Probing Interstellar Grain Alignment with Balloon-borne Submillimeter Observations

by

Natalie Gandilo

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy

Graduate Department of Astronomy and Astrophysics
University of Toronto

Abstract

Probing Interstellar Grain Alignment with Balloon-borne Submillimeter Observations

Natalie Gandilo
Doctor of Philosophy
Graduate Department of Astronomy and Astrophysics
University of Toronto
2015

The role that magnetic fields play in regulating star formation is one of the outstanding issues in star formation theory. Magnetic fields in star-forming clouds in our galaxy can be observed by tracing the polarized emission from interstellar dust grains. However the mechanism by which the dust grains align with the magnetic fields is not fully understood. Grain alignment theories (e.g. radiative alignment torques) make predictions for the observed spectrum of the polarized emission, so observations at multiple wavelengths can be used to test these theories.

The Balloon-borne Large Aperture Submillimeter Telescope for Polarimetry (BLASTPol) observes polarized dust at 250μm, 350μm and 500μm while suspended from a balloon in the stratosphere above Antarctica. BLASTPol maps the dust polarization signal over entire molecular clouds, with enough angular resolution to trace the cloud sub-structures and protostellar cores. BLASTPol uses an attitude reconstruction system consisting of multiple sensors, including star cameras, sun sensors and rate gyroscopes. This system allows BLASTPol to point in-flight at specific regions on the sky, and allows the telescope's attitude to be reconstructed post-flight. A similar system was built for SPIDER, a balloon-borne telescope that observes the polarization of the CMB.

The analysis pipeline used to produce science maps is discussed, and science results are presented from BLASTPol's 2012 observations of the Vela C molecular cloud. The polarization spectrum shows a minimum at 350μm, similar to the measurements of previous experiments observing other molecular clouds. No strong correlation is seen between he shape of the polarization spectrum and either the temperature or density of the dust. Analysis of the maps is ongoing, and future work will focus on the diffuse dust in the cloud, which is more suitable to compare with dust models.

Acknowledgements

Figure 1: Ever notice how you come across somebody once in a while that you shouldn't have messed with? That's me.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Figure 1.1: You're a legend in your own mind.

1.1 The Role of Magnetic Fields in Star Formation

One of the key outstanding questions in our understanding of star formation is whether magnetic fields provide support against the gravitational collapse of molecular clouds and cores. Star formation has been observed to occur at a slower rate than expected, given the timescale for gravitational collapse (Padoan et al. 2014). It is still unknown what role magnetic fields play in halting the collapse, and some argue that turbulence plays a much more dominant role (Mac Low & Klessen 2004). As will be described in Section 1.3, submillimeter polarimetry can be used to map the magnetic fields in star forming clouds, in order to address this question.

The Balloon-borne Large Aperture Submillimeter Telescope for Polarimetry (BLASTPol) was designed to make maps of magnetic fields over entire molecular clouds, with enough angular resolution to trace the fields into dense filaments and star-forming cores. The maps can be used to study whether the direction of magnetic field lines is correlated with the direction of filaments and other sub-structures in the cloud (Ostriker et al. 2001; Soler et al. 2013). Chan-

drasekhar & Fermi (1953) suggested that the dispersion in the magnetic field angles can be used to estimate the magnetic field strength, as a weaker magnetic field will be more tangled by turbulent motions in the gas. High-resolution observations can probe the magnetic field at the scale of a collapsing core, looking for the hourglass-shaped field aligned with an oblate core that would result from magnetically-regulated star formation (Girart et al. 2009).

1.2 Polarized Galactic Dust Foregrounds

BLASTPol observations can also be used to characterize the polarized dust signal in order to better improve models of the dust itself. These models are especially important to experiments searching for B-mode polarization of the Cosmic Microwave Background (CMB), a signal predicted to result from gravity waves during inflation. The B-mode signal is extremely faint, and the polarized signal from galactic dust is the strongest foreground contaminant that needs to be characterized and removed. BLASTPol does not have the mapping speed to obtain high signal-to-noise observations of large regions of the sky mapped by most CMB experiments, where the dust signal is the faintest. Such an observation would have taken up too much of BLASTPol's limited observing time. However the BLASTPol maps targeting molecular clouds do extend beyond the clouds to dust at low extinction ($A_v \leq 3$), where comparisons can be made with theoretical predictions of diffuse dust.

1.3 Polarized Dust Emission

Submillimeter polarimetry is able to trace magnetic fields because the direction of polarization is related to the direction of the magnetic field lines. When there is a net alignment of a population of interstellar dust grains in the direction of local magnetic field lines, this leads to a net polarization. This polarization is observed in the light coming from background stars, as well as in the thermal radiation emitted by the dust. For a dust temperature of 10-20K, this thermal radiation peaks at submillimeter wavelengths. Figure 1.2 shows what happens when the short axis of the grain is aligned parallel to the magnetic field line. In this case, background star light is preferentially extincted by the long axis of the grain, so the observed polarization is aligned parallel to the direction of the local magnetic field lines in the plane of the sky. Conversely, the dust grains radiate most effectively when the E-field is parallel to the long axis of the grain (Draine 2003), so the polarization direction of the thermal emission is perpendicular to the magnetic field direction in the plane of the sky. In this way, polarimetric observations can span ultraviolet to millimeter wavelengths, tracing the direction of magnetic field lines from the outskirts of molecular clouds into regions of high visual extinction.

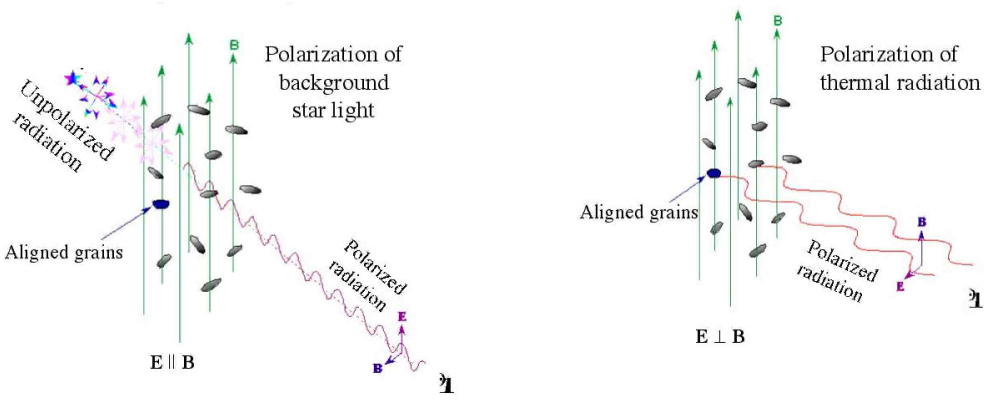

Figure 1.2: The alignment of dust grains with the local magnetic field causes linear polarization in the background star light (left) and in the thermal dust radiation (right).

1.4 Grain Alignment Theory

Although the effect described in the previous section has been observed, the physical mechanism by which the dust grains align in such a manner with the magnetic field lines is still not fully understood. In order to use observations of polarized dust emission to infer magnetic field direction, this alignment mechanism must be well characterized.

The physics of dust grain alignment has been a long-standing question, and various theories have been put forward. The currently favoured theory proposes that alignment occurs through radiative torques (RATs) acting on helical dust grains. RAT theory was first proposed by Dolginov & Mitrofanov (1976), who suggested that helical grains would scatter left- and right-polarized light differently, resulting in the grains being spun up by unpolarized light. This theory was not fully embraced until numerical calculations (Draine & Weingartner 1996, 1997) showed irregularly shaped grains getting spun up by RATs. An analytical theory of RATs was then introduced by Lazarian & Hoang (2007). RAT theory can be observationally tested, because it makes clear predictions of the alignment efficiency that will occur depending on the particular radiation field that the dust experiences.

1.4.1 Polarization Spectrum

Because BLASTPol is a multi-wavelength experiment, it can measure the polarization spectrum, i.e. percentage polarization as a function of wavelength. Figure 6.8 shows the observed polarization spectrum (from far-infrared to millimeter wavelengths) from the combined measurements of different experiments (Vaillancourt & Matthews 2012; Zeng et al. 2013). The observations reveal a spectrum that rises longward of 350μm, but falls at shorter wavelengths.

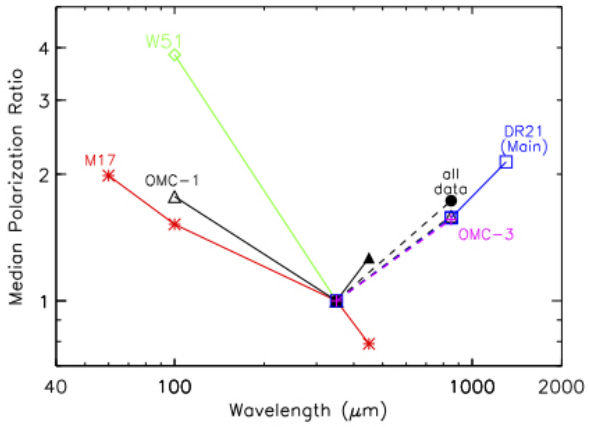

Figure 1.3: Median polarization ratio normalized to 350μm. Different colours represent measurements from different experiments, with labels indicating the molecular cloud that was observed (from Vaillancourt et al. (2013)).

Such a shape cannot be explained by isothermal dust, and indicates dust at multiple temperatures (Hildebrand et al. 1999).

Models of dust grains predict a spectrum that rises with wavelength, with the slope depending on the shape and composition of the aligned grains (Draine & Fraisse 2009; Bethell et al. 2007). However these models pertain to diffuse dust where the physical conditions are very different from bright dense molecular clouds. It has been suggested that the falling part of the spectrum (from $60 - 350\mu$m) comes from warmer dust near embedded sources in the observed molecular clouds. These grains would be highly irradiated and therefore experience a high degree of RAT efficiency.

Because BLASTPol's observing bands bracket the 350μm minimum, it is the first time that the shape of the polarization spectrum can be verified by a single experiment. In addition, because BLASTPol maps encompass entire molecular clouds, the data can test RAT theory by analyzing how the polarization spectrum varies in different radiative environments.

1.5 Why BLASTPol

BLASTPol solves the problems caused by the Earth's atmosphere by employing a high altitude balloon platform. Ground-based observations at submillimeter wavelengths are affected by the fact that the atmosphere absorbs submillimeter light, and varies on large angular scales. This limits ground experiments to observing in narrow wavelength bands, to avoid atmospheric emission lines, and restricts the observations to small angular scales.

The ideal place to put a submillimeter telescope is in space, but a cheaper and faster alternative is to suspend it from a stratospheric balloon. This allows observations to be done above

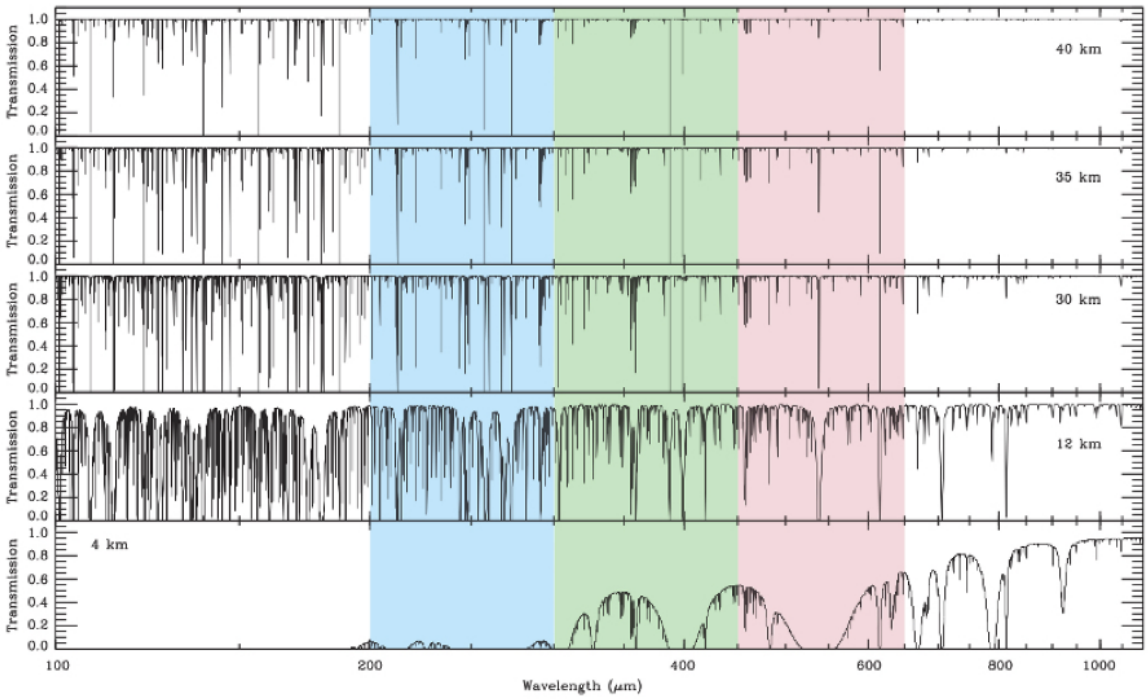

Figure 1.4: Atmospheric transmission versus wavelength for various altitudes from the ground up to 40km. Highlighted in colour are the BLASTPol observing bands. BLASTPol observes at an altitude of $\sim$ 38km (from Marsden (2006)).

99.5% of the atmosphere in a near-space environment, at a fraction of the cost of a satellite mission. Figure 1.4 shows how the atmospheric transmission increases with increasing altitude. At typical balloon float altitude of $\sim$ 38km, the transmission is nearly 100% across most of the wavelength range observed by BLASTPol (from 250μm-650μm).

1.6 My thesis work

My thesis work encompassed the construction, integration, and deployment of BLASTPol for both the 2010 and 2012 flights. Chapter 2 describes the BLASTPol instrument. The primary focus of my work was in building the attitude reconstruction system (Sections 2.7 and 2.8). The systems I built for BLASTPol were used on another telescope, Spider, which flew in 2014 and is described in Chapter 3. Since Spider did not have a test flight, the 2010 and 2012 flights served as opportunities to test its attitude reconstruction system. Chapter 4 describes the flight performance of these systems.

