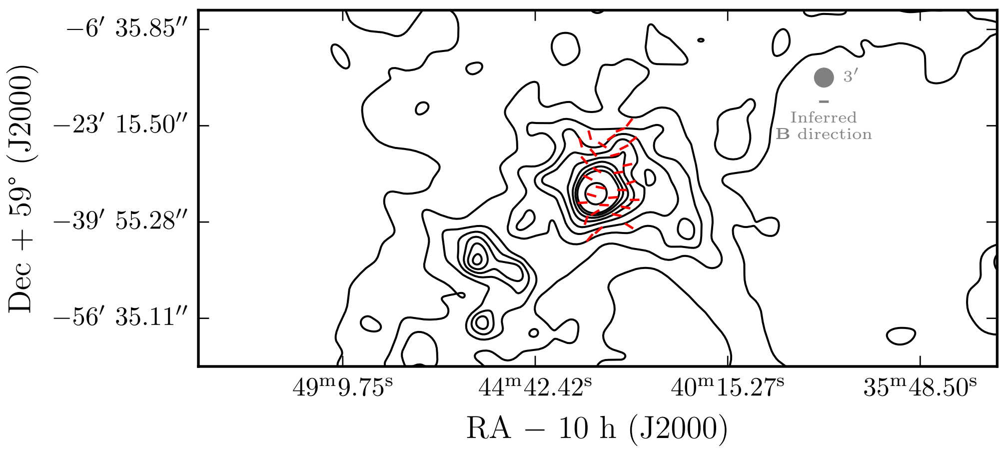


Figure 5.3: Inferred directions of the projected **B**-field in the Carina Nebula, as measured by SPARO at 450 µm (red). The pseudo-vectors have all been drawn the same length, rather than being scaled by $p$, to show the directions more clearly. The contours show the BLASTPol $I$ map at 350 µm.

These figures can be compared with Figures 5.13 and 5.14 (Section 5.6), which show the corresponding plots for the BLASTPol measurement. The inferred **B**-field directions are qualitatively similar in the same location. In [63], a (pseudo-) vector-by-vector comparision is done, showing that there is actually remarkably close agreement between the polarization directions in the locations where they were measured by both experiments.

## 5.4 Map Deconvolution

The BLASTPol maps have been convolved with the multi-lobed, asymmetric PSF described in the previous chapter. Two different methods to deconvolve that PSF from the maps were attempted, with the goal of producing a symmetric Gaussian beam that was the same in each band.

### 5.4.1 Fourier-Space Deconvolution

The first attempt at deconvolution was done in Fourier space. The problem can be stated as follows. The desired symmetric Gaussian beam $s(x, y)$ can be thought of as resulting from the convolution of the asymmetric BLASTPol PSF $b(x, y)$ with a kernel $h(x, y)$:

$$s(x, y) = b(x, y) * h(x, y) \tag{5.1}$$

The goal was to determine the $h(x, y)$ that would produce an output Gaussian beam of a given FWHM. From the convolution theorem for Fourier transforms:

$$\mathscr{F}\{s(x, y)\} = \mathscr{F}\{b(x, y)\} \cdot \mathscr{F}\{h(x, y)\} \tag{5.2}$$

where $\mathscr{F}$ is the Fourier transform operator. Since the convolution is just a product in Fourier space, the required smoothing kernel can be computed by dividing the BLASTPol PSF from the symmetric beam in Fourier space:

$$h(x, y) = \mathscr{F}^{-1}\left\{\frac{\mathscr{F}\{s(x, y)\}}{\mathscr{F}\{b(x, y)\}}\right\} \tag{5.3}$$

In practice, the inverse transformation of Eq. 5.3 need not be carried out. A given input map $M(x, y)$ can be convolved with the kernel in Fourier space through multiplication, and then inverse transformed to produced a deconvolved map $M_d(x, y)$

$$M_d(x, y) = \mathscr{F}^{-1}\{\mathscr{F}\{M(x, y)\} \cdot \mathscr{F}\{h(x, y)\}\} \tag{5.4}$$

The 2D discrete Fourier transform (DFT, Eqs. 5.5 & 5.6) of the beam and source maps was computed using the fast Fourier transform (FFT) algorithm. The specific implementations used were the `scipy.fftpack.fft2` and `scipy.fftpack.ifft2` routines.

$$\tilde{M}(k_x, k_y) = \mathscr{F}\{M(x, y)\} = \sum_{x=0}^{N_x-1} \sum_{y=0}^{N_y-1} M(x, y) e^{-i\left[k_x\left(\frac{2\pi}{N_x}\right)x + k_y\left(\frac{2\pi}{N_y}\right)y\right]} \tag{5.5}$$

$$M(x, y) = \mathscr{F}^{-1}\{\tilde{M}(k_x, k_y)\} = \frac{1}{N_x N_y} \sum_{k_x=0}^{N_x-1} \sum_{k_y=0}^{N_y-1} \tilde{M}(k_x, k_y) e^{i\left[k_x\left(\frac{2\pi}{N_x}\right)x + k_y\left(\frac{2\pi}{N_y}\right)y\right]} \tag{5.6}$$

The input map has dimensions of $N_y \times N_x$, in pixels. Therefore, the coordinates $(k_x, k_y)$ in the DFTs of the maps correspond to angular spatial frequencies of $\mu_x = (2\pi k_x/N_x)$ and $\mu_y = (2\pi k_y/N_y)$.

The beam models shown in Chapter 4 are in telescope coordinates. The PSF in the Carina Nebula map is rotated and smeared due to the changing parallactic angle of the beam on the sky during the various Carina observations. Attempting deconvolution required producing a parallactic-angle-averaged PSF model $b(x, y)$, weighted by integration time:

$$b(x,y) = \frac{1}{N_f} \sum_{i=1}^{N_{\rm obs}} n_i b_i(x, y, \phi_i) \tag{5.7}$$

In this equation, $N_{\rm obs}$ is the number of separate Carina Nebula observations. The $i^{\rm th}$ Carina Nebula observation has $n_i$ data samples (BLASTbus frames), and $N_f = \sum_{i=1}^{N_{\rm obs}} n_i$ is the total number of data frames for Carina. The model $b_i(x, y, \phi_i)$ is the beam model from Chapter 4, rotated by the parallactic angle $\phi_i$ of the beam during the $i^{\rm th}$ observation, as computed by the pointing solution.

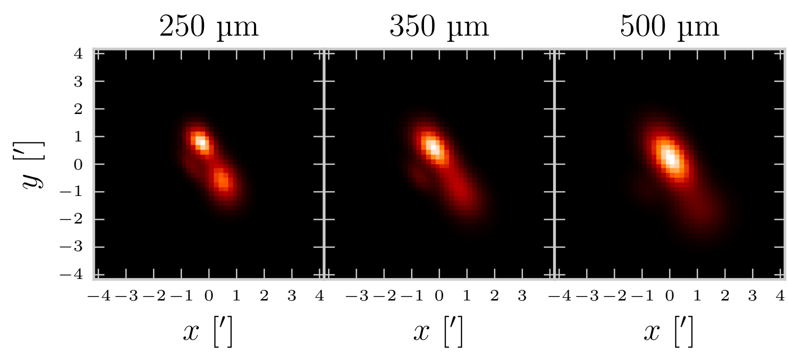


Figure 5.4: BLASTPol PSFs in each band at 10″ resolution, averaged over parallactic angle for the Carina Nebula observations

The resulting averaged beams in each band are shown in Figure 5.4. They appear reasonably similar to point-like sources in the Carina maps in terms of orientation, peak shape, and relative peak intensity. For deconvolution, some improvements were made over the beam modelling described in Chapter 4. A hybrid beam model was used, in which the FWHMs were derived from fitting to Saturn maps, while all other beam model parameters listed in Table 4.1 came from fitting to IRAS 08470-4243 maps. A tertiary peak was added to the 500 µm fit model, which was possible provided its aspect ratio was contrained to be the same as the secondary and primary peaks.

Symmetric 2D Gaussian beam templates were produced with sizes of 2′, 1.5′, 1.25′, and 1′ FWHM. These beams and the BLASTPol PSFs were Fourier transformed, and the convolution kernel was computed (Eq. 5.3). Example results of these operations for 2′ at 350 µm are in Figures 5.5 through 5.7.

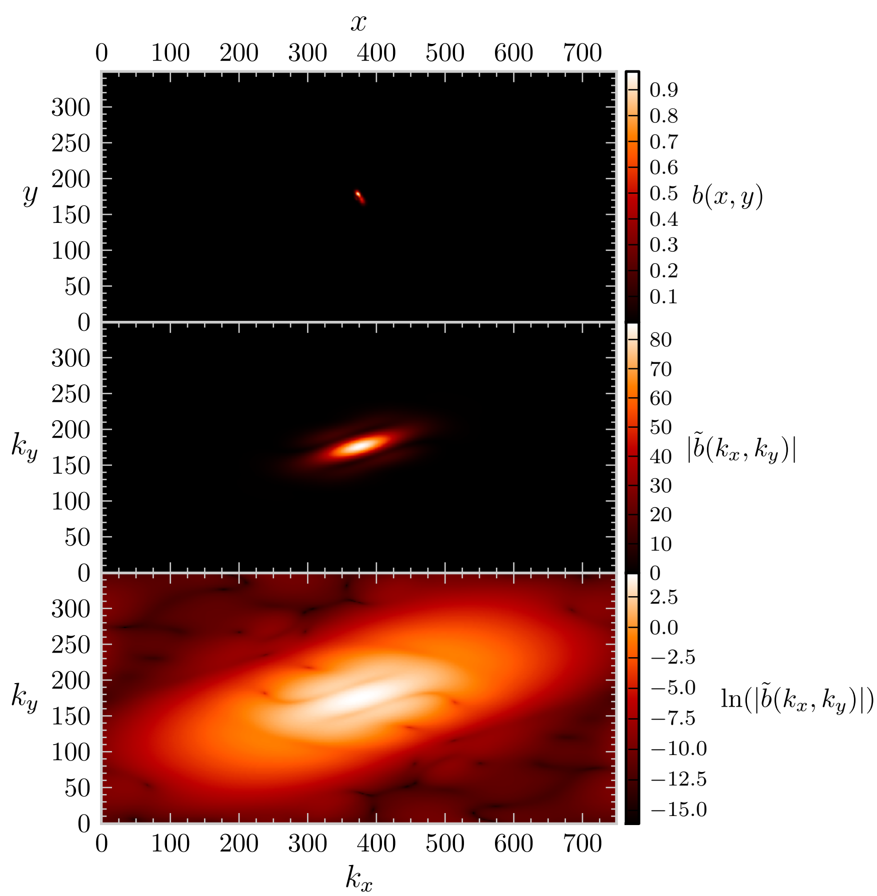


Figure 5.5: *Top* – The BLASTPol PSF for the Carina map at 350 µm. *Middle* – The modulus of the DFT of the PSF on a linear scale. *Bottom* – The modulus of the DFT of the PSF on a log scale.

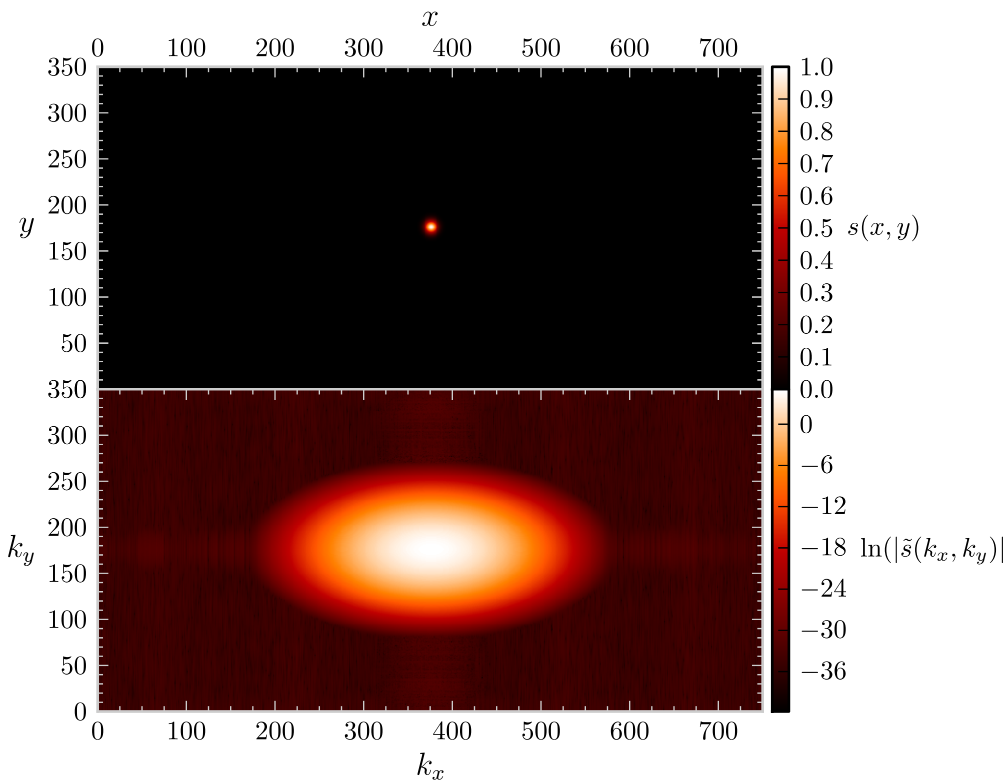


Figure 5.6: *Top* – The symmetric 2′ FWHM Gaussian beam. *Bottom* – The modulus of the DFT of the symmetric Gaussian beam on a log scale.

The deconvolved maps obtained using the kernel (Eq. 5.4) showed some rippling and other image artifacts, with increasing severity as the FWHM of the symmetric Gaussian decreased. These could be reduced by removing certain modes from the kernel. In particular, dark features in the DFT of the BLASTPol PSF (Figure 5.5, bottom), become bright features in the resulting deconvolution kernel, due to the division. To suppress these features, pixels $(k_x, k_y)$ for which $|\tilde{b}(k_x, k_y)| < 0.05 \cdot \max[|\tilde{b}(k_x, k_y)|]$ were set to zero in $\tilde{h}(k_x, k_y)$.