After the flight, I worked on reduction and analysis of the data, focusing in particular on the pointing reconstruction, described in Chapter 5. Chapter 6 presents the results of using BLASTPol data to investigate the polarization spectrum in the Vela C molecular cloud.

Chapter 2

BLASTPol Instrument

Figure 2.1: There's money to be made in a place like this.

BLASTPol is a 1.8-m telescope that observes in three submillimeter wavelengths - 250μm, 350μm and 500μm, with diffration-limited resolutions of $36''$, $42''$ and $60''$ respectively. BLASTPol uses detector arrays that are prototypes of those used in the SPIRE instrument of the Herschel satellite mission. BLASTPol observes while suspended from a stratospheric balloon at an altitude of ~ 40 km, which puts it above $\sim 99.5\%$ of the Earth's atmosphere.

The BLASTPol telescope has had two Long Duration Balloon (LDB) flights, both from McMurdo Station, Antarctica, in 2010 and 2012. BLASTPol is a reconstruction of the BLAST telescope (Pascale et al. 2008), which had two science flights from Sweden in 2005 and Antarctica in 2006. At the termination of the 2006 flight, the telescope's parachute failed to detach from the payload upon reaching the ground, causing BLAST to be dragged 200km and resulting in damage to the telescope. However the important parts of the telescope (detectors, receiver, primary and secondary mirrors) were recovered, allowing the telescope to be re-built with the added capability of polarimetry.

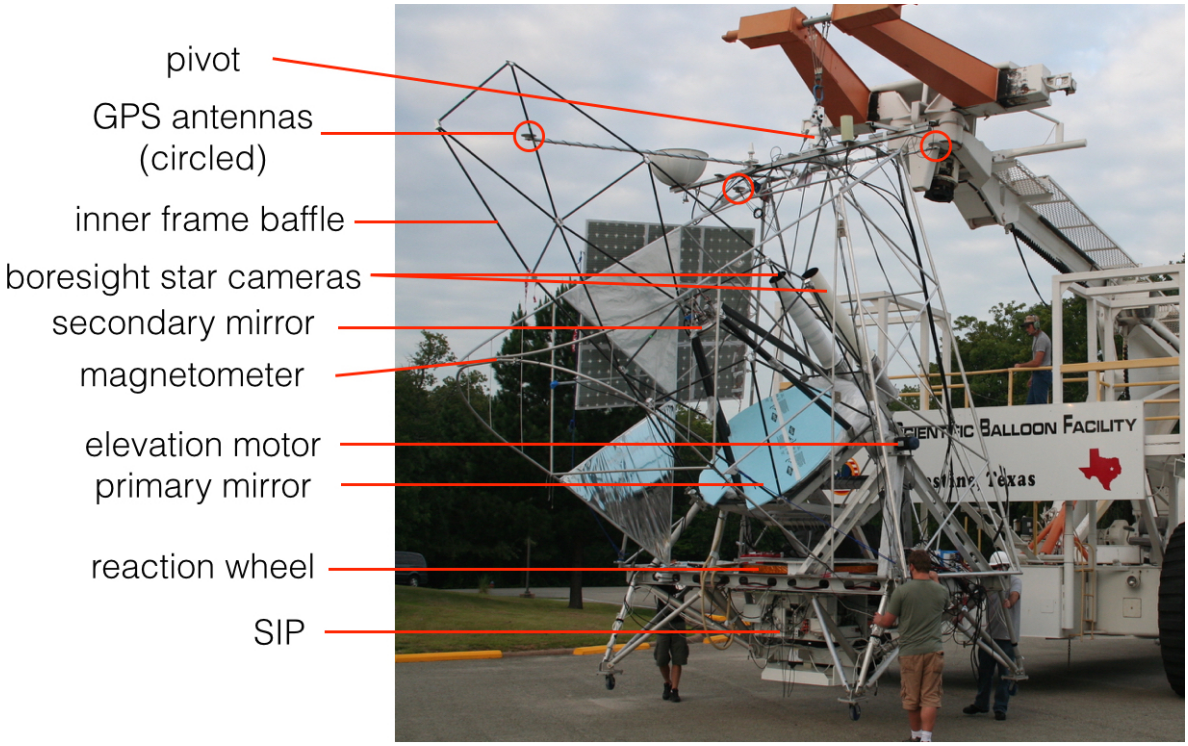

Figure 2.2: BLASTPol on the launch vehicle during integration in Palestine, Texas.

2.1 Gondola

The BLASTPol gondola consists of an outer frame, which connects to the balloon and supports the inner frame and sunshields, an inner frame, which holds the optics and receiver, and sunshields, which prevent the sun from illuminating the telescope. The gondola is capable of pointing the telescope anywhere on the sky by moving the outer frame in azimuth, and the inner frame in elevation.

2.1.1 Outer Frame

The outer frame structure consists of several parts. A rectangular frame holds a floor that is used for mounting the attitude control electronics. Each side of the rectangle holds four beams pointing upwards in a pyramid shape, and the inner frame is mounted on top of these pyramids. Below the floor of the outer frame is mounted the reaction wheel. Each corner of the outer frame has legs to which wheels are attached when BLASTPol is on the ground. Before launch, these are removed and the CSBF solar panels are attached. The CSBF solar panel array is normally a four sided array, to assure operation in the event of the failure of the gondola pointing system. The CSBF Support Instrumentation Package (SIP) is held within beams extending from the floor of the outer frame in a structure called the SIP cage. Before

launch, a ballast hopper is attached to the bottom of the SIP cage. The outer frame is suspended from each corner by steel cables attaching to the pivot.

2.1.2 Inner Frame

The inner frame of the gondola consists of aluminum beams in the shape of a hexagonal box. The inner frame connects to the outer frame at two points, allowing it to move in elevation. At one attachment point the elevation motor is mounted. At the other point is a hollow tube through which cables from inner frame components can pass through to the attitude control system on the outer frame. The primary mirror is mounted into the front hexagonal face of the box. Carbon fiber struts are attached to points of the inner frame around the primary mirror and project forward in a pyramidal shape. The secondary mirror and its actuators are mounted at the top of the pyramid. When the mirror faces forward (ie. points at zero elevation), the cryostat is lowered into the space behind the primary mirror. The inner frame also provides the mounting point for the two boresight star cameras, which sit the top of the inner frame, above the primary mirror. The balancing system, consisting of two tanks connected by a hose and pump, is also mounted to the inner frame. The inner frame also holds the asymmetrical carbon-fiber baffle that was designed to allow BLASTPol to point within 40 degrees of the Sun. The base of the baffle attaches to the inner frame at six tabs around the primary mirror, and extends forward along the boresight by 7 metres.

2.1.3 Sunshield

The sunshield is a large structure of aluminum tubing that mounts to the edges of the floor of the outer frame and surrounds the telescope on all sides other than the front. The top of the sunshield has a hexcel shelf in the middle, upon which the pivot can rest when the telescope is not being suspended. The top of the sunshield is also the attachment point for GPS antennas, and for the CSBF antenna boom, upon which CSBF mounts all of the antennas used to communicate with BLASTPol in flight. The solar panels attach to the sunshields on the sun-facing side. Before launch the sunshields are covered on the sun-facing side by inserting panels of blue foam into the space between the tubes, and covering with aluminized mylar.

2.2 Optics

BLASTPol is an on-axis Cassegrain telescope designed to give an angular resolution of $30''$, $42''$, and $60''$ at 250μm, 350μm, and 500μm, respectively. The optical layout is depicted in Figure 2.3. A 1.8m aluminum parabolic primary mirror (M1) collects light from the sky and

focuses it onto a 40 cm secondary mirror (M2). Light then passes through a hole in the primary mirror, and through a window into the cryostat containing the cold optics (M3, M4, and M5). The cold optics are arranged in an Offner-relay configuration. In this configuration, M3, M4, and M5 are all spherical and share a common center of curvature; M3 and M5 are concave while M4 is convex. Two dichroic filters split the light into the three bands. A calibration lamp is placed in the Lyot stop in M4, and briefly pulsed periodically during the flight. This allows changes in detector responsivity to be tracked.

2.2.1 Focus Mechanism

BLASTPol experiences a changing thermal environment when launched into the stratosphere. Although the telescope is focused before launch, thermal effects on the telescope cause the focus to change relative to what it was on the ground. In order to solve this problem, BLASTPol has an active focus mechanism that allows the focus to be adjusted during the flight. The secondary mirror is equipped with three motorized actuators which can be moved in and out relative to the primary mirror, with a precision of 100μm. Simulations have shown that the relative distance between M1 and M2 can deviate by at most $\lambda/2$ at 250μm without significant degradation of the image.

2.3 Detectors

At the focal plane of the telescope are prototype Herschel SPIRE bolometers, with 149, 88 and 43 detectors at 250, 350 and 500μm respectively. The detectors are maintained at 300mK in a liquid helium/nitrogen cryostat with a ~ 12-day hold time.

The detectors are silicon-nitride micromesh ("spider web") bolometers, coupled to smooth-walled conical feedhorn arrays spaced at $2f\lambda$. A waveguide filter at the bottom of the feedhorn defines the long wavelength cutoff in the detector spectral response. A filter placed in front of the feedhorn array defines the short wavelength cutoff for each band. Figure 2.4 shows the spectral response of the BLASTPol bands.

2.4 Polarimetry

Polarimetry is achieved through photo-lithographed polarizing grids mounted in front of each detector array, with alternating horizontal and vertical polarization angles across each row of detectors (Figure 2.5). This allows a measurement of one Stokes parameter (Q or U) to be made at a few Hz as the source is scanned across the array. In addition, an achromatic half-

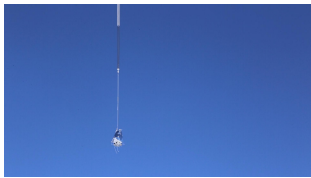

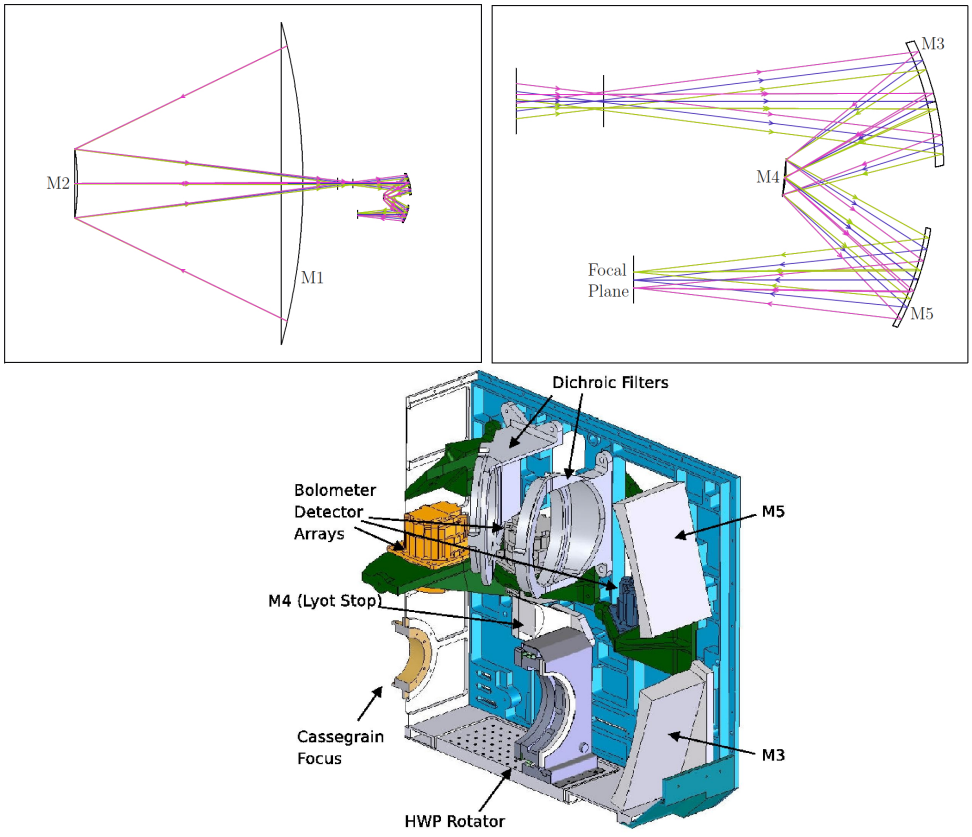

Figure 2.3: Optical layout of BLASTPol. The top figure shows the complete optical system, with the cold optics on the right. The bottom figure shows the cold optics in more detail.

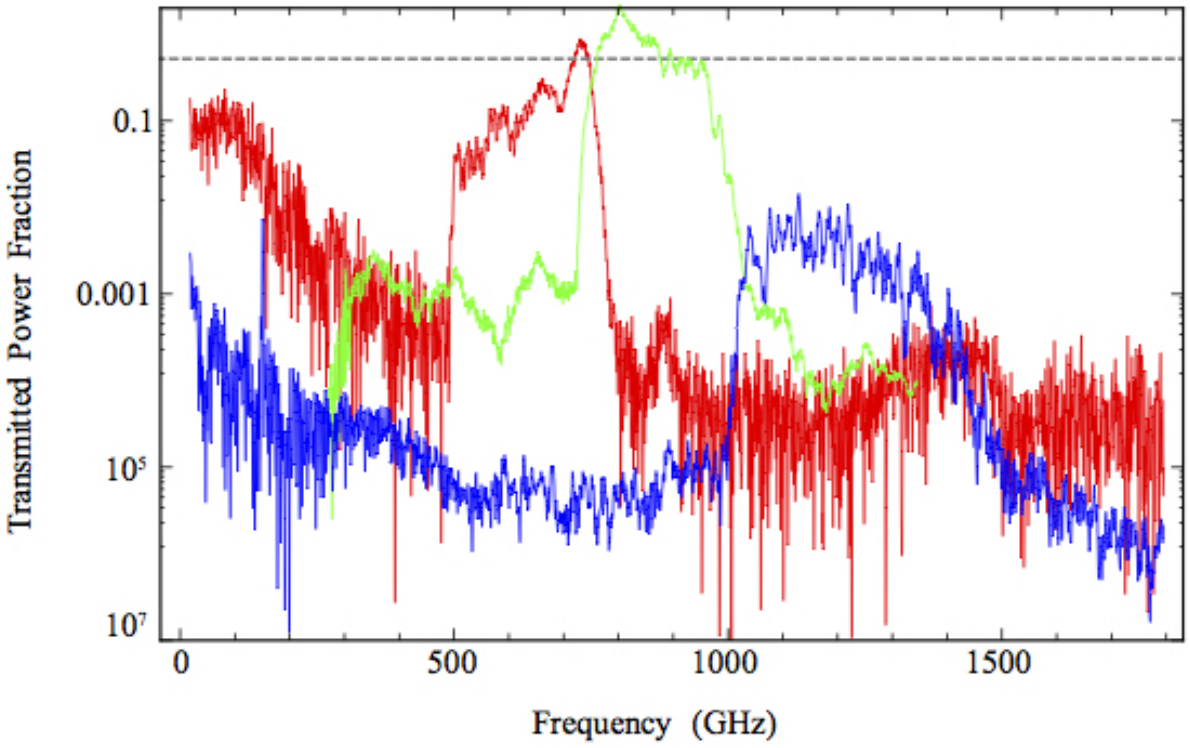

Figure 2.4: Relative spectral response of the three BLASTPol bands (blue=250μm, green=350μm, red=500μm), measured using a Fourier Transform Spectrometer and a nitrogen source.

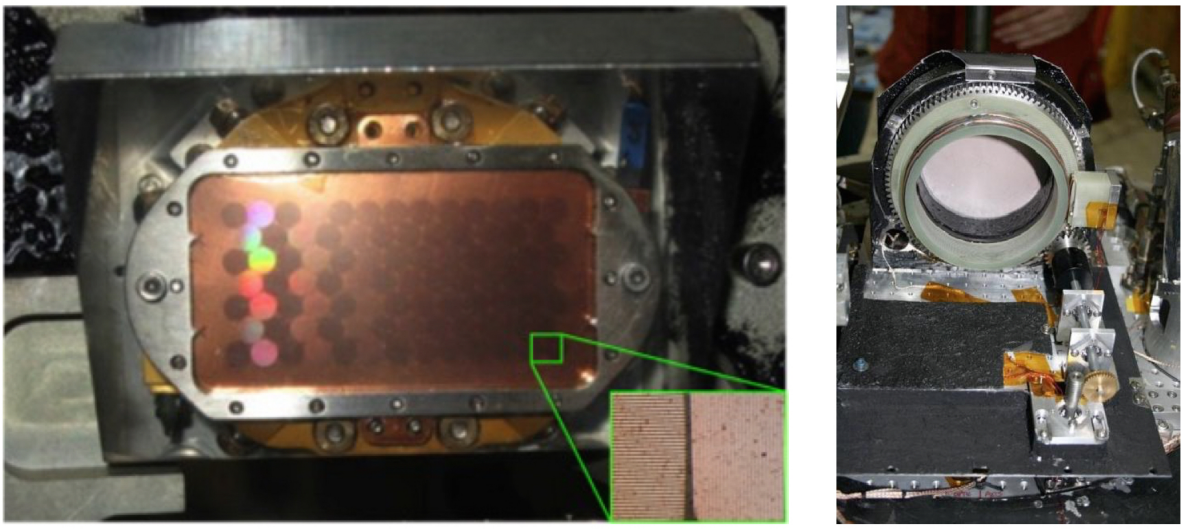

Figure 2.5: Polarimetry. Left: Photo-lithographed olarization grid, with inset showing the grid in front of two adjacent detectors; Right: Half-wave plate and rotation mechanism used to rotate the plate.

wave plate (Moncelsi et al. 2014) is mounted inside the cold optics box and stepped by 22.5° after every scan.

The instrumental polarization (IP), which is the spurious polarization due to the instrument when observing an unpolarized source, was measured pre-flight using an unpolarized submillimeter calibrator. The calibrator consisted of a chopped heated source, placed in front of the receiver window to fill the detector array field of view. The source temperature was varied, and the measurement was repeated at different half-wave plate angles. This measured only the cryostat and cold optics, since Cassegrain telescopes are known to introduce very small levels of IP. The IP calibrator source was also used in parallel with a polarizing grid to measure the polarization efficiency of the instrument, which is the amount of polarization measured when observing a 100% polarized source. For all three bands, the IP was found to be $< 1\%$, and the polarization efficiency was $\sim 80\%$.

2.5 Flight Computers and ACS

BLASTPol is designed to run completely autonomously after launch. All operations are controlled by one of two redundant flight computers. If the in-control flight computer fails, it is rebooted while control automatically switches to the other flight computer. Both flight computers run `mcp` (master control program), which receives data from and sends commands to the various components of the telescope. The flight computers communicate with Data Acquisition System (DAS), which reads out the detector signals, and the Attitude Control System

(ACS), which reads out the gondola systems. Both the DAS and ACS use electronic cards custom designed for BLAST. The cards are stored in crates, with the DAS crate mounted in the space behind the primary mirror and the ACS crate mounted to the gondola floor. Each flight computer writes its own copy of the flight data to a hard drive, which must be recovered after the flight.

2.6 Telemetry

2.6.1 Support Instrumentation Package

On both BLASTPol flights, CSBF provided a Support Instrumentation Package (SIP), which communicates with the BLASTPol flight computers over RS-232 serial to provide forward (commanding) and return (data) telemetry. The SIP is powered from its own array of PV panels, and electrically and thermally isolated from the gondola.

Science Stack

The science stack provides 28 open-collector command outputs, allowing for power switching of 14 BLASTPol subsystems independently of the science instrument. If normal commanding is available and working, the science stack is typically not used. It is only required at the end of the flight during the termination sequence, to power off the BLASTPol flight computers and ACS.

2.6.2 Line-of-Sight

CSBF equips BLASTPol with two L-Band or S-Band transmitters, and directional antennas for communicating with the payload during the period after launch when it is still within line-of-sight (LOS) of the launch site. One transmitter is used to send a NTSC video signal of images of the star camera. As soon as diagnosis of the star camera performance is complete, the video transmitter is powered off and not used again during the flight. The other LOS transmitter is used to send a 1Mb biphase-modulated signal containing data from the flight computers. It is also powered off after LOS.

2.6.3 Over the Horizon

When the payload is beyond view of radio antennas, communication is done through the Tracking and Data Relay Satellite System (TDRSS) and the Iridium satellite network. TDRSS data downlink is available at a continuous 6 kilo-bit rate using the TDRSS omni-directional

antenna. A higher rate, up to 92kbps, is also available through the TDRSS High Gain Antenna (HGA) which can be pointed at available satellites using attitude information provided by a separate GPS array flown by CSBF. The Iridium downlink provides 2kbps in dialup mode, which uses an onboard modem. An additional data stream is available in a 255 byte packet that is transmitted over Iridium every 15 minutes. This data packet is only big enough to contain a small fraction of the total number of BLASTPol data fields.

2.6.4 Ground Station

Commands can be sent from from both the Operations Control Center (OCC) in Palestine, Texas, or from the Remote Operations Control Center (ROCC) at the LDB site in McMurdo. Data is received at the OCC within a few minutes of transmission from the payload. It is then relayed to two redundant BLASTPol ground station computers set up in Palestine. The ground station computer then serves the data over the Internet to be downloaded and viewed by client computers.

2.7 Attitude Control

Azimuthal Control

The BLASTPol motion control system is described in detail in Wiebe (2007) and Fissel (2013), and shown in Figure 2.2. The primary azimuth motor drives a reaction wheel, consisting of an aluminum honeycomb disk with brass disks in the outer perimeter. The reaction wheel is mounted below the floor of the outer frame. The second azimuth motor is an active pivot motor that connects the top of the suspension cables with the balloon flight train and dumps angular momentum to the balloon.

Elevation Control

Elevation control is provided by an elevation motor mounted on the port side of the inner frame (see Figure 2.2). An encoder in the motor measures elevation with sub-arcsecond precision. A lock pin is used to lock the inner frame during launch, and the pin is retracted after launch.

Scan Strategy

Science targets are mapped by defining a region of the sky around the target, and scanning at constant speed back and forth in azimuth ($\sim 0.1 - 0.2°/\mathrm{s}$) while drifting in elevation, in

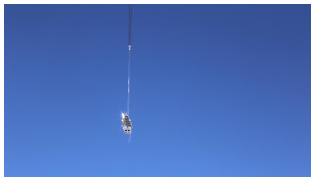

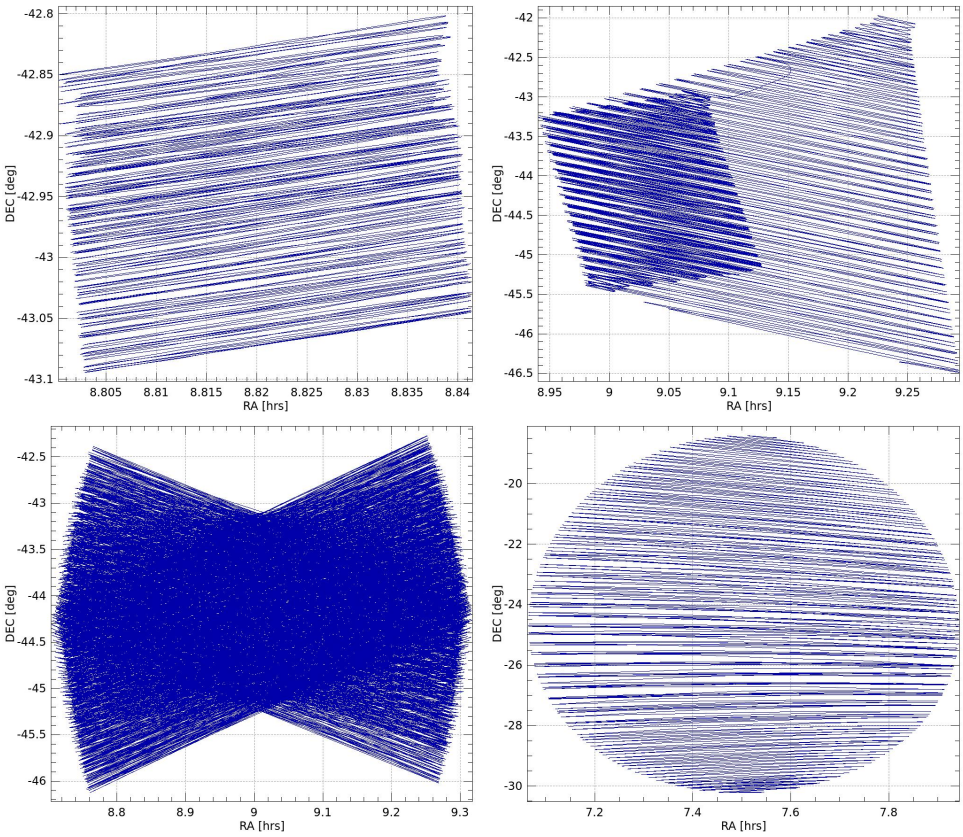

Figure 2.6: Typical scans used by BLASTPol. Top left: box scan; Top right: Extended region (shown over a box scan); Bottom left: bowtie scan; Bottom right: cap scan.

order to fill in the region. Figure 2.6 shows examples of typical BLASTPol scans. Each raster up and down across the target is done at one of the four half-wave plate angles. One "scan" is therefore defined as four up and down rasters, thus covering a complete set of half-wave plate angles. The gondola is also capable of goto modes, in which the telescope boresight remains at a fixed position, specified in equatorial (Right Ascension (RA) and Declination (Dec)) or horizontal (Azimuth and Elevation) coordinates. RA/Dec goto is used, for example, during focusing of the boresight star cameras.

2.8 Pointing Sensors

There are two goals that the pointing sensors are designed to achieve. First, the in-flight pointing needs to be accurate enough to ensure that the detector array observes the desired targets. Second, the post-flight reconstructed pointing solution must be accurate enough to make a map that oversamples the beam of the telescope. BLASTPol has a multitude of pointing sensors that provide absolute attitude solutions, but these are at a much lower sample rate than

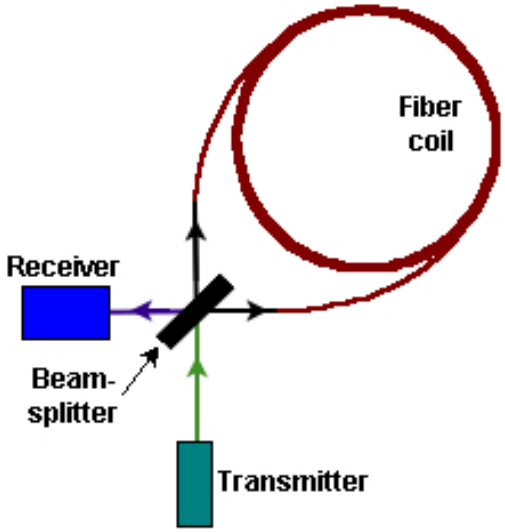

Figure 2.7: Basic schematic of a fiber optic gyroscope. A beam of light is sent in opposite directions around a fiber optic coil. An interferometer detects the phase shift between the two beams induced by the angular velocity of the coil.

the submillimeter detectors. Integrated velocity data from gyroscopes are therefore used to interpolate between absolute positions.

2.8.1 Gyroscopes

Six gyroscopes are mounted in a box on the inner frame of the gondola. For the purpose of redundancy, two gyroscopes are mounted on each of the three orthogonal axes of inner frame motion - elevation, roll, and yaw. BLASTPol used KVH DSP-3000 digital fiber-optic gyroscopes (FOGs). A FOG measures angular velocity about its sensitive axis by acting as a Sagnac interferometer (see Figure 2.7). The Sagnac effect describes the phase shift induced between light waves traveling in opposite directions around a rotating coil. The phase difference, $\Delta\Phi$, detected by interferometry is related to the angular rate of rotation, Ω, of the coil by

$$\Delta\Phi = \frac{4\pi LR}{\lambda c}\Omega, \tag{2.1}$$

where L is the length of the optical path, R is the radius of the coil, λ is the wavelength of the light and c is the speed of light.

The DSP-3000 uses a digital signal processor and an open loop optical circuit of polarization-maintaining fiber. The gyros have an angle random walk of $0.0667^\circ/\sqrt{\text{hr}}$ (or $4.0''/\sqrt{\text{s}}$), which comes from noise in the rate signal. This is independent of other sources of error such as scale

factor error of 1500ppm (=0.15%), and a bias error of $\pm 20''/s$. The angle random walk noise is proportional to the square root of the integration time and therefore limits the length of time one can rely solely on information from the gyros in the absence of another reference. This determines the frequency at which solutions from the star camera need to be available.