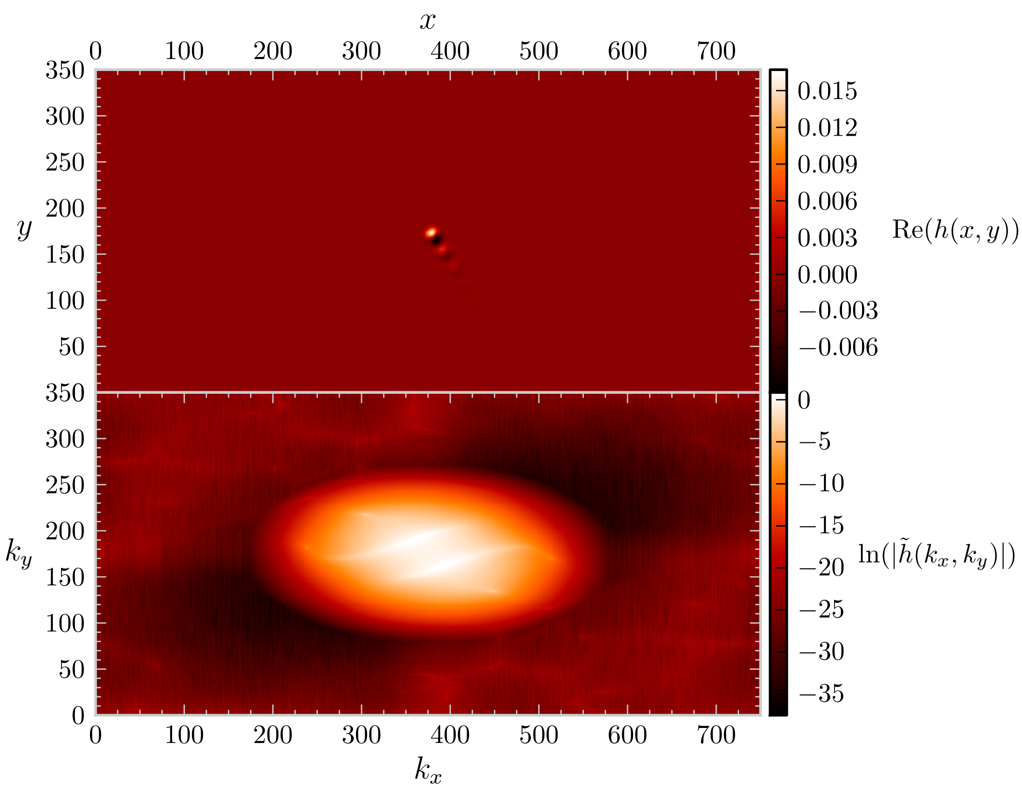


Figure 5.7: *Top* – The real part of the inverse DFT of the Fourier deconvolution kernel for the Carina map. *Bottom* – The modulus of the Fourier deconvolution kernel on a log scale.

Figure 5.8 shows the results of the Fourier deconvolution of the Carina Nebula $I$ map at 350 µm. The technique succeeds in replacing the multi-lobed PSF with a symmetric Gaussian beam, and in recovering information at scales below the full extent of the elongated PSF. However, even with modes in the kernel zeroed, significant ripples and ringing around sources is apparent. These effects are more pronounced in the $Q$ and $U$ maps, and are also worse at 250 µm. This result demonstrates the limitations of the Fourier deconvolution method. A deconvolved map with a symmetric beam, and without significant artifacts, can only be produced at a resolution of $\sim 3'$ FWHM.

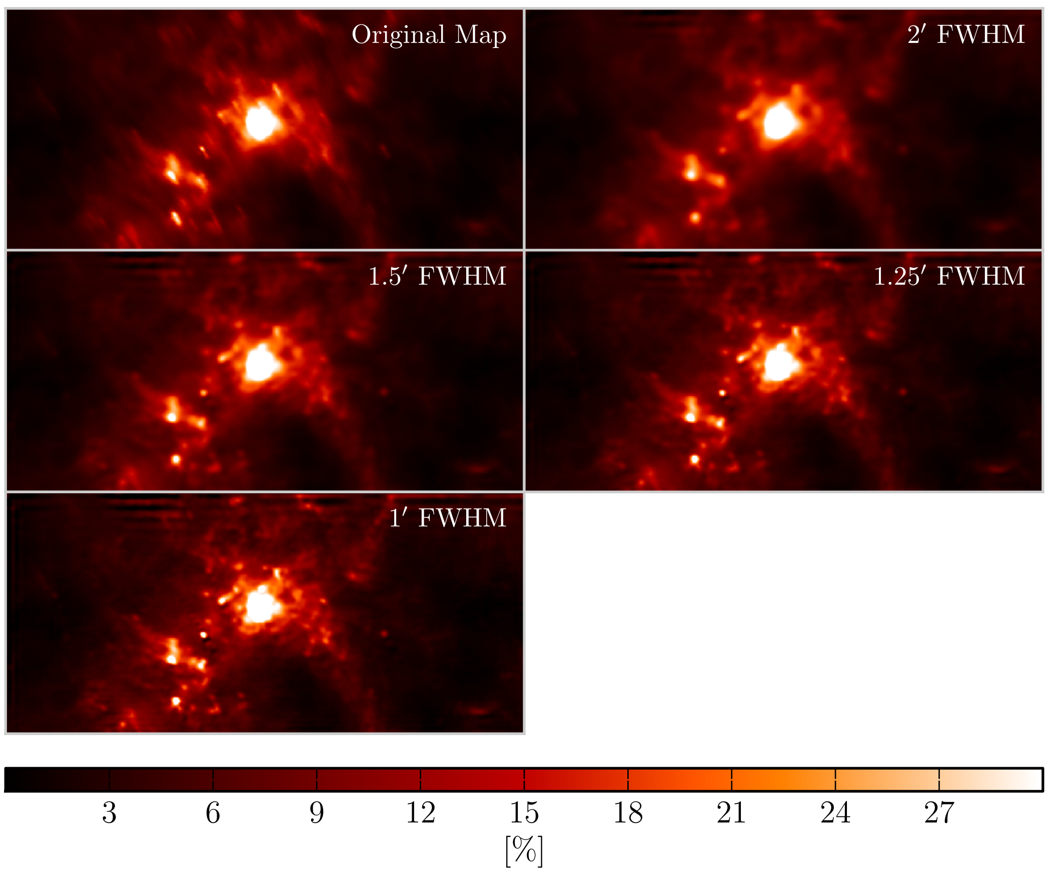


Figure 5.8: Carina Nebula $I$ maps at 350 µm Fourier-deconvolved to $2'$, $1.5'$, $1.25'$, and $1'$ FWHM resolutions. The colour scale shows the intensity as a percentage of the full dynamic range in the original $I$ map.

### 5.4.2 Lucy–Richardson Deconvolution

The second deconvolution method attempted was Lucy–Richardson (L–R) deconvolution [105, 106]. This is an iterative method that attempts to determine the most likely value in each image pixel given a known PSF, and given the values in surrounding pixels.

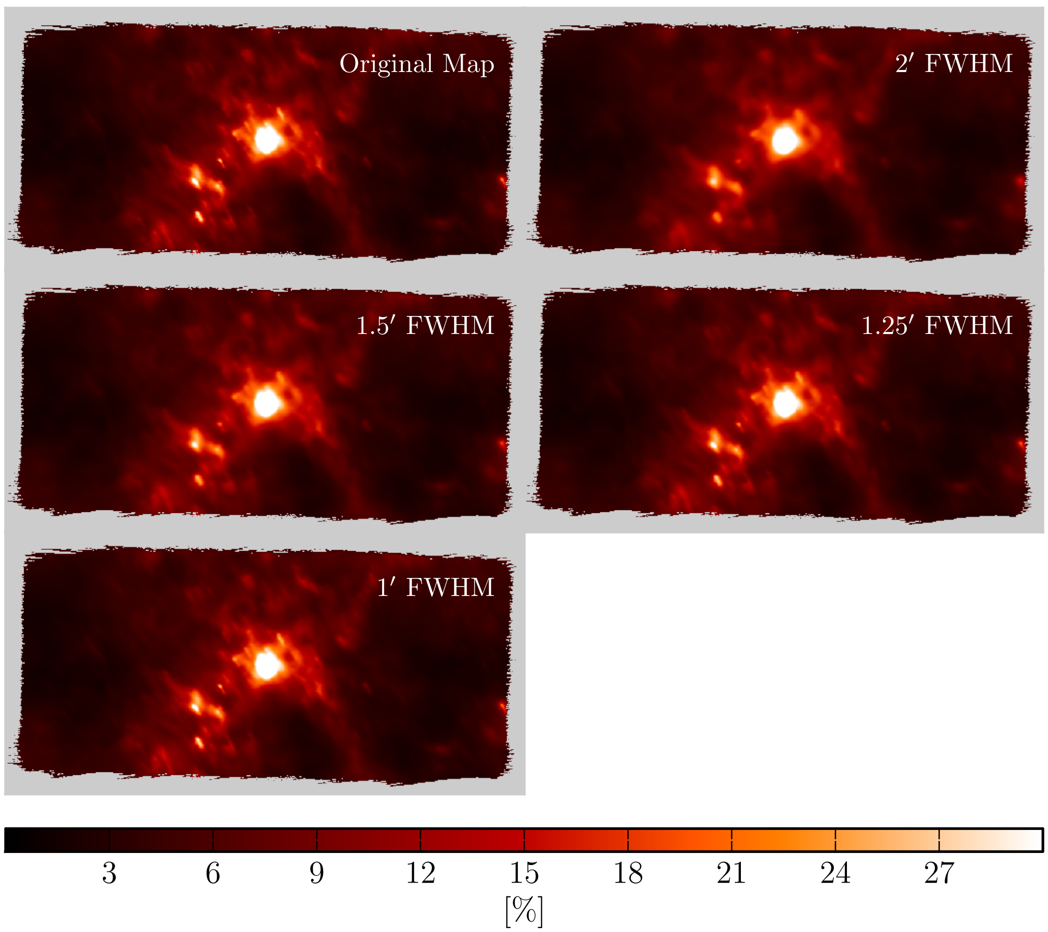


Figure 5.9: Carina Nebula $I$ maps at 350 μm L–R-deconvolved to 2′, 1.5′, 1.25′, and 1′ FWHM resolutions. The colour scale shows the intensity as a percentage of the full dynamic range in the original $I$ map.

The deconvolution used the Python *scikit-image* `skimage.restoration` module's `richardson_lucy()` function. This function was used in two different ways. The first was direct deconvolution of the parallactic-angle-averaged PSF (Figure 5.4) from the Carina map, followed by smoothing with a symmetric Gaussian, in this case of $1.5'$ FWHM. Due to the map size, this required a lot of computation time, even for only 10 iterations, and did not produce a satisfactory result. The second method was to use the L–R function to deconvolve a $50 \times 50$ pixel BLASTPol PSF map from a $50 \times 50$ map of a symmetric Gaussian. The resulting kernel was then used to smooth the Carina maps using the `astropy.convolution.convolve_fft` module.

Figure 5.9 shows the result of applying this second method to the Carina $I$ map at 350 µm, at various resolutions. For FWHMs of $2'$ or less, the method has little effect, suffering from the limitation that the kernel produced is real and positive, and therefore cannot fully-remove structure in the elongated PSF that is larger in extent than the symmetric beam to which the map is being deconvolved. Given the time available, it was determined that L–R deconvolution would only be practical down to $3'$ FWHM, just as with the Fourier method. For analysis, the second L–R method was chosen over the Fourier method, because it seemed to produce final maps free from any image artifacts, when deconvolving to a $3'$ beam.

## 5.5 Map Pre-processing

Before the maps could be used for analysis, a few pre-processing steps had to be applied. The first step was to choose a reference region (bounded by the white outline in Figure 5.10). The reference region is a visually-selected dim region of the $I$ map that is taken to be an area of essentially zero flux. For each of the $I$, $Q$, and $U$ maps, the mean map flux within the reference region was subtracted from the entire map, producing a zero-point-corrected map. The maps were then L–R-deconvolved to $3'$ FWHM as described in the previous section. Figure 5.10 shows the zero-point-corrected TOAST maps in each band: original and deconvolved. The latter were used for analysis.

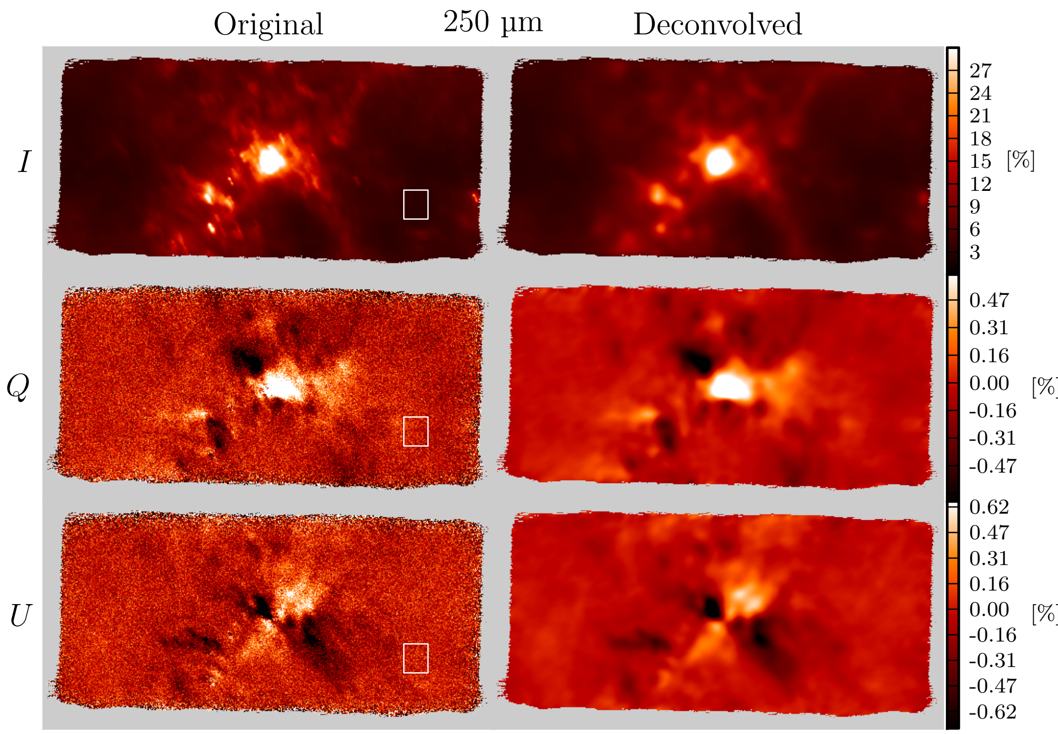


Figure 5.10: Original and L–R-deconvolved Carina Nebula $I$, $Q$, and $U$ maps in each band. The colour scales show the intensity as a percentage of the full dynamic range of the original $I$ map. The white outline bounds the reference region, which was used to correct the zero-point of all the maps (*Continued below*).