The bias of the gyroscopes is sensitive to changes in temperature, and therefore the interior temperature of the gyro box needs to be regulated. The heaters are five 100Ω resistors connected in parallel, and the temperature of the box is measured using a thermally sensitive resistor encased in a thermally conductive epoxy. Each gyro is connected to a mini-relay that allows it to be individually switched on or off. The four signals from each gyro - clock, data, strobe and built-in-test - are amplified through Schmidtt-trigger logic ICs and sent to the on-board computer. The gyro box receives 24V power, which is converted by a DC-DC to 5V for the gyros.

Magnetic shielding

Fiber optic gyroscopes are affected by the Earth's magnetic field due to the Faraday effect. This occurs because the magnetic field causes a rotation in the plane of polarization of the light in the optical fiber, and this results in a phase shift which affects the signal. In order to protect the gyros from this effect, they are wrapped in Metglas Alloy 2714A, a cobalt-based magnetic shielding foil that is only 0.0006 in. ($\sim 15\mu$m) thick. The material has a magnetic permeability, μ, of 1,000,000, where μ=B/H is the ratio of the flux density B to the magnetic field strength H and describes the ability of the material to be magnetized. The effectiveness of the shielding was tested and shown to reduce the effect of a magnetic field on the gyros by a factor of > 10 (Figure 2.8).

Orthogonalization

Once the gyroscopes are mounted inside the box, an orthogonalization process is done. This determines the orientation of the gyros with respect to three mutually orthogonal axes, and this information is then incorporated into the flight code. This process was done by placing the gyro box on a rotary table and turning it around three orthogonal axes of rotation. Let the three unit vectors in this orthogonal reference frame be $\hat{x}$, $\hat{y}$ and $\hat{z}$, and the direction vectors of each gyro be $\hat{1}$, $\hat{2}$, ..., $\hat{6}$. If the box is rotated at angular velocity $\omega = (\omega_x, \omega_y, \omega_z)$ and the angular velocity measured by each gyro is $[g_1, g_2, ..., g_6]$, then

$$\omega_x\hat{x} + \omega_y\hat{y} + \omega_z\hat{z} = g_1\hat{1} + g_2\hat{2} + g_3\hat{3} + g_4\hat{4} + g_5\hat{5} + g_6\hat{6}. \quad (2.2)$$

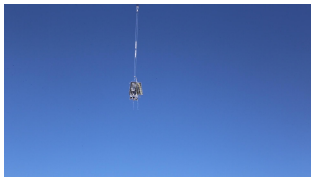

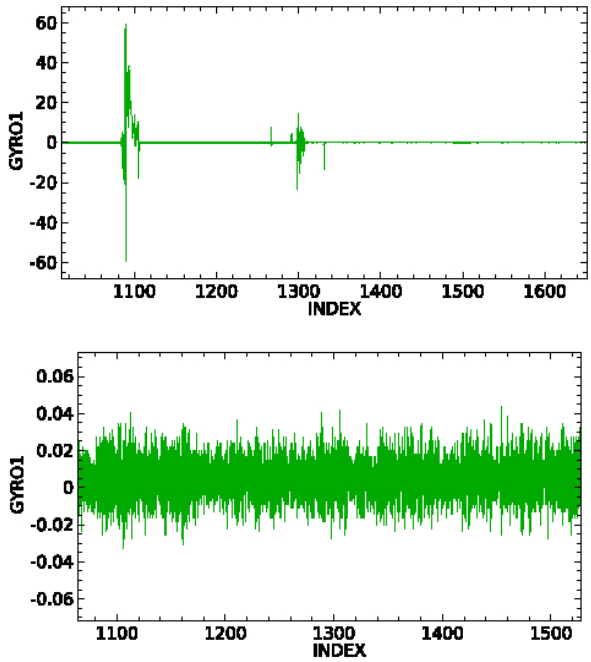

Figure 2.8: Response of gyro to a magnetic field. Top: Angular velocity (in $^\circ/s$) measured by a gyro without shielding while exposed briefly to a magnetic field. A signal up to $60^\circ/s$ is seen at INDEX=1100, Bottom: Angular rate measured with shielding while exposed to the same magnetic field. No response is seen above $0.04^\circ/s$.

So if the box is rotated in the $x-$axis, then $\omega_y = \omega_z = 0$ and the gyro signals are

$$g_i = \omega_x \hat{x} \cdot \hat{i}, \text{ for } i = 1...6. \tag{2.3}$$

If the ratio of one gyro signal versus another is measured, for example g_2 vs. g_1, then

$$\frac{g_2}{g_1} = \frac{\hat{x} \cdot \hat{2}}{\hat{x} \cdot \hat{1}}. \tag{2.4}$$

Therefore by plotting the the gyro signals against each other and performing a linear least-squares fit, and repeating this for the other 2 axes, the orientation of each gyro with respect to each other is obtained.

In order to test three orthogonal axes, three sides of the box were designed to be at exactly 90° to each other, and used to place the box on the rotary stage at three orthogonal orientations. The signals of the gyros mounted orthogonal to the axis of rotation are then plotted against the gyro mounted on the axis of rotation, and a linear least-squares fit is done to determine the orientation of each gyro. When the process was done a second time, the orientations of the gyros were mostly found to be consistent to $< 10''$. After a 90° scan, a $10''$ offset in the position of the gyro would result in an error of $15.7''$, which is much less than the required pointing

accuracy. The non-orthogonality of the box sides was measured to 0.001in accuracy using a dial indicator, and the results incorporated into the rotation matrix of the gyroscopes.

The flight digital signal processor code reads in the six gyroscope signals, and applies a 4-stage boxcar digital filter before writing the field to disk. For in-flight pointing, single-stage filtered signals from the 6 gyros are combined, using their known rotations relative to each other, to produce 3 gyro signals, one for each axis of rotation.

Every time the gyrobox is mounted on the gondola, its orientation in the reference frame of the gondola must be determined. This is done by scanning the gyrobox around 3 orthogonal axes of rotation while on the gondola. First the inner frame is pointed at an elevation of zero, and the gondola is scanned in azimuth 20 times at $2^\circ/s$. These scans are then repeated with the inner frame at an elevation of 90°. Finally, the inner frame is scanned up and down in elevation 20 times at $2^\circ/s$.

During the BLAST 2006 flight, a cosmic ray induced a single bit flip in the ACS gyroscope digital filter, causing a large jump in the measured velocity from the gyroscope. The resulting effect on the in-flight pointing solution caused the gondola to spin out of control. In order to prevent this from happening, in-flight software was written to ignore such non-physical jumps in the gyro readings, as well as to detect an unchanging gyro signal indicating a failed gyro. Failed gyros are then power cycled while their signals are masked, preventing them from contributing to the in-flight pointing solution. Redundant gyros on each axis also reduce the probability of the in-flight pointing solution being affected by gyro failure.

2.8.2 Star Cameras

BLASTPol's main pointing sensors are a pair of star cameras mounted above the inner frame. These star cameras were built for BLAST and used in its 2005 and 2006 science flights. They are described in detail in Rex et al. (2006) and Pascale (2007). For BLASTPol, the star cameras are the best pointing sensor to achieve the necessary pointing accuracy in-flight. They are also the only sensor that can provide a post-flight pointing solution with the required accuracy to make maps that oversample the angular resolution of the telescope.

Computer and software

Each star camera uses its own single board computer to run software separate from `mcp` running on the flight computers. The software that runs on the ISC and OSC is called `netisc`, which was written for previous BLAST flights, and used mostly unaltered for the BLASTPol flights. Each star camera also has a hard drive which is used to store images and `netisc` log files.

`netisc` uses a blob-finding algorithm to detect bright sources in the image. The algorithm first removes a gradient in the image and boxcar filters the image. It then calculates the mean and standard deviation of the image. It identifies sources that are above a signal to noise threshold (which is usually set to ~ 5 but can be changed by command in-flight), and interpolates the flux across several pixels of the "blob", in order to determine its centre.

The star camera software also includes a program that displays video images of the star camera field. The images show boxes around sources identified as stars, labeled in order of flux magnitude. The display also includes information about the image (frame number, whether it was saved, mean pixel value), about the camera (camera name, focus position, exposure time) and about the gondola position and time (local sideral time (LST), latitude and longitude, azimuth and elevation). If a pointing solution was found for the image, the magnitudes of the identified stars are included above the boxes around the stars, and an X appears next to the RA and Dec displayed. Images are periodically saved to disk, at a commandable rate.

Each star camera computer's video output is connected to a video converter, which converts from VGA to NTSC. The signal is sent to a CSBF video transmitter mounted below the gondola floor. The NTSC signal from each camera is transmitted over twisted pair wire to a relay box that sits on the gondola floor. The relay box is used to switch between camera video output signals, since only one signal can be sent to the transmitter. The transmitter then broadcasts the signal to the ground station at LDB when the payload is still in line of sight. For the 2012 flight, the output from each star camera's video converter first passed through a balun before being sent to the relay box, in order to reduce the undesirable effects (image distortion) caused by the unbalanced signal in the twisted wires.

Every time an image is taken, the star camera computer sends a packet to the flight computers over ethernet. It contains the x,y positions of the 15 brightest blobs, their fluxes, the mean value and standard deviation of the image map, the focus position of the lens, the CCD temperature, the time the image was taken, and the duration of the exposure. A thread in the master control program interprets the packet and writes the data to disk. The flight computer also sends electrical trigger pulses to the star cameras, which cause exposures to be taken at the appropriate times. There are two commandable exposure times, fast and slow (usually around 80 and 180 milliseconds, respectively). In the middle of a scan, the fast exposure time is used, and at turnarounds the slow exposure time is used.

`netisc` uses the Pyramid algorithm (Mortari et al. 2001) to identify the stars observed by the star camera blob-finder. It is a 'lost-in-space' algorithm, meaning no a priori knowledge of the pointing is available, so a search is conducted over the entire sky. A faster algorithm was developed for the BLASTPol star cameras, which uses the approximate pointing provided by coarse sensors, and only searches a reduced part of the sky around that position. This was

needed because BLASTPol's in-flight pointing requires the star cameras to identify stars every ~ 2 seconds, and a lost-in-space search would be too time consuming.

The way the Pyramid algorithm works is by matching a four star polygon structure (the "pyramid") with stars in a catalogue. The search is done using an already built database of angular distances between stars, which it compares to the angular distances between the identified blobs in the image. The algorithm starts by identifying all triangles in the image, and matching them to a list of triangles based on the catalogue. It then adds a fourth star and tries to find a matching pyramid using all of the triangles it found. If it finds a unique solution, it then identifies the other stars in the image and outputs their positions and magnitudes. Spurious stars caused by spikes in the CCD are rejected, and the celestial coordinates (RA and Dec) of the centre of the image are calculated using a least-squares fit of the celestial angular distances from the image centre to each star and the distances in the frame in tangent plane coordinates.

Pre-flight Tests

Before flight a list of bad pixels is compiled, and this list is used by the in-flight star camera software in order to exclude these pixels during the blob-finding process. The method used to find these pixels did not appear to find any "dead pixels", which always read zero on all exposures, or "stuck pixels", which always read the maximum value on all exposures. The bad pixel list therefore only includes "hot pixels", which are pixels whose brightness increases unusually rapidly with exposure time. The brightness is due to the pixel having higher than normal rates of charge leakage, so over time the brightness increases.

In order to find these pixels, several dark frames were taken at a range of exposure times: 100ms, 200ms, 300ms, 500ms, 700ms, 1s, 2s, 3s, 5s, 7s, and 10s. An average image is then created by summing the dark frames, weighted by their exposure times. Pixels were then identified whose value was 7σ over the mean of the average image.

In McMurdo, an alignment of the ISC and OSC relative to the submillimeter beam is also done. The purpose is to determine the x and y (horizontal and vertical) offset of each star camera relative to the telescope, which is represented by the center pixel of the 350μm array. This offset is used by the in-flight pointing solution to rotate from the star camera pointing solution reference frame to the telescope reference frame. While the beam mapping of the detector arrays is being performed, a foam board is marked with dots at a known separation, and placed in a known position relative to the submillimeter source being mapped. The star camera takes an image containing both the submillimeter source and the foam board. Then the offset can be calculated given the x, y mounting distance between the camera and the detectors, the distance between the telescope and the source, the distance between the source

Figure 2.9: Setup for aligning the star cameras with the submillimeter beam. A piece of foam with two "+" markings is held a known distance above the submillimeter source being used for detector beam mapping.

and the dots, and the pixel scale of the CCD. This alignment is also repeated once the telescope reaches float using a bright submillimeter source as a pointing calibrator.

2.8.3 Roll star camera

In both the 2010 and 2012 BLASTPol flights, a star camera being tested for SPIDER (see Chapter 3) was attached to the outer frame of the gondola (Figure 2.10). This camera pointed in a direction roughly orthogonal to the boresight star cameras, so its pointing solution could constrain the roll of the telescope around the boresight. It took a single 300ms exposure at each scan turnaround, when the gondola was stationary in azimuth. Its video output was also sent to the video transmitter relay box. It was designed to test all of the functions of the SPIDER boresight camera, and is described in more detail in Section 3.2.4. The roll star camera flown on BLASTPol went on to fly on SPIDER as its boresight star camera.

2.8.4 Coarse Pointing Sensors

BLASTPol uses a combination of coarse sensors to determine its attitude in flight. The purpose of the coarse sensors is to orient the payload roughly in position, as their accuracy is too coarse to point the telescope to the desired $30''$. The coarse attitude is provided as a starting point around which the star cameras can search a catalogue to identify stars in their images.

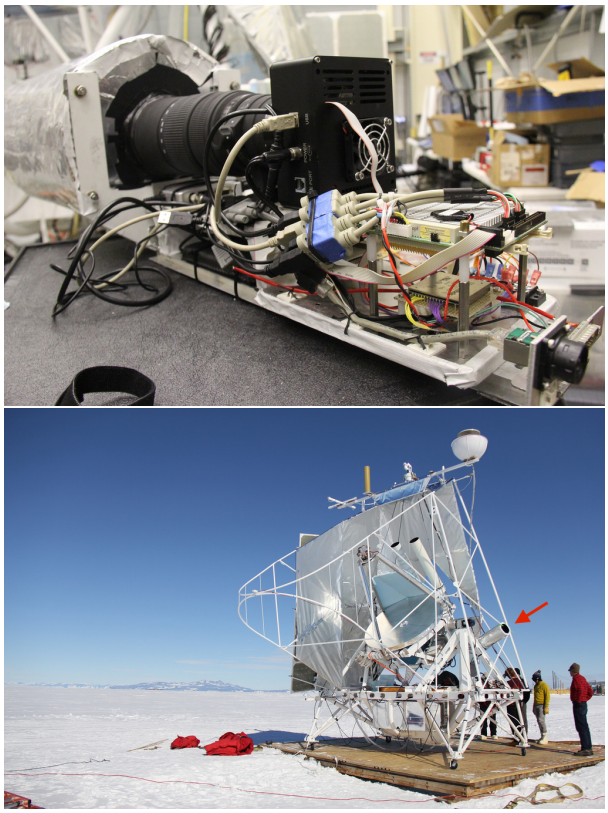

Figure 2.10: Roll star camera. Top: Camera in the lab before being covered in a foam box. Bottom: Camera (indicated by arrow) mounted to the gondola.

GPS

For BLASTPol's 2010 flight, a new (never flown on a balloon experiment) differential GPS unit was used - the Septentrio PolaRx2e@. The GPS provides time, position (latitude, longitude, altitude), speed (horizontal and vertical), and attitude (heading, pitch, roll). The receiver uses an array of three antennas in order to determine the attitude of the payload. A main antenna is used to calculate the location of the receiver, and two auxiliary antennas are used to calculate attitude, by using carrier-phase measurements along the antenna baselines. The main antenna is a Septentrio PolaNt* dual frequency GPS antenna, operating at 1575 ± 10MHz and 1227 ± 10MHz. The auxiliary antennas are PolaNt*_SF single frequency antennas, operating at 1575 ± 2MHz.

The accuracy of the reported attitude can be increased by properly choosing the placement of the GPS antennas. First, the distance between the antennas should be as large as possible, given the size of the payload. Second, the angle between the *main-auxiliary1* and *main-auxiliary2* baselines should be as close to 90° as possible. This is to maximize the projection of the *main-auxiliary2* baseline onto the plane perpendicular to the *main-auxiliary1* baseline, which maximizes the accuracy of the roll determination. Third, the antennas should be placed as far away as possible from any reflective surfaces on the payload (e.g. large CSBF antennas, or mylarized sunshields). This is to reduce the effect of multipath error, which is the error introduced by the signal from a GPS satellite being reflected off some surface before reaching the antenna. The attitude reported by the GPS is the attitude of the antenna array, not the attitude of the payload. Therefore any offsets between the array axes and the payload axes must be entered into the flight code in order to convert GPS attitude into telescope attitude. If the positions of the antennas are known with high accuracy, they can be provided to the receiver, but they are not required to determine attitude.

The receiver communicates with the flight computer over RS-232 serial in a proprietary Septentrio Binary Format (SBF) that is parsed by `mcp`. The settings on the receiver are configured to enable SBF output, set the output packets (listed below), the output rate (10Hz) and the serial settings. The receiver's primary communication port (COM1) outputs the messages, while the secondary port (COM2) can be used on the ground to send commands and debug the unit. The output message blocks are

- `ReceiverTime`: This block provides the current time with a 1-second resolution. The time is provided in both GPS time and UTC time. GPS time is given as Time-Of-Week (TOW) and Week Number count (WNc). TOW is expressed in whole milliseconds from the beginning of the current GPS week. WNc is a continuous week count with Week 0 starting at 00:00:00 UTC on Sunday, January 6th, 1980. UTC time is given in year, month,

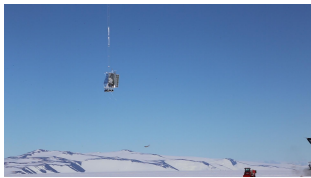

day, hour, minute and second.

- `PVTGeodetic`: This block contains position in latitude, longitude and height (ϕ, λ, h) and velocity in the North, East and "Up" direction (v_n, v_e, v_u). These are given in geodetic coordinates that make use of the WGS-84 ellipsoid.

- `AttitudeEuler`: This block contains the computed Euler angles pitch, roll and heading (ϕ, θ, ψ), calculated using the WGS-84 ellipsoid.

- `AttitudeCovEuler`: This block contains the diagonal elements of the symmetric variance-covariance matrix of the attitude angles:

$$\begin{pmatrix} \sigma_\phi^2 & \sigma_{\phi\theta} & \sigma_{\phi\psi} \\ \sigma_{\theta\phi} & \sigma_\theta^2 & \sigma_{\theta\psi} \\ \sigma_{\psi\phi} & \sigma_{\psi\theta} & \sigma_\psi^2 \end{pmatrix} \tag{2.5}$$

- `AuxAntPositions`: This block contains the antenna system information, including the relative positions and velocities of the two auxiliary antennas with respect to the main antenna, in the local East/North/Up reference frame.

These settings are saved as a configuration file in non-volatile memory, forcing the receiver to start up in that configuration after a device reboot. The GPS receiver time is also used as a Network Time Protocol (NTP) server to synchronize the flight computer clocks, through a NTP daemon running on the flight computers. The star camera computers also run their own NTP daemons, and synchronize their clocks to whichever flight computer is currently in charge.

For the 2010 BLASTPol flight, the Septentrio antennas were bolted to flat plates at the ends of aluminum tubes. Two 0.5m tubes were attached to the top back corners of the sunshields, pointing directly outward. A third 3m pole was cantilevered off the top of the sunshields, pointing forward. This produced an L-shaped antenna array, with 3m between the *main* and *auxiliary1* antenna, 3m between the *main* and *auxiliary2* antenna and 90deg between the *main-auxiliary1* and *main-auxiliary2* baselines. For this antenna spacing, the expected accuracy is 0.1°, 0.2°, and 0.2° in heading, pitch, and roll, respectively.

When the Septentrio GPS was first mounted onto the BLASTPol gondola in 2010, the antennas were not near any reflective surfaces, and the calibration and testing showed no problems. However the forward most antenna was mounted on a long boom that placed it next to the carbon fiber inner frame baffle structure (see Figure 2.2). Later in the campaign this structure was covered in aluminized Mylar, creating a highly reflective surface near the front GPS antenna. At that point the mounting of the GPS antennas could not be changed, so the multipath effect is visible in the resulting performance.

Figure 2.11: Septentrio PolaRx2e@ receiver (left) and PolaNt antenna (right).

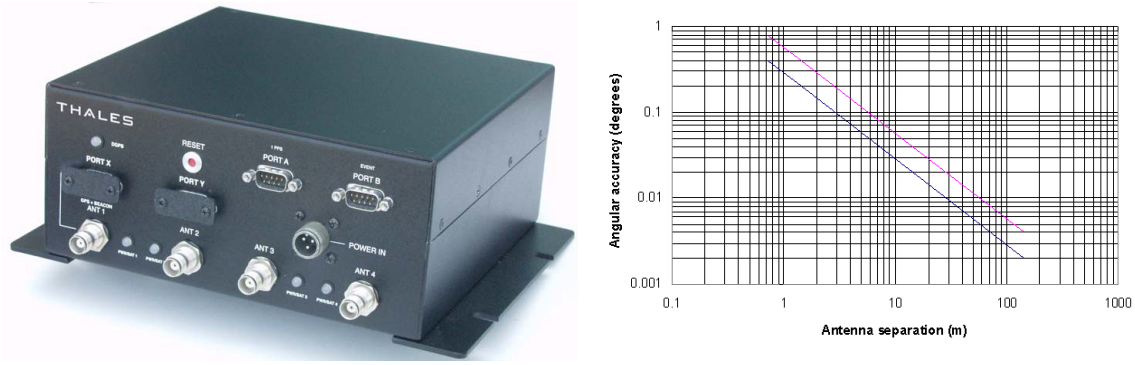

Figure 2.12: Left: Thales ADU5. Right: Accuracy as a function of antenna separation. The lower line represents heading, and the upper line is pitch and roll.

For BLASTPol's 2012 flight, the Septentrio unit was not flown. Instead, GPS information was acquired from the GPS unit flown by CSBF, which is a Thales Navigation ADU5. The ADU5 outputs position and attitude messages in National Marine Electronics Association (NMEA) format at 1Hz. The flight computer is connected over RS-232 serial to the secondary communication port on the ADU5, with the primary communication port being used by CSBF to send GPS data to the SIP. The CSBF antenna array is a 1m cross mounted on top of the sunshields, providing an accuracy of 0.3° in heading and 0.6° in pitch and roll (see Figure 2.12).

Pinhole Sun Sensors

BLASTPol uses multiple pinhole sun sensors, mounted on the outer frame of the gondola, to determine the azimuth of the outer frame based on the position of the sun. For the 2010 flight, two sun sensors were flown - one to be used during observations of Vela and one to be used during observations of Lupus. For the 2012 flight, two additional sensors were flown with a new design that will be used on the 2014 flight of SPIDER.

The sun sensors use Hamamatsu S5991-01 two-dimensional 9×9mm position-sensitive de-

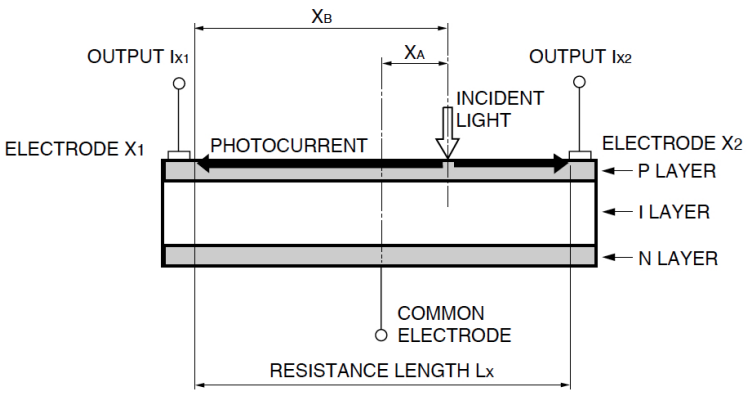

Figure 2.13: Cross section of a PSD

tectors (PSDs). A PSD consists of a uniform P-type resistive layer formed on the surface of a N-type high-resistivity silicon substrate. The P-layer acts as an active area for photoelectric conversion and a pair of output electrodes are formed on both ends of the P-layer for extracting signals. A common electrode is connected to the backside of the silicon substrate (Figure 2.13).

When a light spot hits the PSD, it generates an electric charge that is proportional to the light intensity. The electric charge travels through the resistive layer to the output electrodes, generating a photocurrent that is inversely proportional to the distance between the position of the incident light and the electrode. The PSD used is a tetra-lateral pincushion type two-dimensional PSD (Figure 2.14). In a tetra-lateral PSD, the N-type layer is processed to form a resistive layer, and two pairs of electrodes are formed on the upper surface. One pair (the X electrodes) extracts the X position signal, and the other pair (the Y electrodes) extracts the Y position signal. Because the electrodes are formed on the same surface, they interact with each other at the corners of the active area, causing the position to be distorted (compared to a duo-lateral PSD where the X and Y electrodes are on opposite surfaces of the N-type layer). However a tetra-lateral type PSD features the application of a reverse bias voltage, small dark current, and high-speed response. The pincushion variant is an improvement on the tetra-lateral type that reduces the electrode interaction at the corners, reducing the distortion in the position measurement.

The readout connection of the PSD is shown in Figure 2.15. The output is read through an operational amplifier, with output voltage, Vout, determined by the feedback resistance, Rf, and the input current, Isc. The required output voltage is between 5-10V in order to be within the dynamic range of the BLASTPol analog cards. For the two new PSS's flown in 2012, custom circuit boards were made according to the schematic shown in Figure 2.16, which included input voltage regulation and output signal amplification.

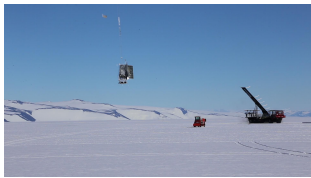

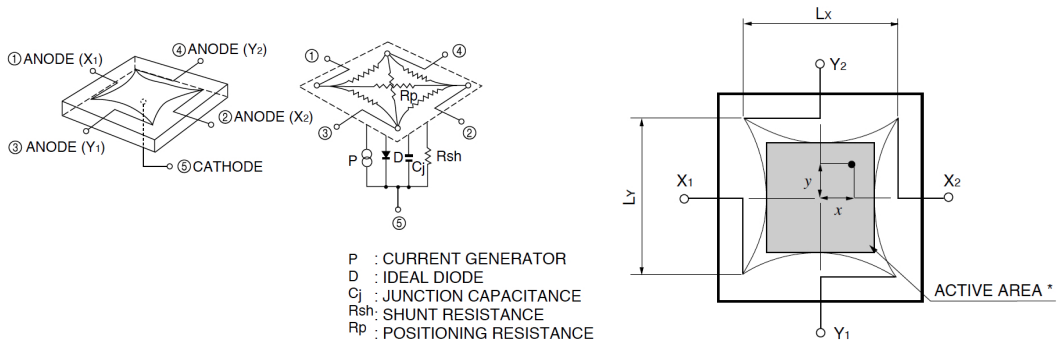

Figure 2.14: Circuit (left) and active area (right) of a tetra-lateral PSD.

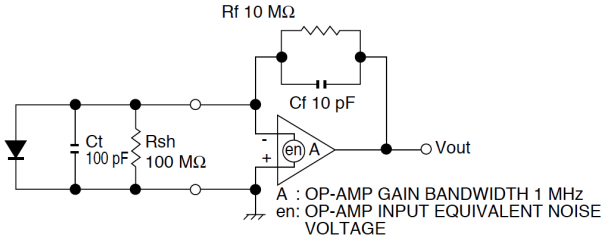

Figure 2.15: PSD connection

The measured currents $(I_1, ..., I_4)$ from the four anodes placed at the corners of the detector give the position of the light spot on the detector according to a conversion formula provided by the manufacturer:

$$\frac{(I_2 + I_3) - (I_1 + I_4)}{I_1 + I_2 + I_3 + I_4} = \frac{2x}{L}$$
$$\frac{(I_2 + I_4) - (I_1 + I_3)}{I_1 + I_2 + I_3 + I_4} = \frac{2y}{L} \tag{2.6}$$

where $L = 10$mm. Sunlight comes in through a 200μm copper precision pinhole from Edmund Optics, mounted 10mm above the centre of the active area of the PSD.