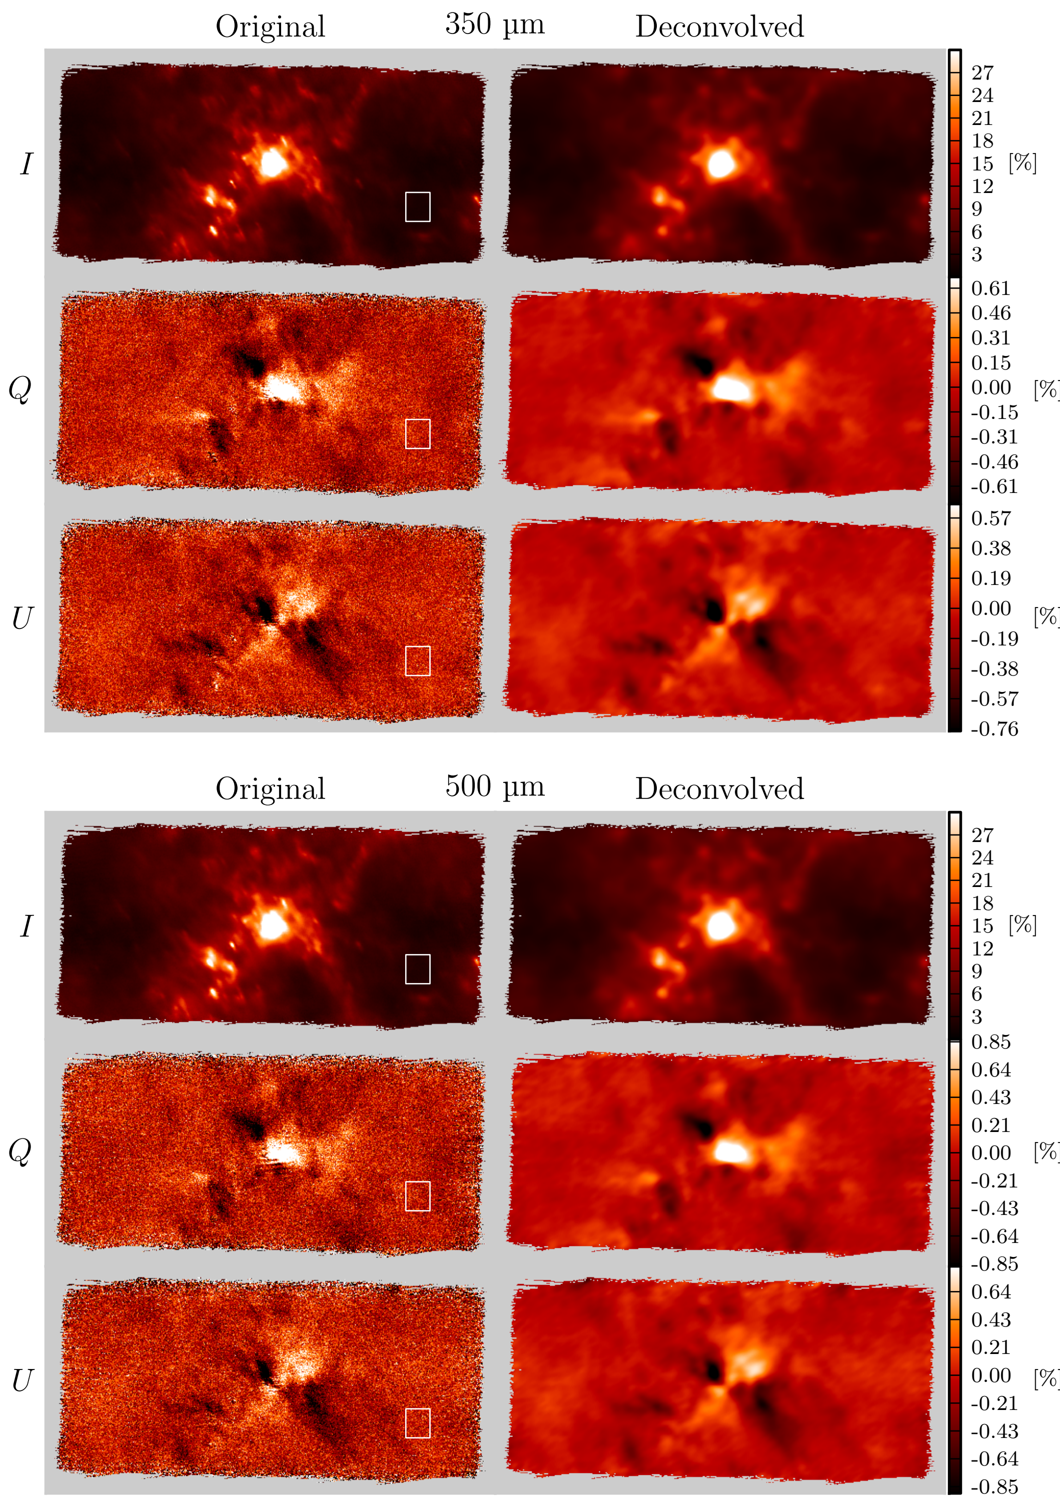


Figure 5.10: (*Continued from above*) Original and L–R-deconvolved maps in each band

## 5.6 Polarization Amplitude and Direction

The final deconvolved $I$, $Q$, and $U$ maps can be used to compute the overall amplitude $P$ of the linear polarization (see also Section 1.6):

$$P \equiv \sqrt{Q^2 + U^2} \tag{5.8}$$

The maps are not yet calibrated into physical units. Therefore the analysis herein does not use the absolute polarization ampltitude. Normalized Stokes parameters $q \equiv Q/I$ and $u \equiv U/I$ are computed instead, and these are used to compute the *fractional* polarization $p$ throughout the source map:

$$p \equiv \sqrt{q^2 + u^2} = \frac{P}{I} \tag{5.9}$$

Assuming that $p$ is associated with the cloud, it can be used to infer the fractional degree of dust grain alignment, and hence something about the physical environment of the dust. The other key piece of information is the plane-of-sky component of the polarization direction $\phi$:

$$\phi \equiv \frac{1}{2}\text{atan2}(u, q) \tag{5.10}$$

Figure 5.11 shows maps of $p$ in each band, displaying the spatial variation of the polarization fraction. Figure 5.12 shows corresponding maps of $\phi$, displaying the spatial variation of the polarization direction. Table 5.1 gives the means, medians, and standard deviations of the $p$ maps, which were computed after applying all data cuts (Section 5.7). The table also gives the standard errors in the means, which were estimated using bootstrapping (Section 5.9).

Another representation of $p$ and $\phi$, in the form of pseudo-vectors, is shown in Figure 5.13. To compute these pseudo-vectors, all data cuts are applied (Section 5.7), and the data are re-binned from the original TOAST $10''$ map pixels into pixels of $1.5'$ (see Section 5.8.2 for details). Figure 5.14 shows the resulting inferred directions of the projected **B**-field.

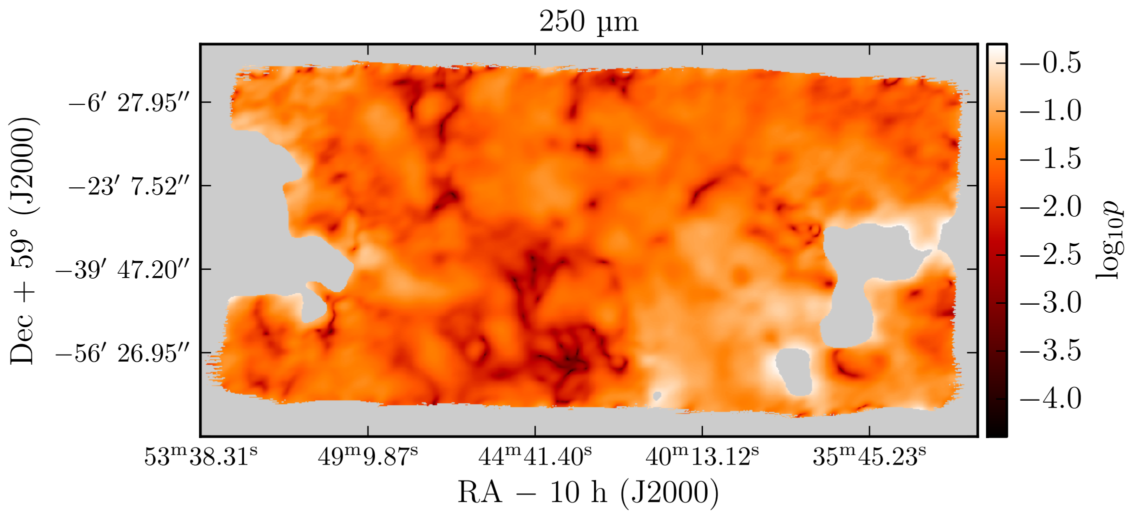


(a) $p$ at 250 µm

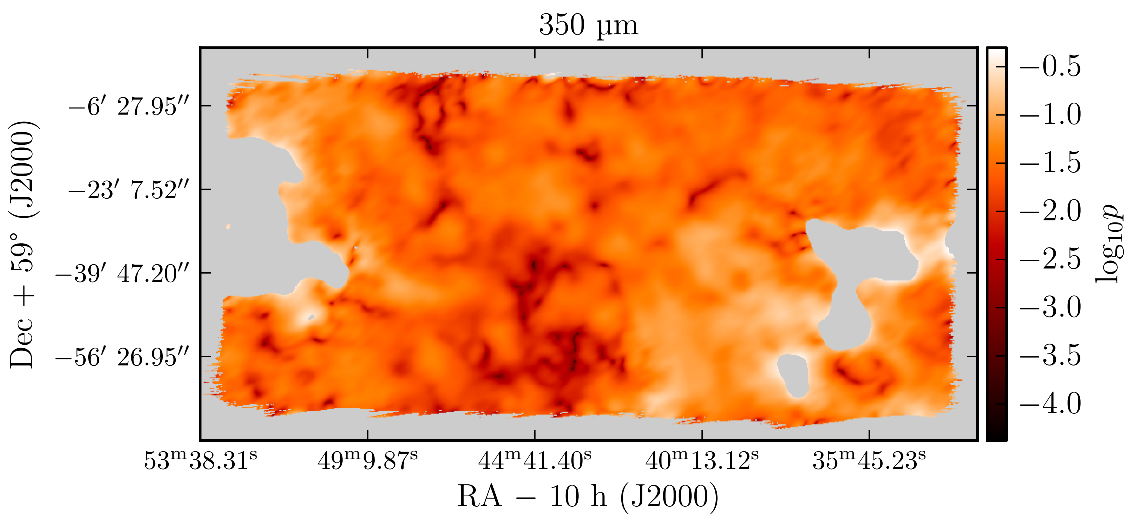


(b) $p$ at 350 µm

Figure 5.11: Fractional polarization $p$ (log scale) in the Carina Nebula in each band (*Continued below*)

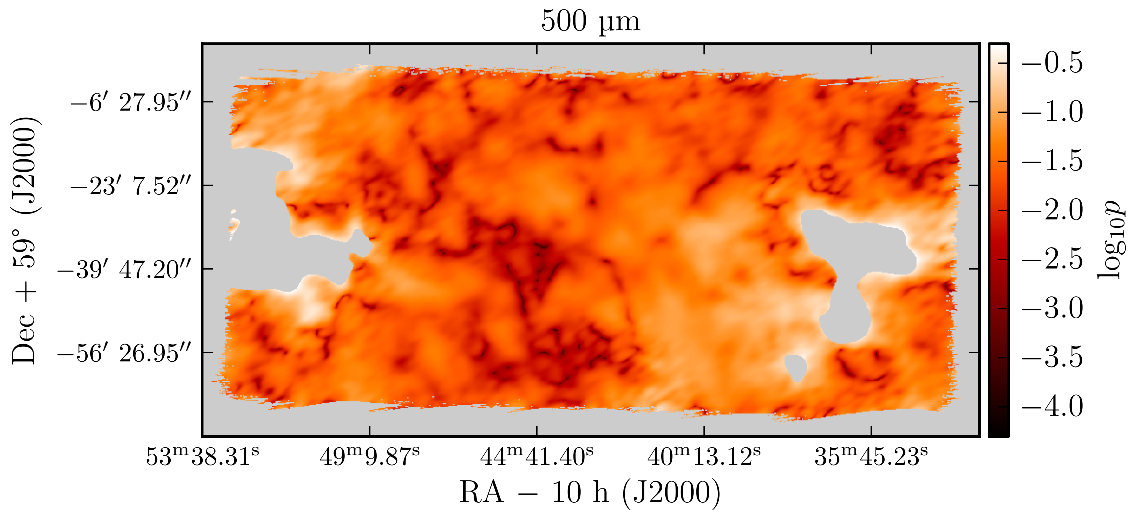


(c) $p$ at 500 μm

Figure 5.11: (*Continued from above*) Fractional polarization $p$ (log scale) in the Carina Nebula in each band

Table 5.1: Statistics of the fractional polarization $p$ in the Carina Nebula in each BLAST-Pol band

| | Mean [%] | Median [%] | Std. Dev. [%] | Std. Error of Mean [%] |
|---|---|---|---|---|
| $p_{250}$ | 6.75 | 5.30 | 5.28 | 0.015 |
| $p_{350}$ | 6.84 | 5.18 | 5.50 | 0.016 |
| $p_{500}$ | 7.06 | 4.96 | 6.63 | 0.019 |

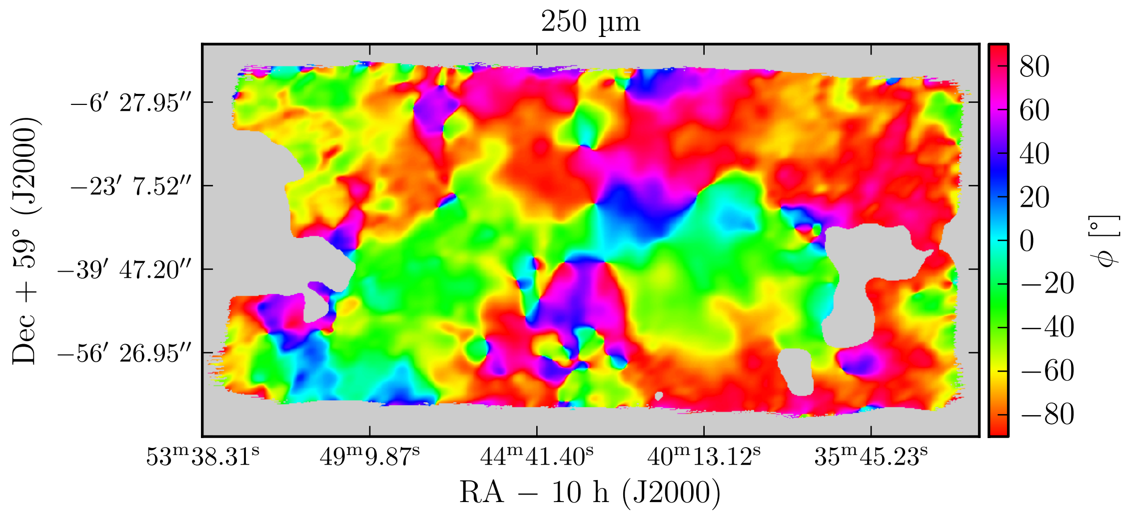


(a) $\phi$ at 250 µm

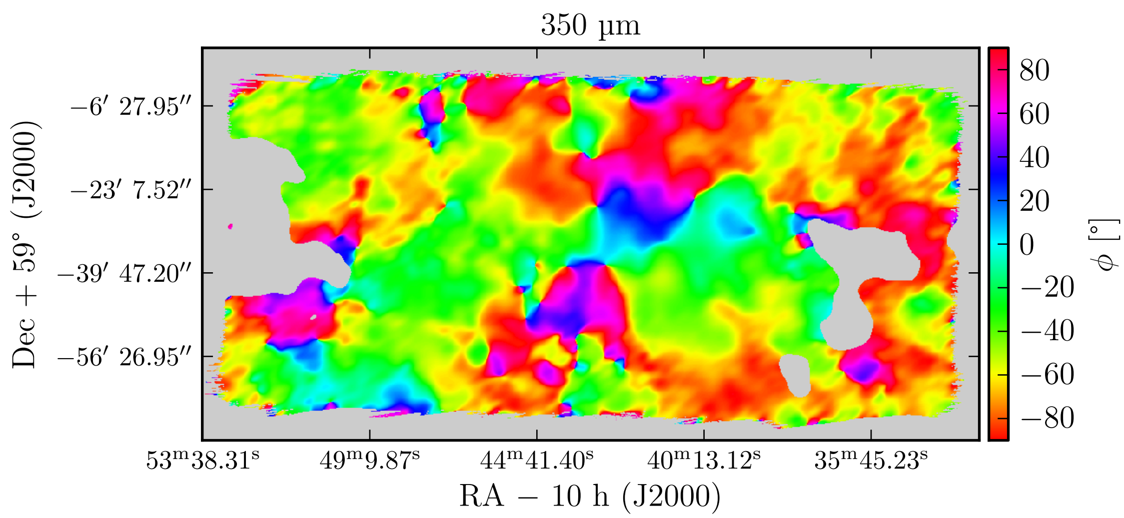


(b) $\phi$ at 350 µm

Figure 5.12: Polarization direction $\phi$ (plotted on a cyclic colour scale) in the Carina Nebula in each band (*Continued below*)