Once the position of the light spot is known, the angle between the sensor and the sun can be calculated, knowing the distance from the sensor to the pinhole (see Figure 2.17). Only detections of a light spot $\leq$4mm from the centre of the sensor are considered, since the accuracy degrades near the edge of the sensor. This gives each sun sensor a field of view of $2 \times \arctan \frac{4mm}{10mm} \simeq 40°$. Each sensor is held in an aluminum case, with the inside painted black (Figure 2.18). The cases are mounted on a boom to avoid areas shaded by the gondola (Figure 2.19). The outer frame is preferred over the top of the sunshields as a mounting point because it is more rigid and less susceptible to warping, which makes the orientation of the sensors more fixed over the course of the flight. The sensors are mounted at different angles in azimuth in order to cover the range over which BLASTPol scans in azimuth relative to the

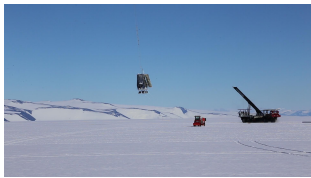

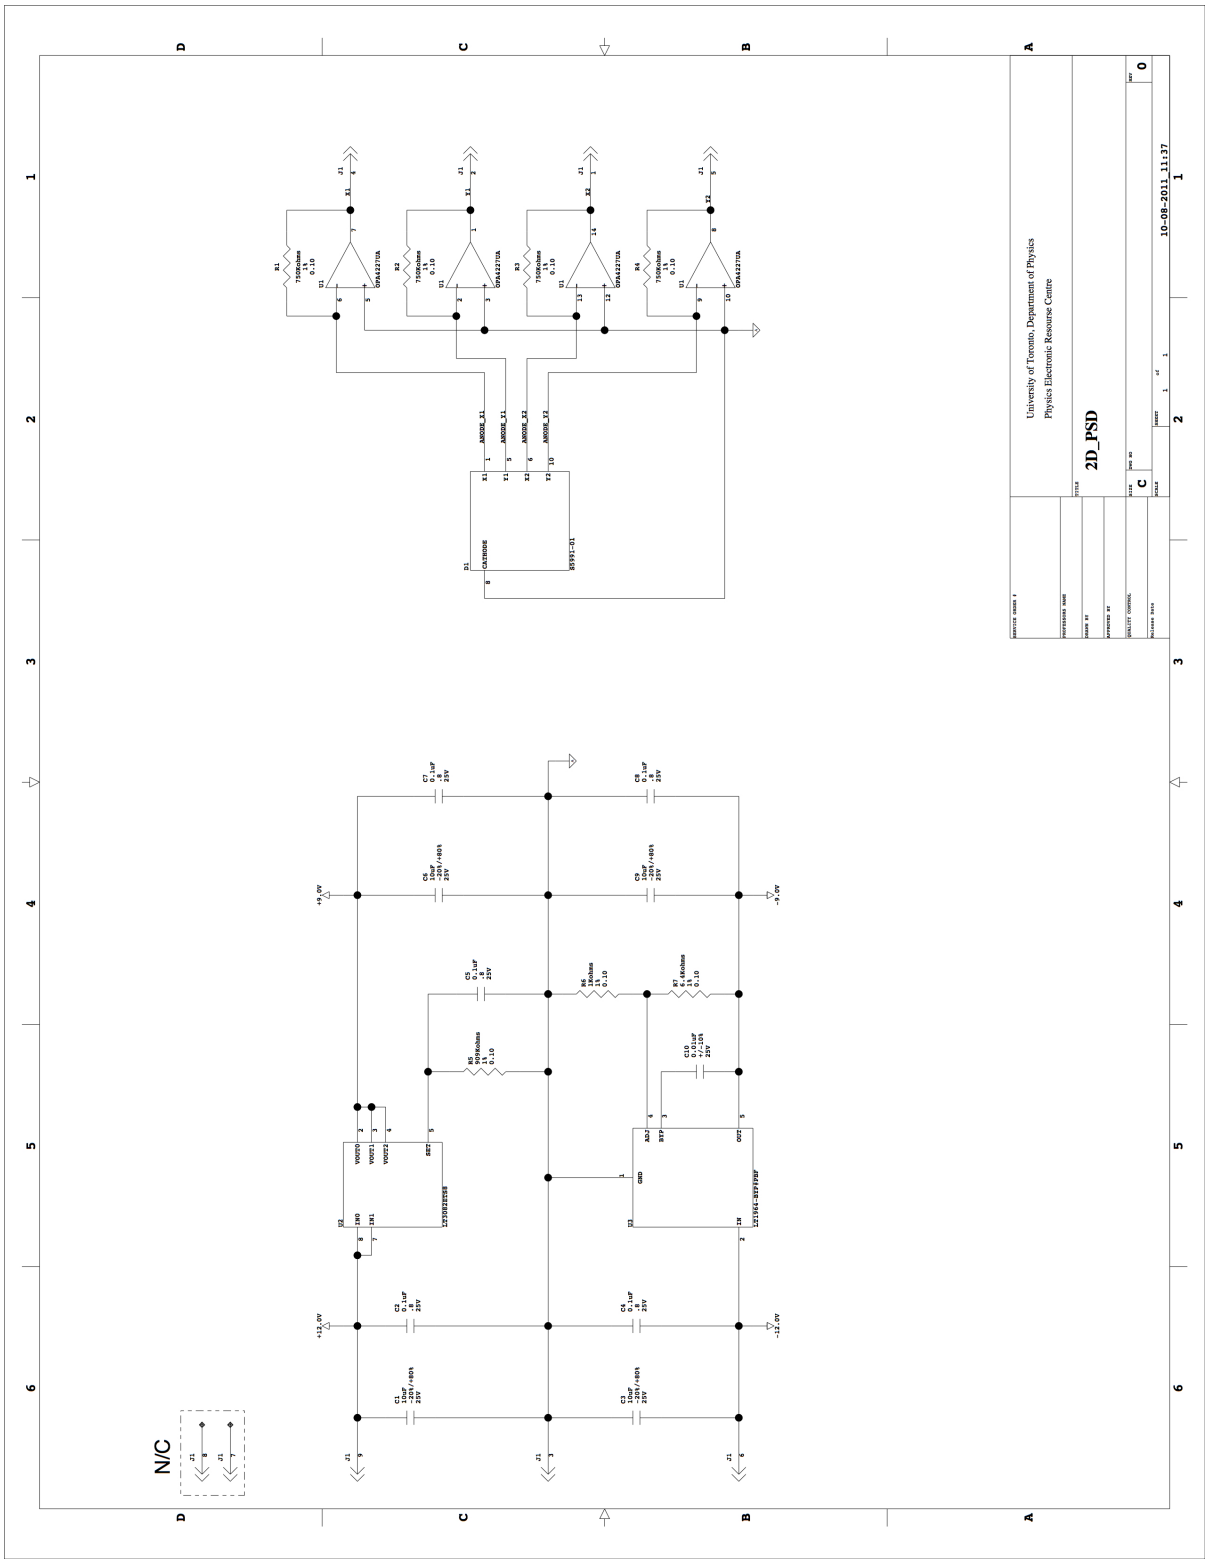

Figure 2.16: Readout schematic.

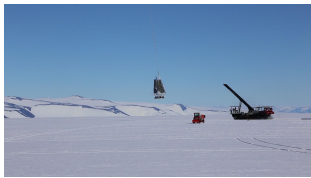

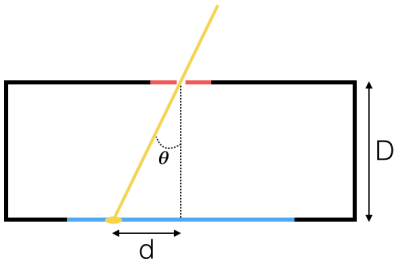

Figure 2.17: Light passing through the pinhole (red) and hitting the PSD (blue). The angle, θ, between the sensor and the sun can be calculated from the distance, d, between the light spot and the centre of the PSD, and the distance, D, between the sensor and the pinhole.

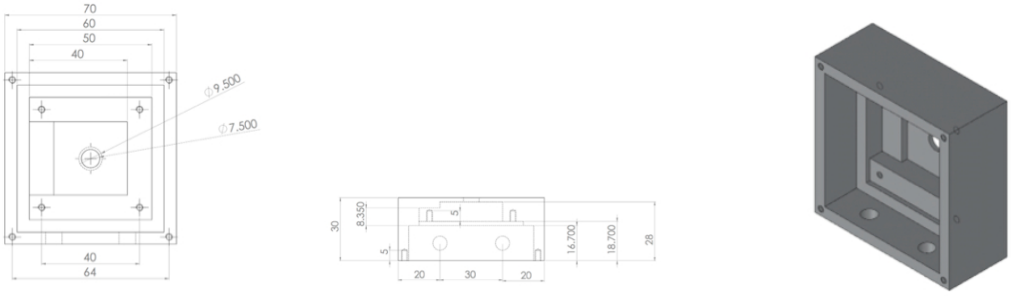

Figure 2.18: PSS case.

sun, and such that they overlap each other's field of view by $5°$. Table 2.1 lists the intended mounting angles of the sun sensors for both BLASTPol flights. The angles given are in azimuth relative to the telescope boresight. The elevation of the sun varies from $15 - 35°$ in Antarctica in December, so each case is mounted at $25°$ elevation.

The in-flight pointing solution calculates an azimuth for any sensor whose total current exceeds a certain threshold. It then combines the azimuths from all sensors, weighted by their signal-to-noise. Various parameters that go into the calculation of sun sensor azimuth can be adjusted during pre-flight calibrations, or during the flight. Inaccuracies in the construction of the sun sensor mount introduce deviations from the nominal 10mm distance between sensor and pinhole, as well as the azimuthal mounting angles listed in Table 2.1. These numbers are adjusted during pre-flight calibrations as described in Section 2.8.4. The minimum current threshold can also be adjusted. Parameter adjustments can also be commanded in-flight, using information gained during ascent of the balloon when the gondola spins in azimuth, or at float during azimuthal slews of the telescope.

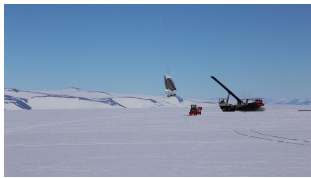

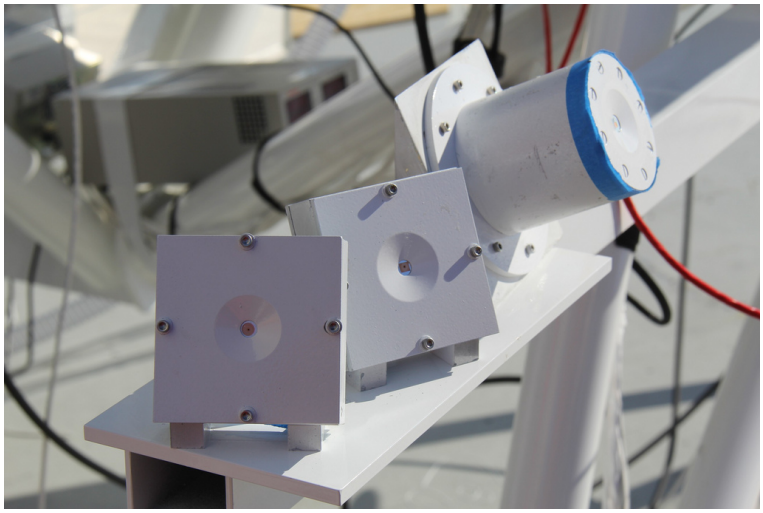

Figure 2.19: PSS boom.

2010	Angle [°]
PSS1 (Lupus)	60
PSS2 (Axehead)	150
2012	
PSS1	50
PSS2	85
PSS3	120
PSS4	155

Table 2.1: PSS mounting angles.

Magnetometer

A magnetometer is mounted at the most forward point of the sunshield (see Figure 2.2). It provides coarse azimuth by detecting changes in the direction of the Earth's magnetic field. The direction of the magnetic field lines at any given location is given by the National Oceanic and Atmospheric Administration (NOAA) World Magnetic Model. The model only accounts for the portion of the magnetic field generated in the Earth's fluid outer core. It does not include effects generated in the crust and upper mantle, or in the magnetosphere and ionosphere. The model is only updated every 5 years, and after 5 years the accuracy of the model is $\sim 10\%$. The model used for the BLASTPol 2010 and 2012 flights was published in December 2009, and expired in December 2014. Near the south pole, the magnetic field lines are highly inclined, so detecting changes in the component parallel to the horizon is less effective.

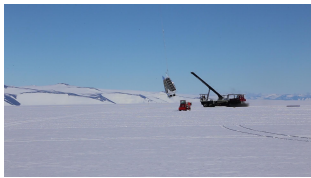

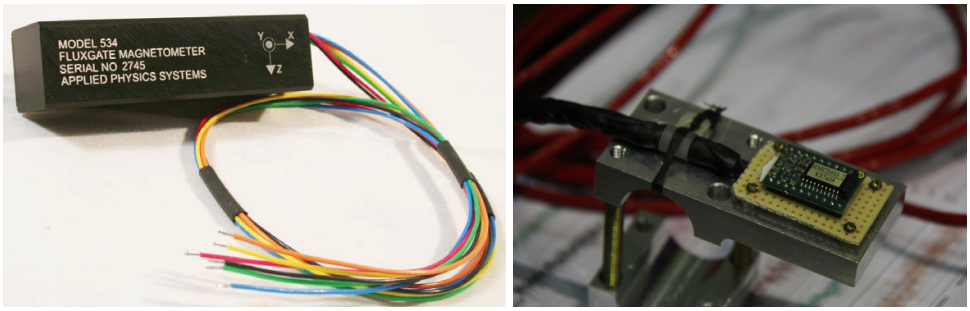

Figure 2.20: Left: 2010 magnetometer. Right: 2012 magnetometer.