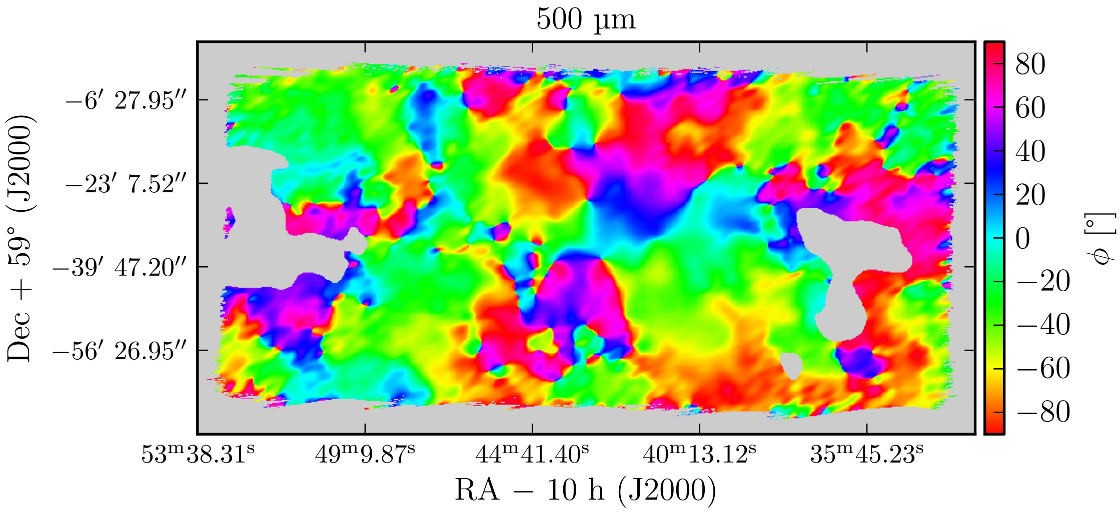


(c) $\phi$ at 500 µm

Figure 5.12: (*Continued from above*) Polarization direction $\phi$ (plotted on a cyclic colour scale) in the Carina Nebula in each band

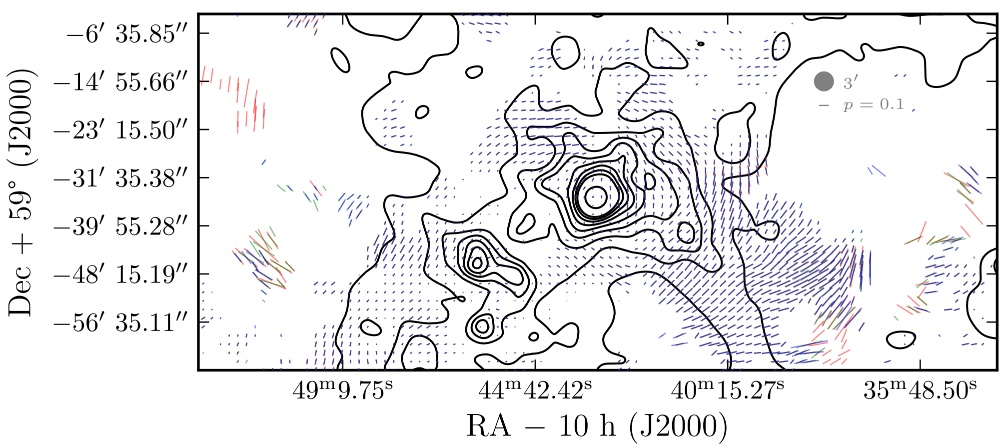


Figure 5.13: Directions and magnitudes of the fractional polarization $p$ in 1.5′ pixels in the Carina Nebula, as measured by BLASTPol at 250 µm (blue), 350 µm (green), and 500 µm (red). The diameter of the grey circle shows the size of the 3′ beam in the deconvolved map. The grey rod length is for 10% polarization. The contours show the BLASTPol $I$ map at 350 µm.

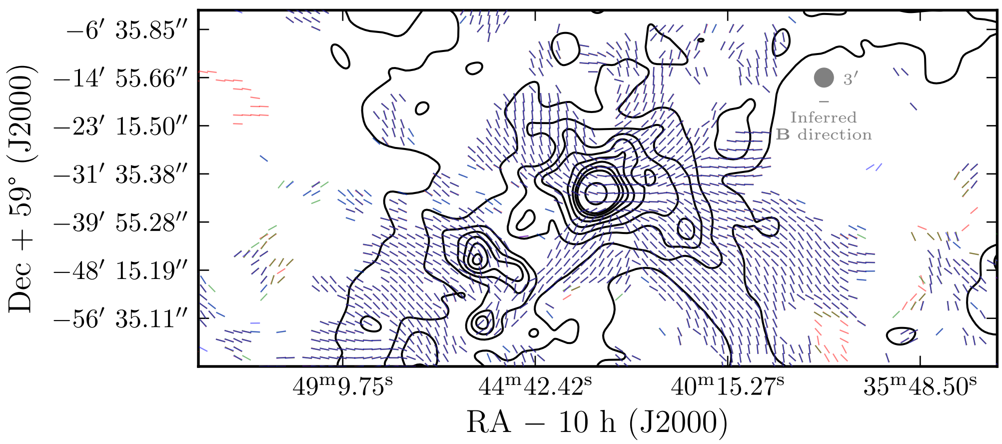


Figure 5.14: Inferred directions of the projected **B**-field in 1.5′ pixels in the Carina Nebula, as measured by BLASTPol at 250 µm (blue), 350 µm (green), and 500 µm (red). The pseudo-vectors have all been drawn the same length, rather than being scaled by $p$, to show the directions more clearly. The contours show the $I$ map at 350 µm.

## 5.7 Data Cuts and Error Estimation

The $p$ and $\phi$ maps shown above have holes that are not in the original TOAST maps. These are from the first two data cuts. First, pixels with $I < \sigma(I_{\mathrm{ref}})$ (the standard deviation of $I$ in the reference region) are excluded. This $I$ cut removes any residual negative and low signal-to-noise values of $I$ that remain after the zero-point correction. Second, pixels with $p > 0.5$ are discarded as having an unphysically-large polarization fraction. The $p$ maps were also divided by the measured BLASTPol instrumental polarization efficiencies: 0.81, 0.79, and 0.82, at 250 µm, 350 µm, and 500 µm respectively.

The next data cut was to exclude low signal-to-noise values of $p$. The practice of previous polarimetric experiments [104, 107] of only including pixels for which $p \geq 3\sigma_p$ was followed here. Therefore, an estimate of the noise variance $\sigma_p^2$ in every map pixel was required. The TOAST map maker produces an estimate of the noise covariances—$\mathrm{Cov}(I, I)$, $\mathrm{Cov}(I, Q)$, $\mathrm{Cov}(I, U)$, $\mathrm{Cov}(Q, Q)$, $\mathrm{Cov}(Q, U)$, and $\mathrm{Cov}(U, U)$—in each map pixel [63]. Thus $\sigma_p^2 = \mathrm{Cov}(p, p)$ can be calculated using the following error propagation formulae:

$$\sigma_q^2 = q^2 \left( \frac{\sigma_I^2}{I^2} + \frac{\sigma_Q^2}{Q^2} - 2\frac{\mathrm{Cov}(I,Q)}{IQ} \right) \tag{5.11}$$

$$\sigma_u^2 = u^2 \left( \frac{\sigma_I^2}{I^2} + \frac{\sigma_U^2}{U^2} - 2\frac{\mathrm{Cov}(I,U)}{IU} \right) \tag{5.12}$$

$$\mathrm{Cov}(q,u) = qu \left( \frac{\sigma_I^2}{I^2} + \frac{\mathrm{Cov}(Q,U)}{QU} - \frac{\mathrm{Cov}(I,Q)}{IQ} - \frac{\mathrm{Cov}(I,U)}{IU} \right) \tag{5.13}$$

$$\sigma_p^2 = \left( \frac{q^2}{p^2}\sigma_q^2 + \frac{u^2}{p^2}\sigma_u^2 + 2\frac{qu}{p^2}\mathrm{Cov}(q,u) \right) \tag{5.14}$$

Similarly:

$$\sigma_\phi^2 = \frac{1}{4}(Q^2 + U^2)^{-2} \left( U^2\sigma_Q^2 + Q^2\sigma_U^2 - 2QU\mathrm{Cov}(Q,U) \right) \tag{5.15}$$

Prior to these calculations, the covariance maps were smoothed with the square of the L–R deconvolution kernel that was used to smooth the Carina maps (Section 5.4.2).

Since $p$ is positive-definite, noise in the $I$, $Q$, and $U$ maps will bias $p$ in the positive direction. Here, a commonly-suggested method of de-biasing $p$ is followed [107]:

$$p_c = \sqrt{p^2 - \sigma_p^2} \tag{5.16}$$

where $p_c$ is the bias-corrected polarization fraction. This correction is applied first, followed by the signal-to-noise cut $p_c \geq 3\sigma_p$.

The final data cut applied is to exclude pixels for which $|\Delta\phi| \geq 10°$. Here, $\Delta\phi$ is the difference in the plane-of-sky component of the polarization angle between any two of the three bands. This $\phi$ cut is an attempt to exclude cases where the observed wavelength-dependence of $p$ is due to the changing inclination angle of the $\mathbf{B}$-field along the line of sight (LOS). Different positions along the LOS are sampled by different wavelengths. While this cut does not guarantee that the field inclination angle does not change with wavelength in the remaining pixels, the rationale as argued in [107] is that in a statistical sense, most of the pixels for which the field inclination angle changes over the LOS will also have an accompanying change in the field angle on the plane of the sky. Any wavelength-dependence of $p$ in pixels remaining after the $\phi$ cut is interpreted as an actual property of the polarization spectrum in a single location, arising from grain alignment physics. This is the subject investgated in the remainder of this chapter.

## 5.8 Polarization Ratios

A natural question that can be addressed by this data set is whether the fractional polarization differs significantly between wavelength bands. If so, a further question is how these differences vary spatially across the map. These variations could indicate differing degrees of grain alignment in different environments.

### 5.8.1 Histograms of the Polarization Ratios

Previous studies of submillimetre polarization in Galactic molecular clouds [107, 108] have examined the distribution of the ratio of $p$ between two different bands. A similar analysis for the BLASTPol observations of Carina is presented in Figure 5.15, which shows histograms of the ratios $p_{350}/p_{250}$ and $p_{500}/p_{350}$. The distributions are shown for the cases of all the data ($I$ cut only, in blue), for only pixels with $p_c \geq 3\sigma_p$ (in green), and for pixels meeting that criterion and also having $|\Delta\phi| < 10°$ (in red). The original $10''$ map pixels are used for this analysis. Table 5.2 summarizes some statistics of these distributions. Following [107], to characterize the width of these asymmetric distributions, which have a long positive tail, the median absolute deviation (MAD) is defined:

$$\mathrm{MAD} \equiv \mathrm{median}(|x_i - x_m|) \tag{5.17}$$

where $x_i$ are the data points in the distribution, and $x_m$ is their median value.

Table 5.2: Statistics of the fractional polarization ratios between the BLASTPol bands in Carina

| | № of Points | Mean | Median | Std. Dev. | Median Absolute Dev. |
|---|---|---|---|---|---|
| $p_{350}/p_{250}$ | | | | | |
| All Data | 304 832 | 1.094 | 0.986 | 0.844 | 0.147 |
| Only $p_c \geq 3\sigma_p$ | 263 309 | 1.021 | 0.984 | 0.276 | 0.125 |
| Also $\lvert\Delta\phi\rvert < 10°$ | 114 600 | 0.994 | 0.985 | 0.164 | 0.077 |
| $p_{500}/p_{350}$ | | | | | |
| All Data | 300 850 | 0.986 | 0.880 | 1.322 | 0.205 |
| Only $p_c \geq 3\sigma_p$ | 226 767 | 0.948 | 0.912 | 0.345 | 0.158 |
| Also $\lvert\Delta\phi\rvert < 10°$ | 112 219 | 0.953 | 0.938 | 0.246 | 0.110 |

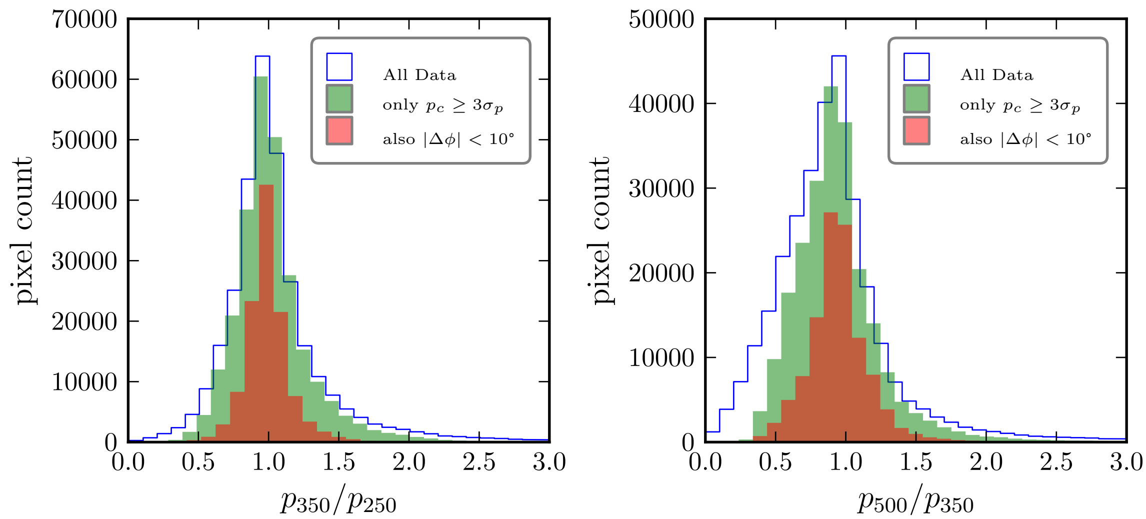


Figure 5.15: Histograms of the polarization ratios between BLASTPol bands in the Carina Nebula. The bin width is 0.1. *Left* – The ratio $p_{350}/p_{250}$. *Right* – The ratio $p_{500}/p_{350}$.

### 5.8.2 Spatial Variation of the Polarization Ratios

Figures 5.16 and 5.17 show the variation of the fractional polarization ratios $p_{350}/p_{250}$ and $p_{500}/p_{350}$, superimposed on contours of the deconvolved Carina $I$ map at 350 µm. For these plots, all data cuts have been applied, and the $p$-ratio maps were re-binned into $1.5'$ pixels by averaging the $p$-ratio values in every $9 \times 9$ set of $10''$ pixels in the original map. A pixel value is only displayed if more than half of these $9 \times 9$ pixels have data in them. There is no evidence of a drop in polarization ratio towards flux peaks, as seen in [108] in the Orion Molecular Cloud. However, the $p_{500}/p_{350}$ map does appear to have a quadrupolar structure centred on $\eta$ Carinae, similar to what appears in the $U$ and $\phi$ maps. There also appear to be areas of a higher polarization ratio (in both maps) at the edges of the South Pillars region.

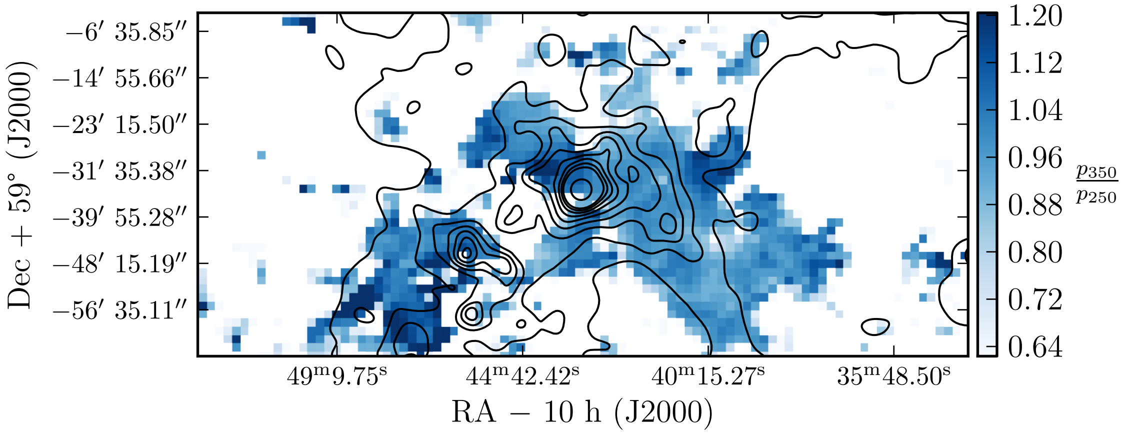


Figure 5.16: Spatial variation of $p_{350}/p_{250}$ in Carina in 1.5′ pixels. The contours (5% to 50% with 5% spacing, and also 75%) show the $I$ map at 350 µm. The colour scale spans $\pm 1.5\sigma$ around the mean ratio.

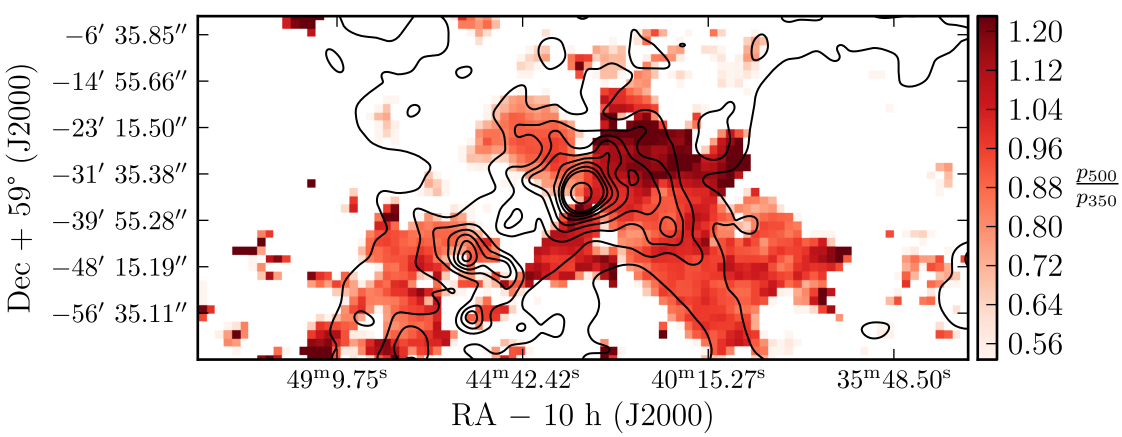


Figure 5.17: Spatial variation of $p_{500}/p_{350}$ in Carina in 1.5′ pixels. The contours (5% to 50% with 5% spacing, and also 75%) show the $I$ map at 350 µm. The colour scale spans $\pm 1.5\sigma$ around the mean ratio.

## 5.9 Polarization Spectra

Enough information now exists to examine the polarization fraction, averaged over the whole cloud, as a function of wavelength. This *polarization spectrum* is shown in Figure 5.18 for three different cases: all the data (meaning $I$ cut only, in blue), and then two successive data cuts requiring $p_c \geq 3\sigma_p$ (in green) and $|\Delta\phi| < 10°$ between any two of the three bands (in red). As is the convention in the field, in a given polarization spectrum, the point plotted for the $i^{\text{th}}$ band is median$(p_i/p_{350})$.

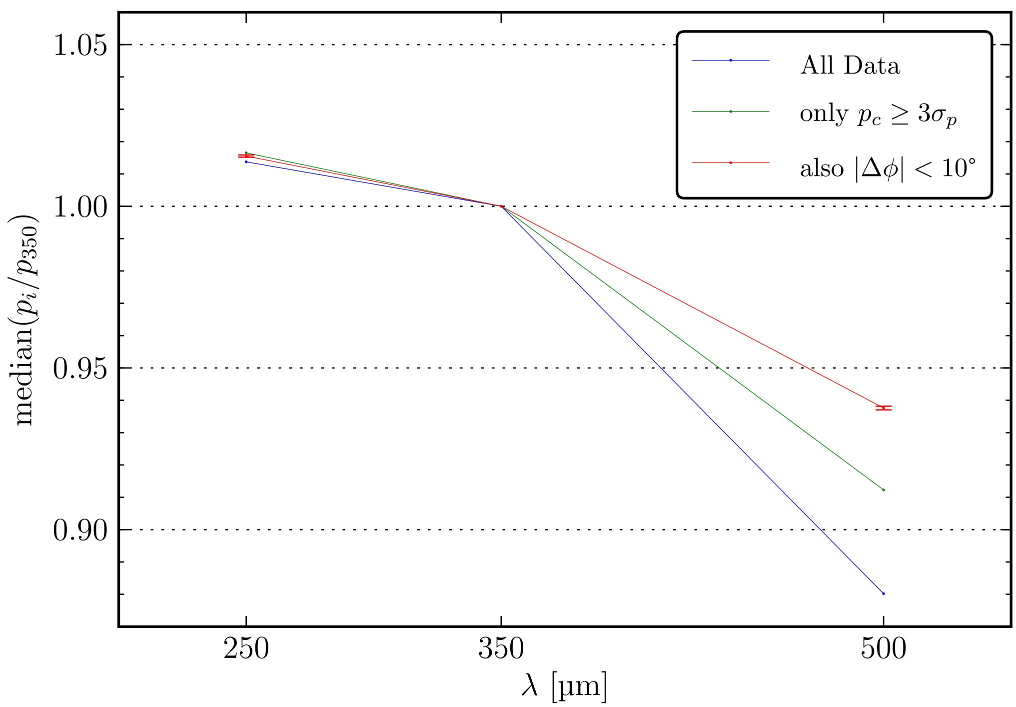


Figure 5.18: Fractional polarization vs. wavelength in the Carina Nebula, normalized to 350 µm. For each plot, the data point in the $i^{\text{th}}$ band is the median of the ratio $p_i/p_{350}$. This polarization spectrum is shown for the case of using all the data (blue), selecting only pixels for which $p_c \geq 3\sigma_p$ (green), and then also requiring that those pixels have $|\Delta\phi|$ between any two bands be $< 10°$ (red). The error bars in the red case show the standard errors of the median ratios, which were estimated using bootstrapping (see text).

The error bars shown for the red case are estimates of the standard errors of these median ratios, computed using bootstrapping. For each median ratio, $N$ data points remained after applying all the data cuts (see Table 5.2). Bootstrapping entails randomly sampling $N$ points from the data, with replacement. Each of the polarization ratio maps for $p_{500}/p_{350}$ and $p_{250}/p_{250}$ was resampled 10 000 times in this manner, by simply using the `numpy.random.rand()` function to generate $N$ random indices in the range $[0, N-1]$ for each data array. For each resampling, the median ratio was computed. The standard error was taken to be the standard deviation of these 10 000 medians. The results were $\mathrm{median}(p_{500}/p_{350}) = 0.9376 \pm 0.00056$, and $\mathrm{median}(p_{250}/p_{350}) = 1.0155 \pm 0.00035$.

The data cuts produce significant changes in the polarization spectrum, relative to the estimated error. For the rejection of pixels with low signal-to-noise ratio in $p$, this makes sense. However, the $\phi$ cut also produces just as much variation. An interesting question is whether the pixels that have been rejected due to large variation in $\phi$ between bands contain such variation due to noise, map artifacts, or real physics. To evaluate this, the data were divided into six bins in signal-to-noise ratio ($p/\sigma_p$ at 250 µm) such that each bin contained approioximately the same number of points. For each bin, a histogram of $|\Delta\phi| \equiv |\phi_{250} - \phi_{350}|$ was generated. This analysis used the data before $p$ and $\phi$ cuts, corresponding to 307 587 points. The result is shown in Figure 5.19.

The $|\Delta\phi|$ distributions are wider in the lower signal-to-noise bins, and become progressively narrower as $p/\sigma_p$ increases. This is quantified by the MAD. A reduced chi-squared statistic was also computed for each bin:

$$\chi_r^2 = \frac{1}{N-1}\sum_{i=1}^{N}\frac{(|\Delta\phi|_i - |\Delta\phi|_m)^2}{(\sigma_\phi^2)_i} \tag{5.18}$$

Here, $|\Delta\phi|_i$ is the polarization angle difference in the $i^{\text{th}}$ pixel, and $|\Delta\phi|_m$ is the median polarization angle difference. The noise variance of $\phi$ in the $i^{\text{th}}$ pixel is computed for each band using Eq. 5.15, and the variances for the two bands are added to compute $(\sigma_\phi^2)_i$. The chi-squared values of order unity suggest that most of the variation in $\phi$ between bands can be accounted for by noise, particularly in the case of pixels that were rejected by the $\phi$ cut.

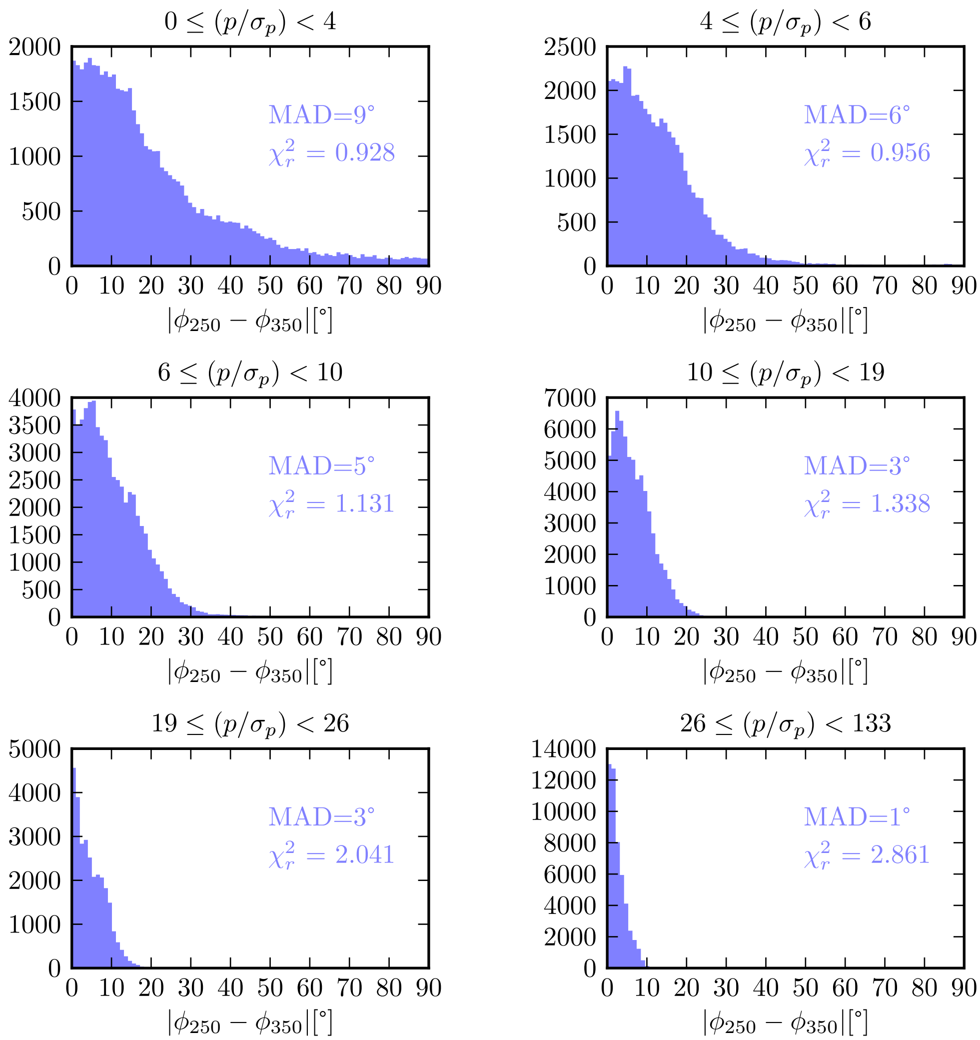


Figure 5.19: Histograms of $|\Delta\phi| = |\phi_{250} - \phi_{350}|$ in the Carina Nebula for six different $p/\sigma_p$ bins. For each plot, the median absolute deviation (MAD) of the distribution is shown, along with the value of the chi-squared statistic (see text).