Both BLASTPol flights used 3-axis analog magnetometers, which provide three output voltages proportional to the magnetic field in three orthogonal directions. They are shown in Figure 2.20. The magnetometer flown in 2010 was an Applied Physics Model 534 3-axis fluxgate analog magnetometer. A fluxgate magnetometer consists of a small, ferromagnetic core wrapped by two coils of wire. An alternating electrical current (AC) passed through one coil causes an alternating magnetic field that induces an AC in the second coil. The AC in the second coil is measured and changes due to external magnetic field changes can be detected.

The BLASTPol 2012 flight used a HMC2003-T Honeywell magnetometer. It has three magneto-resistive sensors sensitive to magnetic fields along three perpendicular axes. The magneto-resistive sensors use a thin film of Permalloy (nickel-iron) deposited on silicon as a resistive strip. The Permalloy is magnetized in a particular direction, and its resistance to current depends on the angle between the direction of magnetization and the direction of current flow. In the presence of the Earth's magnetic field, the magnetization direction is deflected, causing a change in resistance. This produces a corresponding change in voltage across the output of a wheatstone bridge. The outputs are voltages proportional to magnetic field strength in the three axes.

The magnetometer must be mounted using only non-magnetive materials (e.g. nylon or copper fasteners). For both BLASTPol flights, the magnetometer was mounted at the forward-most point of the gondola, which is a lower part of the sunshield that extends upwards from the bottom of the outer frame and shields the telescope from ground reflections. For the 2012 flight, the sensor was mounted onto a circuit board and bolted to an aluminum plate, all of which was wrapped in blue foam and shielded in aluminized mylar.

Clinometer

BLASTPol uses two biaxial clinometers from Applied Geomechanics. Both are variants of the 904-T "Clinometer Pak". One clinometer is mounted to the bottom of the gondola, to measure pitch and roll of the outer frame. It is the Model 904-TH high-gain version, with a $\pm 10^{\circ}$ span.

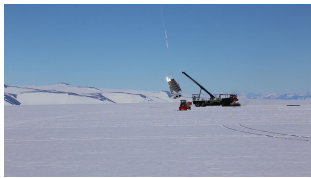

Figure 2.21: Clinometer.

The second clinometer is mounted to the inner frame as an inner frame elevation sensor. It is the Model 904-TS standard version, and has a $\pm 25^\circ$ span. The clinometer uses a liquid-filled electrolytic transducer as the sensing element. The sensor cannot detect pendulations of the gondola, only constant tilt about two orthogonal tilt angles. Its response to tilt also has a temperature dependency. Overall it provides an accuracy of $\sim 0.1^\circ$.

Elevation Encoder

Elevation of the inner frame relative to the outer frame is provided by absolute encoders. BLASTPol uses a Kollmorgen C053A as an elevation motor, which has a built-in encoder with sub-arcsecond accuracy, although it is read out with $20''$ resolution. Because the encoders measure the elevation angle of the inner frame relative to the outer frame, they do not account for pendulations in the outer frame. Therefore they do not measure the absolute elevation of the telescope boresight, and this limits their accuracy as elevation sensors to $\sim 0.1^\circ$.

2.8.5 Pre-flight Pointing Sensor Calibration

Once they are electrically and mechanically integrated with the rest of the telescope, the coarse sensors must be calibrated. This involves measuring their mounting angles relative to the telescope boresight, as well as any gain factors or offsets in their attitude readings. These calibrations must be done outside in order for the sun sensors and GPS to work, and they must be done off the launch vehicle in order to reduce metallic interference with the magnetometer and multipath interference with the GPS. Preliminary calibrations are done during the pre-flight integration at CSBF in Palestine, Texas in order to verify the correct order-of-magnitude and sign of the attitude being reported by each sensor. Final sensor azimuthal alignment numbers and measurement gains are calculated at the Long Duration Balloon site in McMurdo by taking the telescope to a wooden platform (the "dance floor") away from the payload buildings (Figure 2.22). The platform needs to be constructed as level as possible in order to not intro-

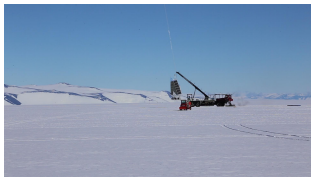

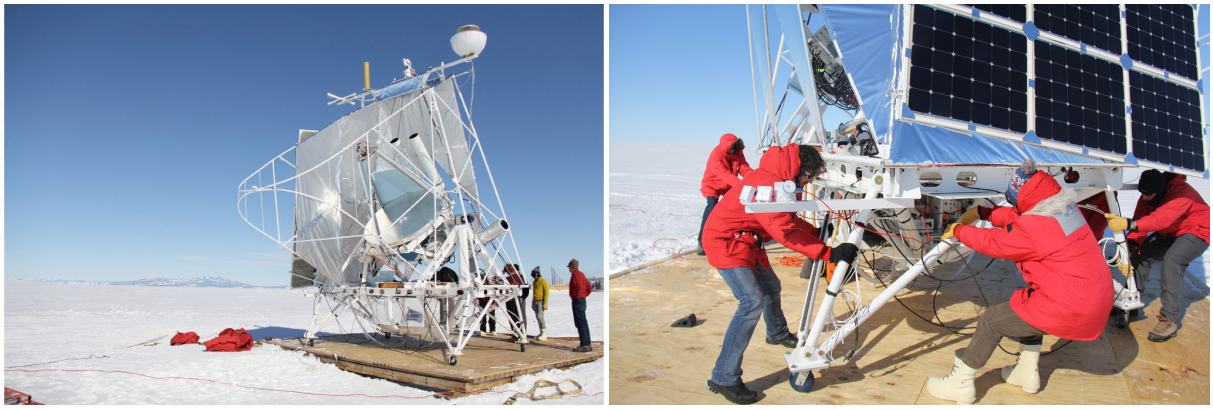

Figure 2.22: BLASTPol on the dance floor (left). BLASTPol being rotated by hand in order to calibrate pointing sensors (right).

duce pitch and roll errors into the azimuthal calibration of the sensors. An absolute azimuth reference is provided by imaging the bright star Canopus with the star cameras. Sensor gains and offsets are then calibrated by rotating the gondola 360° by hand while it is on the dance floor. The pinhole sun sensors have an additional calibration factor which is the distance, d, between the pinhole and the sensor. While it is designed to be 10mm, it is expected that the actual value will differ from that, so the value used in the flight code is adjusted based on the readings on the dance floor. By using the azimuth reading of the GPS as a reference, the parameters are tuned to calibrate the gain and offset of the sun sensor measurement of azimuth rotation.

2.8.6 In-flight pointing solution

A separate pointing solution is calculated for each sensor. It is a weighted sum of the sensor's old pointing solution, p_{old}, and the new reading from the sensor, x:

$$p = \frac{w_1 p_{old} + w_2 x}{w_1 + w_2} + g + t \tag{2.7}$$

where $w_1 = \sigma_{sys}^2$ is the systematic variance of the sensor and $w_2 = w_{\rm samp}$ is the sample weight given to each new reading of the sensor. These two weights are defined according to the accuracy of the sensor. g is the change in angle given by integrating the gyros over the time since the sensor's last reading, and t is a trim angle that can be added by command to adjust the sensor's solution. The variance in the pointing solution is then given by

$$\sigma^2 = w_1 + \frac{1}{w_1 + w_2}. \tag{2.8}$$

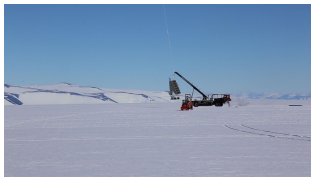

The separate solutions for each sensor, p_i, then get combined into a pointing solution, P:

$$P = \frac{\sum(\sigma_i^2 + w_{1,i})p_i}{\sum(\sigma_i^2 + w_{1,i})}, \tag{2.9}$$

A separate P_{AZ} and P_{EL} are calculated by combining either azimuth or elevation sensor readings. A sensor can also be vetoed by command during the flight if it appears to be malfunctioning. All un-vetoed sensors are used to calculate an estimate of the offset in the gyroscopes, which slowly varies over time.

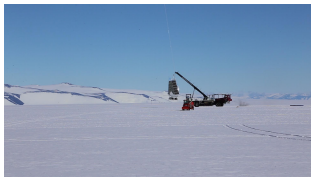

Chapter 3

SPIDER

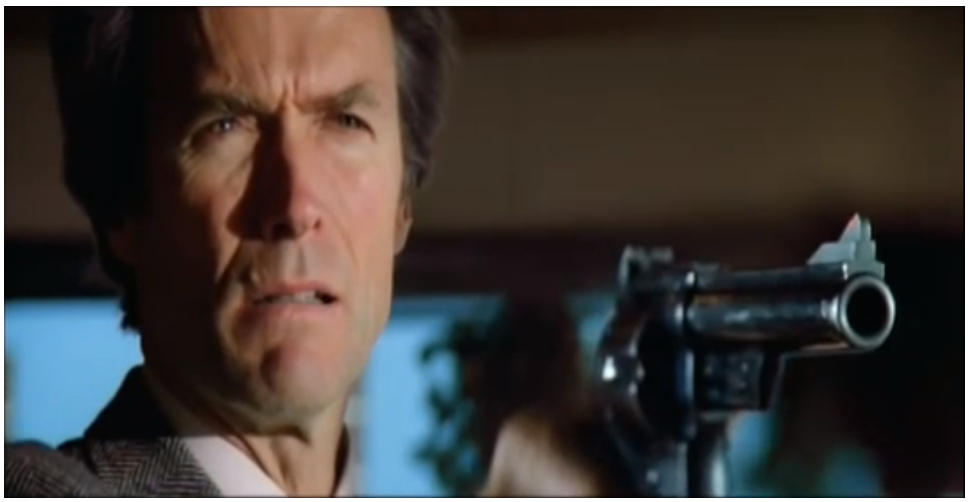

Figure 3.1: Go ahead. Make my day.

SPIDER (Filippini et al. 2010) is a balloon-borne experiment designed to measure the polarization of the Cosmic Microwave Background (CMB) in order to look for the signature of gravitational waves from inflation. The theory of inflation proposes that there was a brief period after the Big Bang when the universe expanded by a factor of $> 10^{25}$ in $\sim 10^{-33}$ seconds. The existence of an inflationary period would explain many of the observed properties of the universe, such as its spatial flatness and large scale isotropic and homogeneous nature. It would also provide a source for the large scale structure in the universe, with quantum fluctuations in the "inflaton" field seeding the primordial density fluctuations that eventually grew through gravitational collapse.

While cosmological observations are consistent with the theory of inflation, there is no direct evidence that this inflationary expansion occurred. Inflation theory predicts that gravitational waves caused by tensor perturbations from inflation left a distinct divergence-free

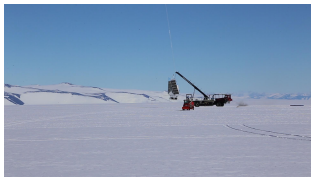

"B-mode" polarization pattern in the CMB. A curl-free "E-mode" CMB polarization results from scalar perturbations, and its power spectrum has been measured by several experiments (Bennett et al. 2013; Ade et al. 2014). Small scale B-modes have also been detected, and are the result of gravitational lensing of E-modes. A large-scale B-mode signal is considered a "smoking gun" for inflation, and the ratio, r, of tensor-to-scalar perturbations is directly related to the energy scale, E, of inflation. If inflation happened at the scale of Grand Unified Theory (GUT) physics, $E > 10^{16}$GeV, corresponding to $r > 0.01$, then its signature might be observable by modern experiments. Recently the BICEP2 experiment reported detection of a B-mode signal at a level of $r = 0.2$ (Ade et al. 2014), but a joint analysis of BICEP2 and Planck data concluded that a larger fraction of the signal was attributable to dust than previously thought. However it remains uncertain exactly how much of the signal is due to inflationary gravity waves versus polarized galactic foregrounds (BICEP2/Keck et al. 2015).

A number of ground-based, airborne, and future satellite experiments are attempting to detect the primordial B-mode signal. SPIDER is designed to take advantage of the low atmospheric background and large sky coverage available from an Antarctic LDB flight. The stratospheric platform allows observations at higher frequencies, which are necessary to constrain the polarized dust signal. The first SPIDER LDB flight took place in January 2015, and observed at 90GHz and 150GHz. The second will be in December 2017, adding detectors at 280GHz. After two flights, SPIDER will place a 3σ upper limit on r of 0.03.

3.1 SPIDER Telescope and Pointing Requirements

SPIDER consists of six telescopes, contained in a single liquid helium cryostat. Each SPIDER telescope insert is a monochromatic refracting telescope, based on the same design as the BICEP2 experiment. Each telescope illuminates a focal plane of antenna-coupled transition edge sensors (TESs). For the 2014 flight, three focal planes had detectors at 90GHz, in a 12x12 array, and three had detectors at 150GHz, in a 16x16 array. The beam full width half max at 90 GHz is 49 arcminutes, and at 150GHz it is 30 arcminutes. The field of view of each array is around 10°. SPIDER mapped roughly 10% of the sky, in a region of low galactic foregrounds, by scanning sinusoidally in azimuth with a peak velocity of 6°/s, and stepping in elevation.

In order to point the detector arrays at the desired part of the sky during the flight, the in-flight pointing solution needs to be accurate to $\sim 1^{\circ}$. Post-flight pointing reconstruction needs to be accurate to $< 10'$, based on simulations which show that at that level of pointing uncertainty the effect is negligible compared to the B-mode signal expected if $r = 0.1$ (MacTavish et al. 2008).

SPIDER's mode of operation is similar to BLASTPol, i.e. it scans an outer frame in azimuth

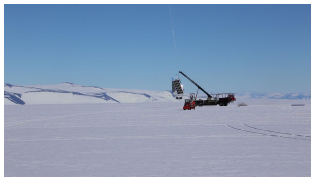

and an inner frame in elevation, in order to map a certain square patch of the sky. Therefore many of the pointing systems used on SPIDER are based on BLASTPol systems, with design modifications to account for the differences in the two experiments (see Shariff et al. (2014) for details).

3.2 Gondola

3.2.1 Structure

The SPIDER gondola is designed to support and shield the cryostat and point the telescopes at the desired observing region on the sky.

Inner frame

The inner frame of SPIDER consists of the cryostat, which is mounted to elevation drives on the gondola. The outer surface of the cryostat vacuum vessel is used as a mounting point for inner frame components, including the boresight star camera.