## 5.10 Discussion

In the submillimetre regime, $\lambda \gg a$, with $a$ being the characteristic dust grain radius of 0.1 µm to 1 µm. Therefore, no variation of the polarization spectrum with wavelength is expected for a single dust grain population [107]. Models have attempted to explain the observed variation of the polarization spectra in molecular clouds in the far-IR and beyond ($\lambda \geq 50$ µm) using two or more dust grain populations with differing degrees of alignment, and that also differ in some other grain property that affects the emitted radiation, such as temperature, grain size, or emissivity [109].

Previous studies [107, 108] have found a V-shaped polarization spectrum in Galactic molecular clouds, in which the fractional polarization at shorter wavelengths falls towards a minimum near 350 µm, before rising again at larger wavelengths. This result is usually interpreted in the context of multiple dust grain populations. However, it is not reproduced here. The BLASTPol data produce a purely falling spectrum in the Carina Nebula. This result is similar to that reported for M17 in [110], in which the median fractional polarization at 450 µm was actually lower than at 350 µm. In that paper (their Figure 5, for example), a qualitative picture of the effects of embedded sources on the polarization spectrum is presented to explain this result. The explanation is given in the context of the RAT model of grain alignment (Section 3.1.1): cool dust far from internal radiation sources is expected to contribute a positive-slope component to the polarization spectrum. Therefore, going from longer to shorter wavelengths, the spectrum is initially falling. However, warm dust heated by internal sources such as young stellar objects begins to contribute at shorter wavelengths, causing the spectrum to turn over and begin rising again at the shortest submillmetre wavelengths.

As explained by Zeng *et al.* in [110], under the RAT mechanism, the most heavily-irradiated and hence hottest grains would be the best-aligned, explaining why the spectrum continues rising with decreasing wavelength. The warm, irradiated dust near embedded sources therefore contributes a negative-slope component to the polarization spectrum. They note that the region of M17 they observe is compressed by winds and radiation from a central cluster of OB stars. Therefore, the negative-slope component dominates, and the minimum of the polarization spectrum is shifted to longer wavelengths, beyond 450 µm (in contrast to other clouds, where the minimum is near 350 µm). They attribute their monotonically-falling spectrum to this effect. This explanation could hold in the Carina Nebula, large parts of which are similarly influenced by radiation from massive stars. According to this view, BLASTPol has measured a monotonically-falling spectrum all the way to 500 µm because the polarized dust emission is dominated by warm dust heated by radiation sources internal to the cloud. The minimum in the polarization spectrum, then, has potentially been shifted longward of the BLASTPol bands.

This data set is the most detailed measurement of submillimetre polarization in the Carina Nebula to date. It has provided evidence that the fractional polarization vs. wavelength may be different in this region than in other more quiescent molecular clouds. This difference appears even when comparing to BLASTPol measurements of other targets, such the Vela C Molecular Ridge [97]. The potential effects of the radiative environment of the cloud hinted at here are worthy of further study.

# Chapter 6

# Conclusion

The experimental techniques presented herein will prove useful to future balloon-borne astrophysical experiments. The first flight of SPIDER has demonstrated the effectiveness of new pointing control system elements, including the elevation drive (Section 2.7) and velocity control mode for the pivot motor (Section 2.6.2.2). The latter has been demonstrated in flight to produce superior pointing stability compared to pivot torque mode (Section 2.10) and to allow a nearly 3000 kg payload to scan in azimuth with a peak speed of $\sim$5 deg/s, and a peak angular acceleration of 0.5 deg$\cdot$s$^{-2}$. Therefore, these new pointing system elements will be retained in the gondola that will be constructed for the second flight of SPIDER. A second flight is already funded, and could occur as early as the 2017–2018 austral summer. As of this writing, a new flight cryostat is under construction for this purpose. The goal of the second flight is to supplement the 94 GHz and 150 GHz receivers with a set of higher-frequency channels: 220 or 280 GHz, for example. These channels will enable better characterization of the polarized foreground emission from Galactic dust. Barring an actual detection, data from the first two flights should set a $3\sigma$ upper limit of $r < 0.03$ [42, 43].

The pivot velocity control mode has also already been adopted in the pointing control system of the Balloon-borne Imaging Testbed (BIT). BIT is a prototype instrument consisting of a three-axis gondola carrying a 0.5 m visible-light telescope coupled to CCD cameras for imaging and pointing. It is designed to test the feasibility of achieving sub-arcsecond pointing stability from a balloon. This level of control is a prerequisite to the development of a 2 m balloon-borne visible and near-UV wide-field imaging telescope. As of this writing, an overnight flight of BIT is scheduled for September 2015, launching from Timmins, Ontario, Canada.

Similarly, the power system design presented herein (Section 2.5), which has now proven to be reliable in three separate balloon flights, will be employed again in the second flight of SPIDER. It will likely also be used to power ULDB successors to BIT. The power breakout board developed for BIT uses the power switching and current sensing circuits that are presented herein (Appendix A.1).

Detailed modelling of the *BLASTPol12* beam has contributed crucially to the BLAST-Pol data analysis effort (Chapter 4). The beam templates for each band have been used in low-level data processing, including centroiding to determine pointing offsets, and array flat-fielding. They have also been used as part of the BLASTPol simulation pipeline, and for map deconvolution (Section 5.4). The first attempts at deconvolution did not entirely succeed at recovering information at angular scales below the extent of the elongated in-flight PSF. However, there are plans to revisit this problem in a second round of data analysis. One potential area for improvement is the beam template itself. The BLASTPol collaboration has been allocated time to observe IRAS 08470-4243 in the submillimetre using the Large APEX Bolometer Camera (LABOCA) on the APEX (Atacama Pathfinder Experiment) telescope. High-resolution images of this object would allow beam templates derived from BLASTPol observations of it to be refined. A second area for improvement is the L–R deconvolution technique, which could potentially produce a better result, given further investigation.

Analysis of the fractional polarization $p$ in the Carina Nebula has found a monotonically-falling spectrum with increasing wavelength, over the BLASTPol bands (Section 5.9). A preliminary interpretation of this result is that the polarized emission is dominated by warm dust irradiated by stellar sources within the molecular cloud. It is consistent with the RAT model of grain alignment that the hottest, most-irradiated grains have the highest fractional degree of alignment, and hence the polarization fraction is highest at the shortest wavelengths. However, the data cannot rule out an alternative explanation for this result, in which the polarization is not associated with the dense molecular cloud at all, but is emitted by clouds of warmer and more diffuse polarized material in the foreground. The maps of the polarization direction $\phi$ (Figure 5.12) appear to have several overlapping areas of different $\phi$, hinting at the possibility of multiple clouds along the line of sight with differing polarization amplitude and direction. However, it is unclear whether this is a real effect. For future work, a velocity study using spectroscopic data of Carina could help disentangle the contributions from multiple sources along the line of sight, if present.

# Appendix A

# Circuit Schematics

## A.1 Power Breakout Board (PBOB)

### A.1.1 Type 1: Latching

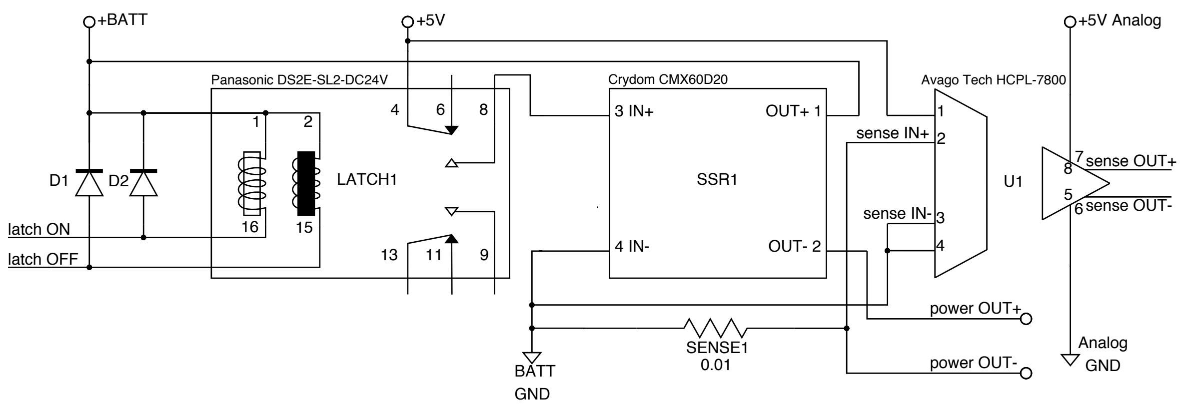


Figure A.1: The latching circuit designed for the PBOB

The first type of circuit designed for the PBOB is shown in Figure A.1. The set (white) or reset (black) coil of the latching relay (LATCH1) is pulsed low by a digital output from either the ACS or the SIP. Applying a pulse to the latch ON line connects the input of a Crydom solid state relay (SSR1) to +5 V. This closes SSR1 (pin 1 connects to pin 2), which connects power OUT+ to the source voltage. The circuit shown in this example provides switched battery power out, which was the most common case on the payload[1]. A 10 mΩ sense resistor is included in the current path of the output power circuit. The voltage across the resistor is supplied to the input of an isolation amplifier (U1) which provides an output signal proportional to this input, referenced to BLASTbus analogue ground. On BLASTPol, this type of circuit switched power to the ACS, DAS, dGPS, and LOS transmitters. On SPIDER, this type of circuit switched power to the ACS, MCCs, Sync Box, MCEs, HK DAS, RSC, BSC, dGPS, LOS transmitters, HWP motors, and SFT motorized valve.

### A.1.2 Type 2: Non-Latching

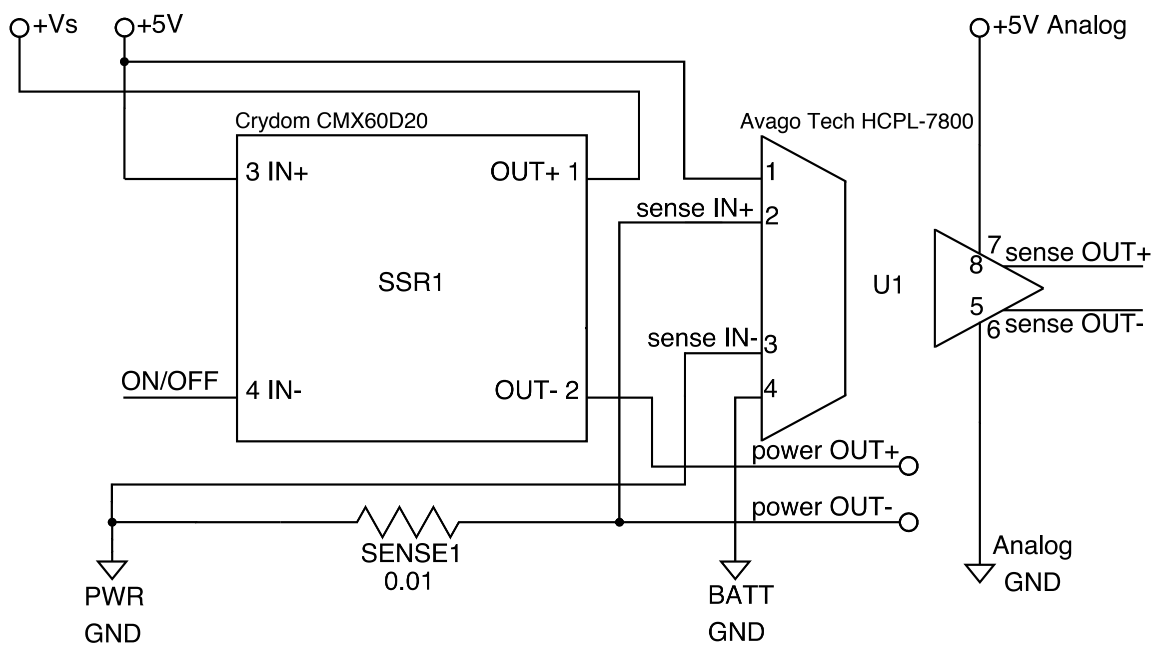


Figure A.2: The non-latching circuit designed for the PBOB

Figure A.2 shows the non-latching circuit, which includes a solid state relay (SSR1), but no mechanical latching relay. The input of SSR1 is driven directly by a BLASTbus

[1]The PBOB included an auxiliary version of this circuit, for which the power supply being switched was configurable.

digital open-collector output. The voltage level at pin 4 (IN−) must be maintained in order to maintain the relay state. Therefore, the power state of subsystems that are switched by this type of PBOB circuit cannot be controlled by the SIP, which provides only voltage pulses. The power state of these subsystems is entirely dependent on the state of the BLASTbus digital outputs in the ACS. Therefore, the Type 2 circuit is typically used for lower-priority subsystems for which this is acceptable. On BLASTPol, circuits of Type 2 switched power to the actuators (elevation lock motor, secondary mirror, and HWP), the magnetometers, and the inclinometers. The latter two required analogue supply voltages other than battery power. These low power pointing sensors also did not have their current sensed individually: it was tied into the main current draw of the ACS. Therefore, the circuits for these systems did not have sense resistors or isolation amplifiers. On Spider, this type of circuit once again switched power to the magnetometer and inclinometers, as well as the HK preamplifier crate, and the elevation lock motor.

In this example, the relay is shown switching a generic power supply $V_s$, whose ground reference is PWR GND. The +5 V used to power the inputs of the SSR and isolation amplifier is referenced to battery ground, just as in the case of Type 1. The +5 V powering the output stage of the isolation amplifier is referenced to BLASTbus analogue ground.

### A.1.3 Type 3: High-Current Latching

For the subsystems drawing the highest current, the pointing motors, the switching circuit of Figure A.3 was developed. In place of the Crydom solid state relays, a KG Technologies K100 mechanical latching relay (LATCH1) was chosen. This component is rated for 100 A and can handle bi-directional current, making it suitable for inductive loads. The BLASTbus digital outputs can only source up to 50 mA: not enough current to drive the 125 Ω relay coils directly at 24 V. Instead, the digital outputs drive the inputs of two optocouplers (SSR1 and SSR2) whose outputs drive the set and reset coils of the K100.

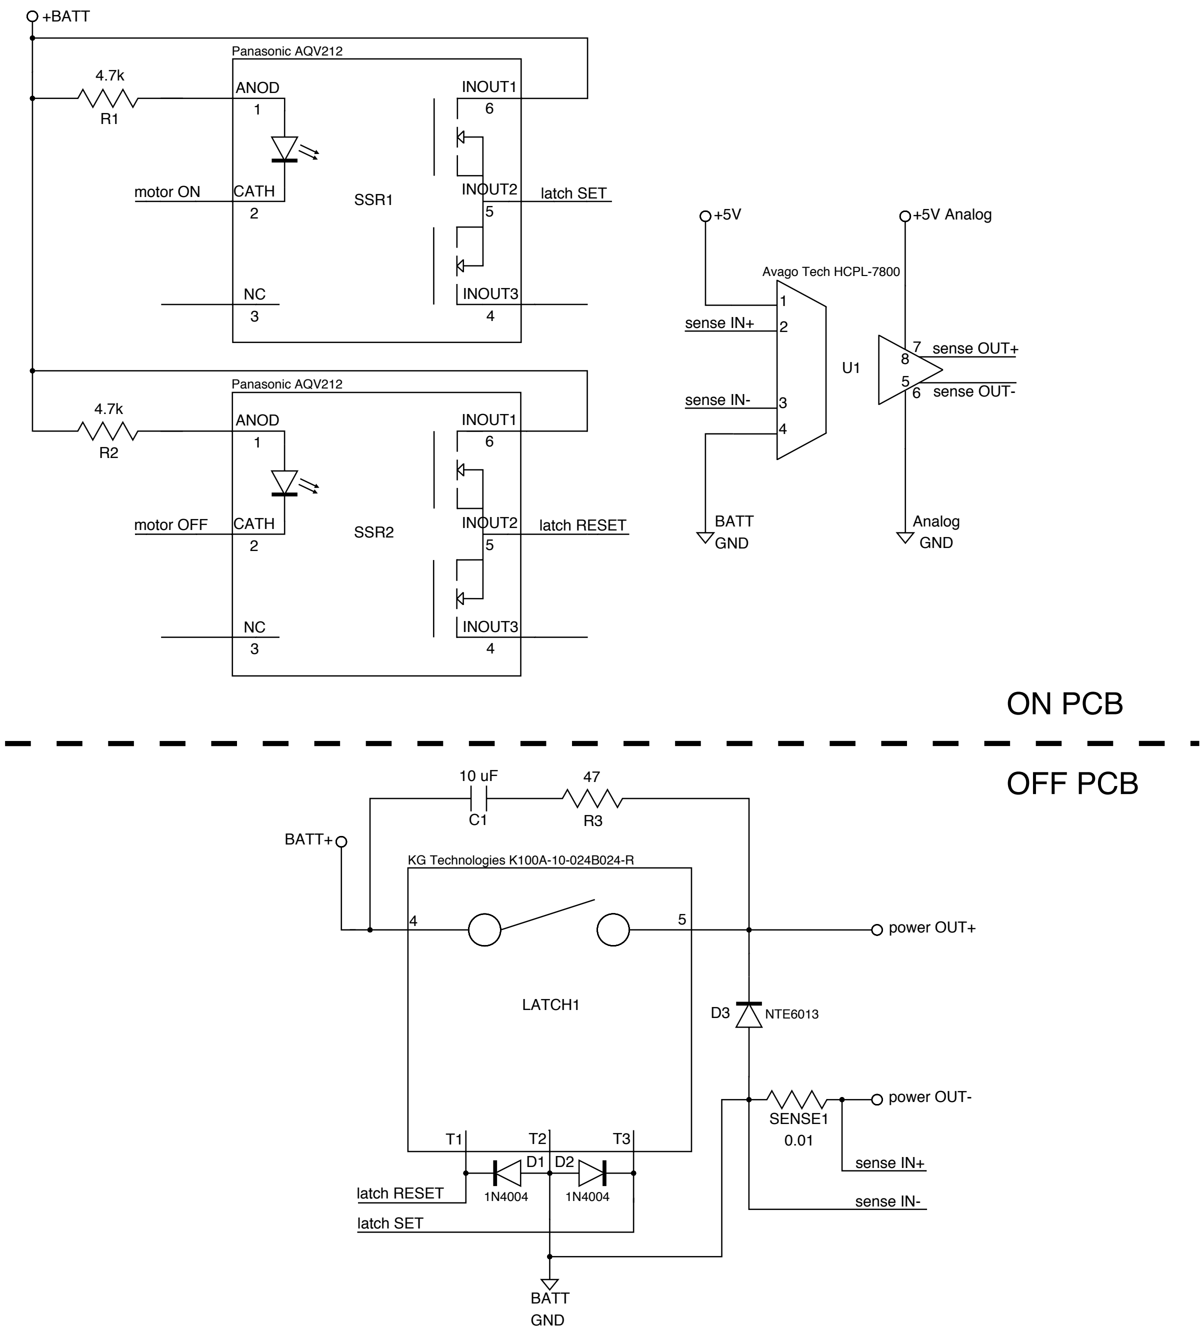


Figure A.3: The high-current latching circuit designed for the PBOB

To avoid having PCB traces carry the large motor current, the relay, sense resistor, and flyback diodes are not located on the PBOB. They are mounted on a separate FR-4 fibreglass sheet also located within the ACS. The optocouplers and the isolation amplifier for current sensing are located on the PCB.

The BLASTPol PBOB had three instances of this circuit, one for each of the elevation, pivot, and reaction wheel drive motors. For Spider, the elevation drive's K100 relay was replaced with the cascade of the small Panasonic latching relay and the Crydom solid state relay from Type 1. This change was made because the K100 latching relay is rated for only 100 000 mechanical cycles, making it marginal for the originally-planned operating mode in which the Spider El drive would be powered on for every El step, and powered off otherwise. Eliminating the high-current relay was possible in this case, because Spider's stepper motors did not draw nearly as much current as the direct drive brushless DC torque motor that was used for the BLASTPol elevation drive.

## A.2 Elevation Drive Delay Timer Circuit

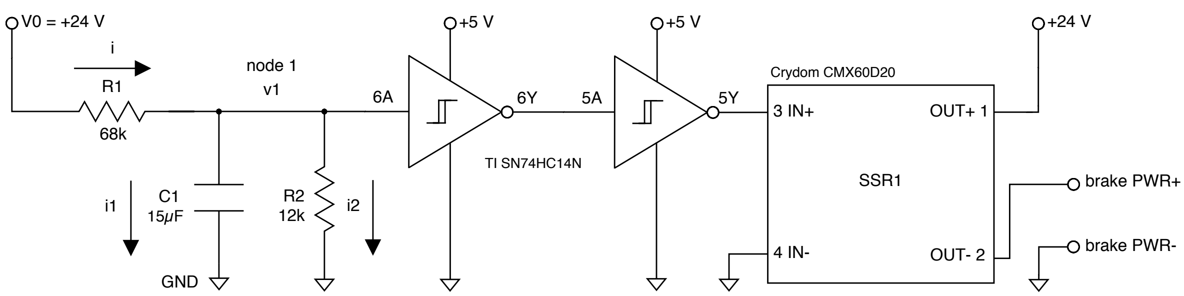


Figure A.4: The El drive delay timer circuit

As discussed in Section 2.7.2, a delay timer circuit was introduced to prevent the Spider elevation drive inline brakes from disengaging until the Cool Muscle stepper motors were powered on and applying holding torque. This delay was implemented in hardware to ensure that the brakes could never be disengaged while the motors were unpowered. The circuit is shown in Figure A.4. The design used a relatively simple RC circuit to set the time constant. For each brake, the 24 VDC was supplied by the same Vicor DC-DC that powered the associated Cool Muscle stepper motor. While the 24 V was applied to the motor immediately, the brake was not powered until the voltage at node 1 rose high enough to close the SSR (pin 1 connects to pin 2), powering the brake to disengage it.

The rising voltage at node 1 was the input to a Schmitt trigger buffer. The hysteresis provided by this component ensured that a sustained logic high level was provided to the input of the SSR. Since the SN74HC14 hex Schmitt trigger inverter chip is what was readily available, two inverters were simply cascaded to restore the original polarity. The 5 VDC to power the logic IC was provided by a TDK-Lambda CC1R5-2405-SF-E DC-DC converter.

Applying Kirchhoff's Current Law to node 1, and assuming negligible input current to the buffer stage:

$$i = i_1 + i_2 \tag{A.1}$$

$$\frac{V_0 - v_1}{R_1} = C_1 \frac{dv_1}{dt} + \frac{v_1}{R_2}$$

Rearranging leads to the differential equation

$$\frac{dv_1}{dt} + \frac{v_1}{RC_1} = \frac{V_0}{R_1 C_1}, \tag{A.2}$$

where $R \equiv R_1 \parallel R_2$. Solving results in the voltage

$$v_1(t) = V_0 \frac{R_2}{R_1 + R_2} [1 - \exp(-t/RC_1)] \tag{A.3}$$

The steady-state voltage is as expected, given the voltage divider arrangement. The time constant of the circuit is

$$\tau = RC_1 = \frac{R_1 R_2}{R_1 + R_2} C_1 \tag{A.4}$$

Given the component values above, $\tau = 0.153$ s, and the steady-state voltage is 3.6 V. The buffer was measured to turn on at $v_1 = 2.64$ V. Therefore, from Eq. A.3, the time delay before the brake was powered was 202 ms.

# Bibliography


[1] Hubble, E. P. "A Relation Between Distance and Radial Velocity Among Extra-galactic Nebulae." In: *Proceedings of the National Academy of Sciences* 15, 3 (Mar. 1929), pp. 168–173.

[2] Wagoner, R. V., Fowler, W. A., and Hoyle, F. "On the Synthesis of Elements at Very High Temperatures". In: *The Astrophysical Journal* 148 (Apr. 1967), pp. 3–49.

[3] Penzias, A. A. and Wilson, R. W. "A Measurement of Excess Antenna Temperature at 4080 Mc/s." In: *The Astrophysical Journal* 142 (July 1965), pp. 419–421.

[4] Fixsen, D. J. et al. "The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set". In: *The Astrophysical Journal* 473 (Dec. 1996), pp. 576–587.

[5] Perlmutter, S. et al. "Measurements of $\Omega$ and $\Lambda$ from 42 High-Redshift Supernovae". In: *The Astrophysical Journal* 517 (June 1999), pp. 565–586.

[6] Guth, A. H. "Inflationary universe: A possible solution to the horizon and flatness problems". In: *Physical Review D* 23 (Jan. 1981), pp. 347–356.

[7] Hubble, E. P. "Extra-galactic Nebulae." In: *The Astrophysical Journal* 64 (Dec. 1926), pp. 321–369.

[8] *Planck* Collaboration. "*Planck* 2013 results. XVI. Cosmological parameters". In: *Astronomy & Astrophysics* 571, A16 (Nov. 2014).

[9] Einstein, A. "Die Feldgleichungen der Gravitation (The Field Equations of Gravitation)". In: *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)* (1915), pp. 844–847.

[10] Robertson, H. P. "Kinematics and World-Structure". In: *The Astrophysical Journal* 82 (Nov. 1935), pp. 284–301.

[11] Walker, A. G. "On Milne's theory of world structure". In: *Proceedings of the London Mathematical Society, Series 2* 42 (1936), pp. 90–127.

[12] Dodelson, S. *Modern Cosmology*. Academic Press, 2003.

[13] Longair, M. S. *Galaxy Formation*. 2nd ed. Springer-Verlag, 2008, p. 429.

[14] Linde, A. D. "A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems". In: *Physics Letters B* 108 (Feb. 1982), pp. 389–393.

[15] Albrecht, A. and Steinhardt, P. J. "Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking". In: *Physical Review Letters* 48 (Apr. 1982), pp. 1220–1223.

[16] Liddle, A. R. "An Introduction to Cosmological Inflation". In: *High Energy Physics and Cosmology, 1998 Summer School*. Ed. by Masiero, A., Senjanovic, G., and Smirnov, A. 1999, p. 260.

[17] Linde, A. D. "Inflationary Cosmology". In: *Inflationary Cosmology*. Ed. by Lemoine, M., Martin, J., and Peter, P. Vol. 738. Lecture Notes in Physics. Springer-Verlag, 2008. Chap. 1, p. 12.

[18] Fixsen, D. J. "The Temperature of the Cosmic Microwave Background". In: *The Astrophysical Journal* 707 (Dec. 2009), pp. 916–920.

[19] Hu, W. and White, M. “A CMB polarization primer”. In: *New Astronomy* 2 (Oct. 1997), pp. 323–344.

[20] Zaldarriaga, M. “The Polarization of the Cosmic Microwave Background”. In: *Measuring and Modeling the Universe*. Ed. by Freedman, W. L. Vol. 2. Carnegie Observatories Astrophysics Series. Cambridge University Press, 2004, p. 309.

[21] Zaldarriaga, M. and Seljak, U. “All-sky analysis of polarization in the microwave background”. In: *Physical Review D* 55 (Feb. 1997), pp. 1830–1840.

[22] *Planck* Collaboration. “*Planck* early results. I. The *Planck* mission”. In: *Astronomy & Astrophysics* 536, A1 (Dec. 2011).

[23] *Planck* Collaboration. “*Planck* 2013 results. XV. CMB power spectra and likelihood”. In: *Astronomy & Astrophysics* 571, A15 (Nov. 2014).

[24] Smoot, G. F. et al. “Structure in the COBE Differential Microwave Radiometer First-year Maps”. In: *The Astrophysical Journal Letters* 396 (Sept. 1992), pp. L1–L5.

[25] Netterfield, C. B. et al. “A Measurement by BOOMERanG of Multiple Peaks in the Angular Power Spectrum of the Cosmic Microwave Background”. In: *The Astrophysical Journal* 571 (June 2002), pp. 604–614.

[26] Hinshaw, G. et al. “Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results”. In: *The Astrophysical Journal Supplement Series* 208, 19 (Oct. 2013).

[27] Sievers, J. L. et al. “The Atacama Cosmology Telescope: cosmological parameters from three seasons of data”. In: *Journal of Cosmology and Astroparticle Physics* 2013.10, 060 (Oct. 2013).

[28] George, E. M. et al. “A Measurement of Secondary Cosmic Microwave Background Anisotropies from the 2500 Square-degree SPT-SZ Survey”. In: *The Astrophysical Journal* 799, 177 (Feb. 2015).

[29] Leitch, E. M. et al. "Measurement of polarization with the Degree Angular Scale Interferometer". In: *Nature* 420 (Dec. 2002), pp. 763–771.

[30] Chiang, H. C. et al. "Measurement of Cosmic Microwave Background Polarization Power Spectra from Two Years of BICEP Data". In: *The Astrophysical Journal* 711 (Mar. 2010), pp. 1123–1140.

[31] BICEP2 Collaboration. "Detection of $B$-mode Polarization at Degree Angular Scales by BICEP2". In: *Physical Review Letters* 112, 241101 (June 2014).

[32] Flauger, R., Hill, J. C., and Spergel, D. N. "Toward an understanding of foreground emission in the BICEP2 region". In: *Journal of Cosmology and Astroparticle Physics* 2014.08, 039 (Aug. 2014).

[33] BICEP2/*Keck* and *Planck* Collaborations. "Joint Analysis of BICEP2/*Keck Array* and *Planck* Data". In: *Physical Review Letters* 114, 101301 (Mar. 2015).

[34] Hanson, D. et al. "Detection of $B$-Mode Polarization in the Cosmic Microwave Background with Data from the South Pole Telescope". In: *Physical Review Letters* 111, 141301 (Sept. 2013).

[35] The POLARBEAR Collaboration. "A Measurement of the Cosmic Microwave Background $B$-mode Polarization Power Spectrum at Sub-degree Scales with POLARBEAR". In: *The Astrophysical Journal* 794 (Oct. 2014), p. 171.

[36] Chinone Y., POLARBEAR Collaboration. via Kermish, Z., private communication.

[37] *Planck* Collaboration. "*Planck* intermediate results. XXX. The angular power spectrum of polarized dust emission at intermediate and high Galactic latitudes". In: *ArXiv e-prints* (Sept. 2014). arXiv: 1409.5738 `[astro-ph.CO]`.

[38] *Planck* Collaboration. "*Planck* intermediate results. XXII. Frequency dependence of thermal emission from Galactic dust in intensity and polarization". In: *Astronomy & Astrophysics* 576, A107 (Apr. 2015).

[39] Vieregg, A. G. et al. “Optical characterization of the *Keck Array* polarimeter at the South Pole”. In: *Millimeter, Submillimeter and Far-Infrared Detectors and Instrumentation for Astronomy VI*. Vol. 8452. Proceedings of the SPIE. Sept. 2012.

[40] MacTavish, C. J. et al. “SPIDER Optimization: Probing the Systematics of a Large-Scale $B$-Mode Experiment”. In: *The Astrophysical Journal* 689 (Dec. 2008), pp. 655–665.

[41] O’Dea, D. T. et al. “SPIDER Optimization. II. Optical, Magnetic, and Foreground Effects”. In: *The Astrophysical Journal* 738, 63 (Sept. 2011).

[42] Fraisse, A. A. et al. “SPIDER: probing the early Universe with a suborbital polarimeter”. In: *Journal of Cosmology and Astroparticle Physics* 2013.04, 047 (Apr. 2013).

[43] Farhang, M. et al. “Primordial Gravitational Wave Detectability with Deep Small-sky Cosmic Microwave Background Experiments”. In: *The Astrophysical Journal* 771, 12 (July 2013).

[44] Palmer, C. E. “The Stratospheric Polar Vortex in Winter”. In: *Journal of Geophysical Research* 64, 7 (1959).

[45] Guðmundsson, J. E. *Probing Early Universe Cosmologies with* SPIDER *& Planck HFI*. PhD thesis. Princeton University, 2014.

[46] NASA Columbia Scientific Balloon Facility. *“Scientific Balloons”*. http://www.csbf.nasa.gov/balloons.html. Accessed: 2015-05-27.

[47] Soler, J. D. *In Search of an Imprint of Magnetization in the Balloon-borne Observations of the Polarized Dust Emission from Molecular Clouds*. PhD thesis. University of Toronto, 2013.

[48] Hecht, E. “Note on an Operational Definition of the Stokes Parameters”. In: *American Journal of Physics* 38 (Sept. 1970), pp. 1156–1158.

[49] Shariff, J. A. et al. "Pointing control for the SPIDER balloon-borne telescope". In: *Ground-based and Airborne Telescopes V*. Vol. 9145. Proceedings of the SPIE. June 2014.

[50] Runyan, M. C. et al. "Design and performance of the SPIDER instrument". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*. Vol. 7741. Proceedings of the SPIE. July 2010.

[51] Rahlin, A. S. et al. "Pre-flight integration and characterization of the SPIDER balloon-borne telescope". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII*. Vol. 9153. Proceedings of the SPIE. June 2014.

[52] Takahashi, Y. D. et al. "Characterization of the BICEP Telescope for High-precision Cosmic Microwave Background Polarimetry". In: *The Astrophysical Journal* 711 (Mar. 2010), pp. 1141–1156.

[53] Aikin, R. W. et al. "Optical performance of the BICEP2 Telescope at the South Pole". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*. Vol. 7741. Proceedings of the SPIE. July 2010.

[54] Ade, P. A. R. et al. "A review of metal mesh filters". In: *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III*. Vol. 6275. Proceedings of the SPIE. June 2006.

[55] Filippini, J. P. et al. "SPIDER: a balloon-borne CMB polarimeter for large angular scales". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*. Vol. 7741. Proceedings of the SPIE. July 2010.

[56] Kuo, C. L. et al. "Antenna-coupled TES bolometer arrays for CMB polarimetry". In: *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy IV*. Vol. 7020. Proceedings of the SPIE. Aug. 2008.

[57] Battistelli, E. S. et al. “Automated SQUID tuning procedure for kilo-pixel arrays of TES bolometers on the Atacama Cosmology Telescope”. In: *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy IV*. Vol. 7020. Proceedings of the SPIE. Aug. 2008.

[58] Bryan, S. A. et al. “Modeling and characterization of the SPIDER half-wave plate”. In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*. Vol. 7741. Proceedings of the SPIE. July 2010.

[59] Bryan, S. A. *Half-Wave Plates for the* SPIDER *Cosmic Microwave Background Polarimeter*. PhD thesis. Case Western Reserve University, 2014.

[60] Guðmundsson, J. E. et al. “Thermal architecture for the SPIDER flight cryostat”. In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*. Vol. 7741. Proceedings of the SPIE. July 2010.

[61] Wiebe, D. V. *BLAST: A Balloon-borne, Large-Aperture, Submillimetre Telescope*. PhD thesis. University of Toronto, 2008.

[62] Benton, S. J. et al. “BLASTbus electronics: general-purpose readout and control for balloon-borne experiments”. In: *Ground-based and Airborne Telescopes V*. Vol. 9145. Proceedings of the SPIE. June 2014.

[63] Benton, S. J. *Mapping Submillimetre Polarization with BLASTPol*. PhD thesis. University of Toronto, 2014.

[64] Battistelli, E. S. et al. “Functional Description of Read-out Electronics for Time-Domain Multiplexed Bolometers for Millimeter and Sub-millimeter Astronomy”. In: *Journal of Low Temperature Physics* 151 (May 2008), pp. 908–914.

[65] Soler, J. D. et al. “Design and construction of a carbon fibre gondola for the SPIDER balloon-borne telescope”. In: *Ground-based and Airborne Telescopes V*. Vol. 9145. Proceedings of the SPIE. June 2014.

[66] ASTM E927 - 10. *Standard Specification for Solar Simulation for Photovoltaic Testing*. http://www.astm.org/doiLink.cgi?E927. Accessed: 2015-04-01.

[67] Doerffel, D. and Abu-Sharkh, S. "A critical review of using the Peukert equation for determining the remaining capacity of lead-acid and lithium-ion batteries". In: *Journal of Power Sources* 155 (Apr. 2006), pp. 395–400.

[68] Fissel, L. M. *Probing the Role Played by Magnetic Fields in Star Formation with BLASTPol*. PhD thesis. University of Toronto, 2013.

[69] Gandilo, N. N. et al. "Attitude determination for balloon-borne experiments". In: *Ground-based and Airborne Telescopes V*. Vol. 9145. Proceedings of the SPIE. June 2014.

[70] Górski, K. M. et al. "HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere". In: *The Astrophysical Journal* 622 (Apr. 2005), pp. 759–771.

[71] Rahlin, A. S. *Toward the First Flight of the* SPIDER *Balloon-Borne Telescope*. PhD thesis. Princeton University, 2015.

[72] Pascale, E. et al. "The Balloon-borne Large Aperture Submillimeter Telescope: BLAST". In: *The Astrophysical Journal* 681 (July 2008), pp. 400–414.

[73] Griffin, M. J. et al. "The *Herschel*-SPIRE instrument and its in-flight performance". In: *Astronomy & Astrophysics* 518 (July 2010), p. L3.

[74] Netterfield, C. B. et al. "BLAST: The Mass Function, Lifetimes, and Properties of Intermediate Mass Cores from a 50 deg$^2$ Submillimeter Galactic Survey in Vela ($\ell \approx 265^\circ$)". In: *The Astrophysical Journal* 707 (Dec. 2009), pp. 1824–1835.

[75] Wiebe, D. V. et al. "BLAST Observations of Resolved Galaxies: Temperature Profiles and the Effect of Active Galactic Nuclei on FIR to Submillimeter Emission". In: *The Astrophysical Journal* 707 (Dec. 2009), pp. 1809–1823.

[76] Devlin, M. J. et al. "Over half of the far-infrared background light comes from galaxies at $z \geq 1.2$". In: *Nature* 458 (Apr. 2009), pp. 737–739.

[77] Marsden, G. et al. "BLAST: Resolving the Cosmic Submillimeter Background". In: *The Astrophysical Journal* 707 (Dec. 2009), pp. 1729–1739.

[78] Pascale, E. et al. "BLAST: A Far-Infrared Measurement of the History of Star Formation". In: *The Astrophysical Journal* 707 (Dec. 2009), pp. 1740–1749.

[79] Viero, M. P. et al. "BLAST: Correlations in the Cosmic Far-Infrared Background at 250, 350, and 500 μm Reveal Clustering of Star-forming Galaxies". In: *The Astrophysical Journal* 707 (Dec. 2009), pp. 1766–1778.

[80] *BLAST! Balloon-borne Large Aperture Submillimeter Telescope*. (Documentary Film). DevlinPix Productions. 2008.

[81] Fissel, L. M. et al. "The balloon-borne large-aperture submillimeter telescope for polarimetry: BLAST-Pol". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*. Vol. 7741. Proceedings of the SPIE. July 2010.

[82] Pascale, E. et al. "The balloon-borne large-aperture submillimeter telescope for polarimetry-BLASTPol: performance and results from the 2010 Antarctic flight". In: *Ground-based and Airborne Telescopes IV*. Vol. 8444. Proceedings of the SPIE. Sept. 2012.

[83] Galitzki, N. et al. "The Balloon-borne Large Aperture Submillimeter Telescope for Polarimetry-BLASTPol: performance and results from the 2012 Antarctic flight". In: *Ground-based and Airborne Telescopes V*. Vol. 9145. Proceedings of the SPIE. July 2014.

[84] Angilè, F. E. et al. "The Balloon-borne Large Aperture Submillimeter Telescope for Polarimetry (BLASTPol) 2012: Instrument and Flight Performance". *In Preparation*. 2014.

[85] Hall, J. S. “Observations of the Polarized Light From Stars”. In: *Science* 109 (Feb. 1949), pp. 166–167.

[86] Hiltner, W. A. “On the Presence of Polarization in the Continuous Radiation of Stars. II”. In: *The Astrophysical Journal* 109 (Feb. 1949), pp. 471–478.

[87] Hildebrand, R. H. et al. “A Primer on Far-Infrared Polarimetry”. In: *Publications of the Astronomical Society of the Pacific* 112 (Sept. 2000), pp. 1215–1235.

[88] Lazarian, A. “Tracing Magnetic Fields with Aligned Grains”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 106 (July 2007), pp. 225–256.

[89] Crutcher, R. M. “Magnetic Fields in Molecular Clouds”. In: *The Annual Review of Astronomy and Astrophysics* 50 (May 2012), pp. 29–63.

[90] Mouschovias, T. C. and Ciolek, G. E. “Magnetic Fields and Star Formation: A Theory Reaching Adulthood”. In: *The Origin of Stars and Planetary Systems.* Ed. by Lada, C. J. and Kylafis, N. D. Vol. 540. NATO Science Series C. Springer Netherlands, 1999, pp. 305–340.

[91] Mac Low, M.-M. and Klessen, R. S. “Control of star formation by supersonic turbulence”. In: *Reviews of Modern Physics* 76 (Jan. 2004), pp. 125–194.

[92] Moncelsi, L. et al. “Empirical modelling of the BLASTPol achromatic half-wave plate for precision submillimetre polarimetry”. In: *Monthly Notices of the Royal Astronomical Society* 437 (Jan. 2014), pp. 2772–2789.

[93] Bock, J. J. et al. “Silicon nitride micromesh bolometer arrays for SPIRE”. In: *Advanced Technology MMW, Radio, and Terahertz Telescopes*. Vol. 3357. Proceedings of the SPIE. July 1998.

[94] Rex, M. et al. “BLAST autonomous daytime star cameras”. In: *Ground-based and Airborne Instrumentation for Astronomy.* Vol. 6269. Proceedings of the SPIE. June 2006.

[95] Markley, F. L. "Attitude Error Representations for Kalman Filtering". In: *Journal of Guidance, Control, and Dynamics* 26.2 (Mar. 2003), pp. 311–317.

[96] Harman, R. R. *Wilkinson Microwave Anisotropy Probe (WMAP) Attitude Estimation Filter Comparison*. Tech. rep. 20060002447. Greenbelt, Maryland, USA: NASA Goddard Space Flight Center, Jan. 2005. URL: http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20060002447.pdf.

[97] Gandilo, N. N. *Probing Interstellar Grain Alignment with Balloon-borne Submillimeter Observations*. PhD thesis. University of Toronto, 2015.

[98] Matthews, T. G. et al. "Lupus I Observations from the 2010 Flight of the Balloon-borne Large Aperture Submillimeter Telescope for Polarimetry". In: *The Astrophysical Journal* 784 (Apr. 2014), p. 116.

[99] Kisner, T. "Time Ordered Astrophysics Scalable Tools (TOAST)". https://theodorekisner.com/software/toast/index.html. Accessed: 2014-09-16.

[100] Allen, D. A. and Hillier, D. J. "The Shape of the Homunculus Nebula around $\eta$ Carinae". In: *Proceedings of the Astronomical Society of Australia* 10 (1993).

[101] Smith, N. "The Structure of the Homunculus. I. Shape and Latitude Dependence from $H_2$ and [Fe II] Velocity Maps of $\eta$ Carinae". In: *The Astrophysical Journal* 644 (June 2006), pp. 1151–1163.

[102] Smith, N. and Brooks, K. J. "The Carina Nebula: A Laboratory for Feedback and Triggered Star Formation". In: *Handbook of Star Forming Regions, Volume II*. Ed. by Reipurth, B. Astronomical Society of the Pacific, Dec. 2008.

[103] Smith, N. et al. "*Spitzer Space Telescope* observations of the Carina nebula: the steady march of feedback-driven star formation". In: *Monthly Notices of the Royal Astronomical Society* 406 (Aug. 2010), pp. 952–974.

[104] Li, H. et al. "Results of SPARO 2003: Mapping Magnetic Fields in Giant Molecular Clouds". In: *The Astrophysical Journal* 648 (Sept. 2006), pp. 340–354.

[105] Lucy, L. B. “An iterative technique for the rectification of observed distributions”. In: *The Astronomical Journal* 79 (June 1974), pp. 745–754.

[106] Richardson, W. H. “Bayesian-Based Iterative Method of Image Restoration”. In: *Journal of the Optical Society of America* 62 (Jan. 1972), pp. 55–59.

[107] Vaillancourt, J. E. and Matthews, B. C. “Submillimeter Polarization of Galactic Clouds: A Comparison of 350 μm and 850 μm Data”. In: *The Astrophysical Journal Supplement Series* 201, 13 (Aug. 2012).

[108] Vaillancourt, J. E. et al. “New Results on the Submillimeter Polarization Spectrum of the Orion Molecular Cloud”. In: *The Astrophysical Journal Letters* 679 (May 2008), pp. L25–L28.

[109] Hildebrand, R. H. et al. “The Far-Infrared Polarization Spectrum: First Results and Analysis”. In: *The Astrophysical Journal* 516 (May 1999), pp. 834–842.

[110] Zeng, L. et al. “The Submillimeter Polarization Spectrum of M17”. In: *The Astrophysical Journal* 773, 29 (Aug. 2013).