Outer frame

The outer frame of the gondola is made up of $\sim$ 70mm diameter carbon fiber tubes and aluminum nodes. The ends of the tube are glued to aluminum inserts using ScotchWeld 2216 epoxy adhesive. The inserts are then bolted to the aluminum nodes. The bottom plane of the outer frame holds 1inch thick aluminum/aluminum honeycomb panels, which serve as the floor of the gondola. The floor provides a mounting point for several components of the gondola control system, including the flight computers, the ACS crate, the serial hub, the rotating star cameras, and the gyrobox. Below the floor of the outer frame are the reaction wheel and SIP cage. The bottom of the SIP cage is used to mount the CSBF solar panels and ballast hopper.

In the lab, the gondola sits on a custom made aluminum cart, with pneumatic tires, allowing the entire experiment to be moved by hand. Nodes at the bottom corners of the outer frame serve as mounting points for the sunshields. The sunshields are constructed from 30mm diameter carbon fiber tubes and aluminum nodes. However the structural requirements are much less than for the gondola outer frame, since the only mass that the sunshields need to support are various antennas, the solar panel arrays, and the aluminized Mylar shielding. The entire sunshield frame can be assembled separately and mounted onto the gondola as one structure.

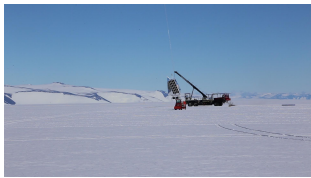

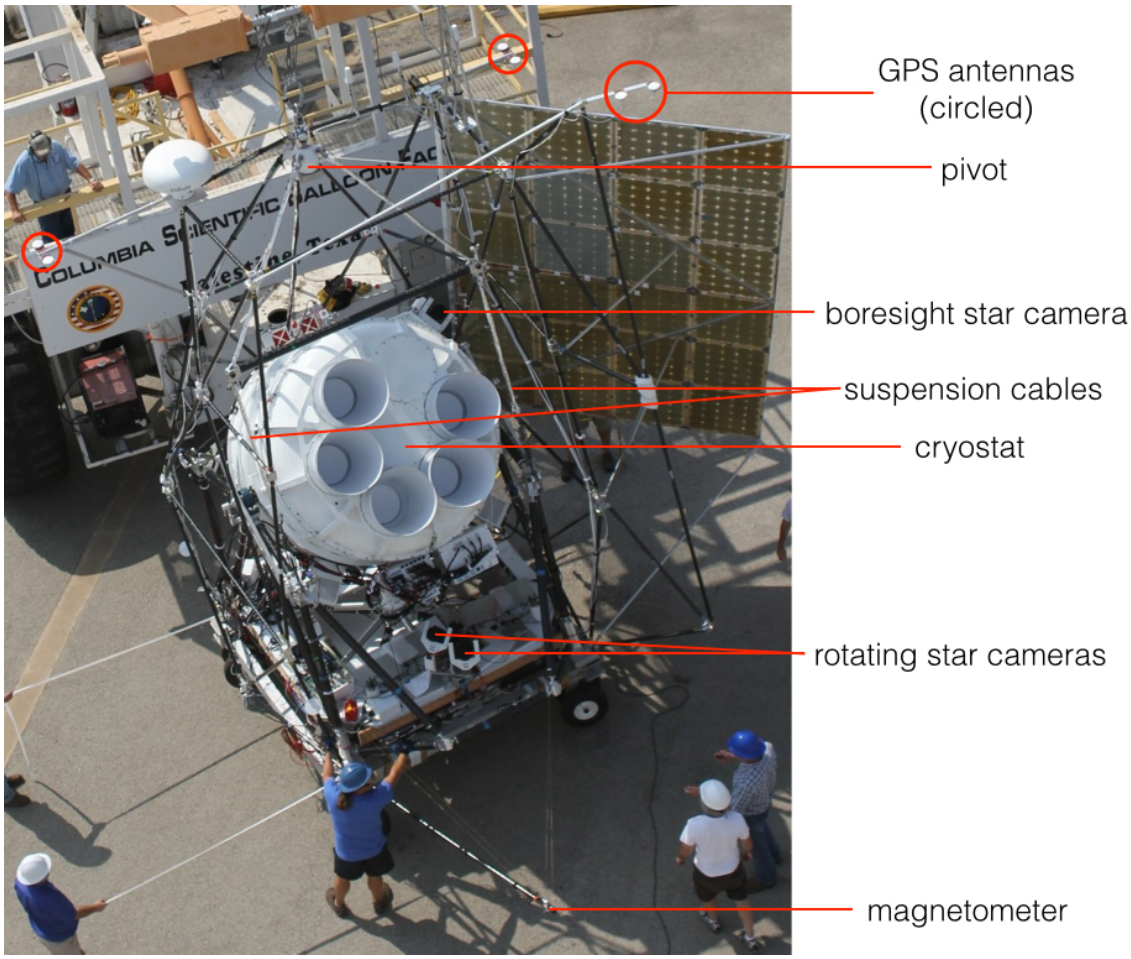

Figure 3.2: SPIDER on the launch vehicle during integration in Palestine, Texas.

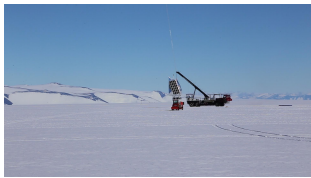

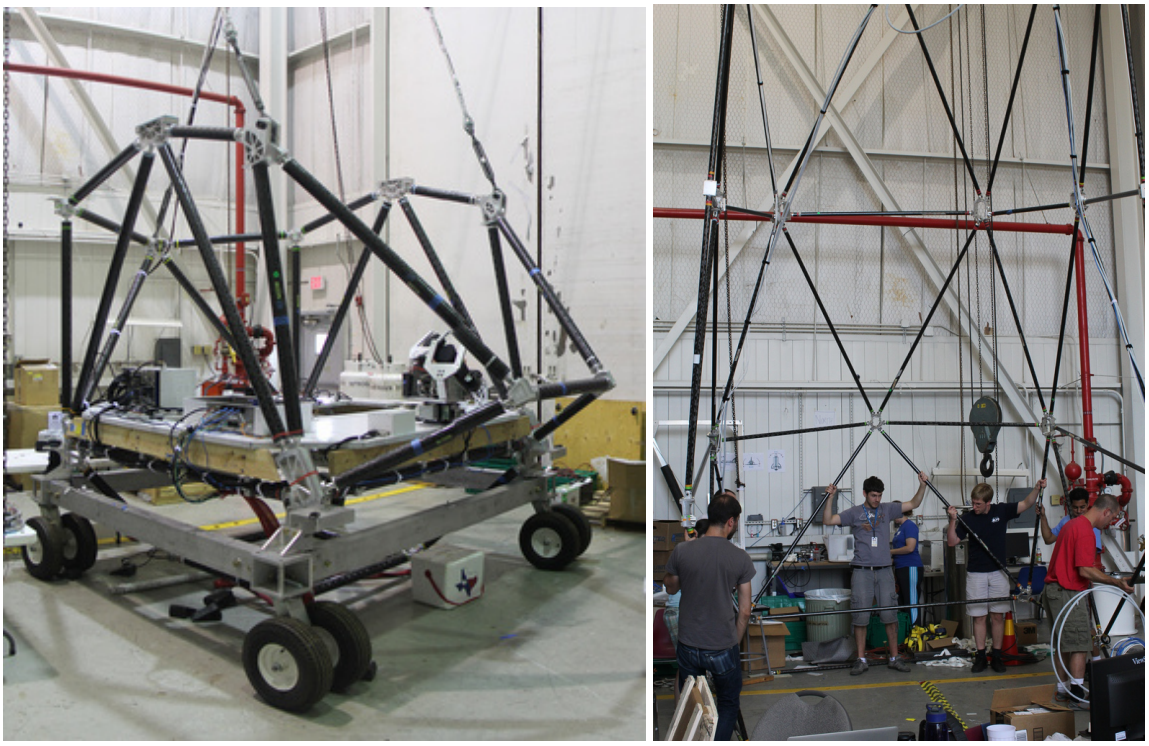

Figure 3.3: SPIDER outer frame and cart (left), and sun shield frame (right).

The gondola is suspended from two nodes in the front and one in the rear using suspension ropes made from braided Technora®, a strong and light weight fiber that maintains its strength during repeated flexing and stretching. It also has low creep and low thermal shrinkage, meaning it maintains its length, which is important for the stability of the gondola. The ropes are covered in aluminized mylar to protect them from UV degradation. They are attached to the gondola via turnbuckles, which can be adjusted to balance the outer frame and pivot. The lifting angle of the two front ropes is increased to 90 degrees by the use of a spreader bar. This removes the horizontal tension in the ropes, and reduces the lateral force on the outer frame of the gondola. Like the rest of the gondola frame, the spreader bar is a carbon fiber tube with aluminum inserts at the ends through which Two shorter ropes then connect the spreader bar to the pivot.

3.2.2 Motion Control

Azimuthal Control

The outer frame moves in azimuth using a multi-spoke reaction wheel, mounted below the gondola floor, and a pivot that interfaces with the balloon flight train (Figure 3.4).

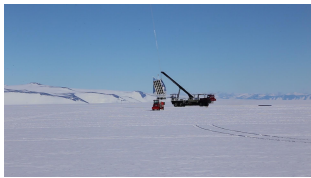

Figure 3.4: SPIDER azimuth motors. Left: Pivot; Right: Reaction wheel.

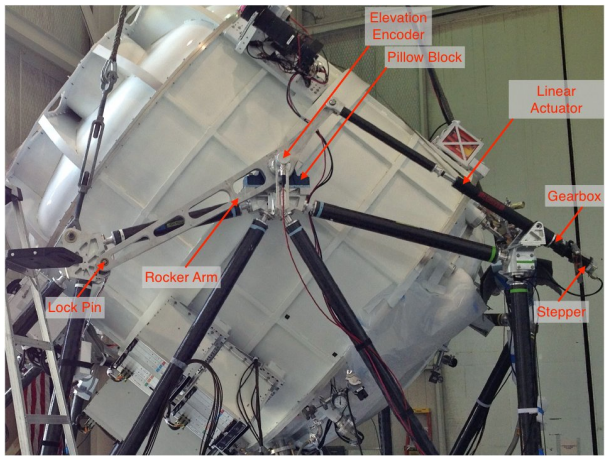

Figure 3.5: SPIDER elevation drive.

Elevation Control

Figure 3.5 shows the various components of the elevation drive. The inner frame moves in elevation using port and starboard rocker arms, each connected to a linear actuator driven by a stepper motor. The rocker arms interface with the cryostat at two points, at which elevation encoders are mounted to measure motion in elevation. Lock pins are used to lock the rocker arm in place.

Scan Strategy

SPIDER scans in azimuth with a sinusoidal velocity profile, reaching a peak velocity of $6^\circ/\mathrm{s}$ in the middle of the scan. The scan width is 90° peak to peak. Rather than drifting continuously as in BLASTPol, SPIDER's elevation is only stepped at every n scan turnarounds, where n is a commandable number. SPIDER has only one scan target, which is all the observable sky from

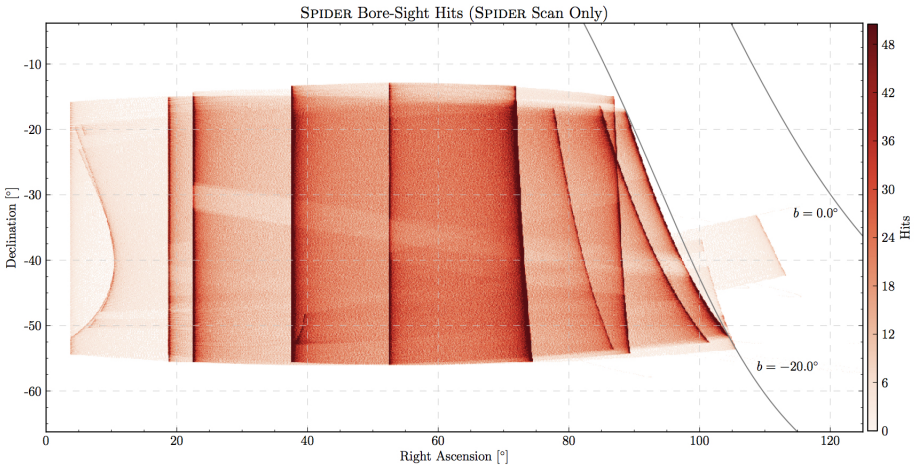

Figure 3.6: Hits map for SPIDER in-flight scans (in RA and Dec).

Antarctica, given the limitations of avoiding the Sun, the galaxy, the balloon and the ground, with additional mechanical constraints of the elevation drive and any other mechanical interferences. The observable region is a box that covers $\sim 10\%$ of the sky, with corners defined by coordinates in RA and Dec, and connected by great-circle arcs. The resulting hits map over 24 solar hours over this region is shown in Figure 3.6.

3.2.3 Coarse Pointing Sensors

Since the SPIDER detector array is so large ($\sim 10^{\circ}$), the pointing solution provided by the coarse sensors has sufficient accuracy to point the telescope in flight. The star cameras are only used for post-flight pointing reconstruction. All of the coarse sensors used on SPIDER were test-flown on BLASTPol, either in 2010 or 2012.

GPS

SPIDER uses the same Septentrio PolaRx2e@ GPS receiver unit that was flown on BLASTPol in 2010. Instead of the Septentrio PolaNt antennas, SPIDER uses Antcom G5Ant-3AT1 antennas provided by CSBF. They are smaller and have higher band rejection compared to the Septentrio antennas. The antennas are mounted to phenolic wood plates attached to aluminum tubes on the sunshield frame. Two antennas are mounted at the back of the sunshields on tubes that point directly outward. The third antenna is mounted on a tube that points outward from the tip of the sunshield wing. This creates an approximately L-shaped array with a *main-auxiliary1* baseline of 4.6m, and a *main-auxiliary2* baseline of 4.2m, with $\sim 90^{\circ}$ degrees between the two

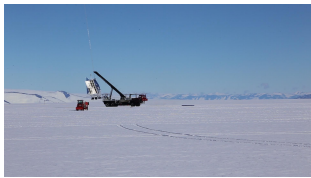

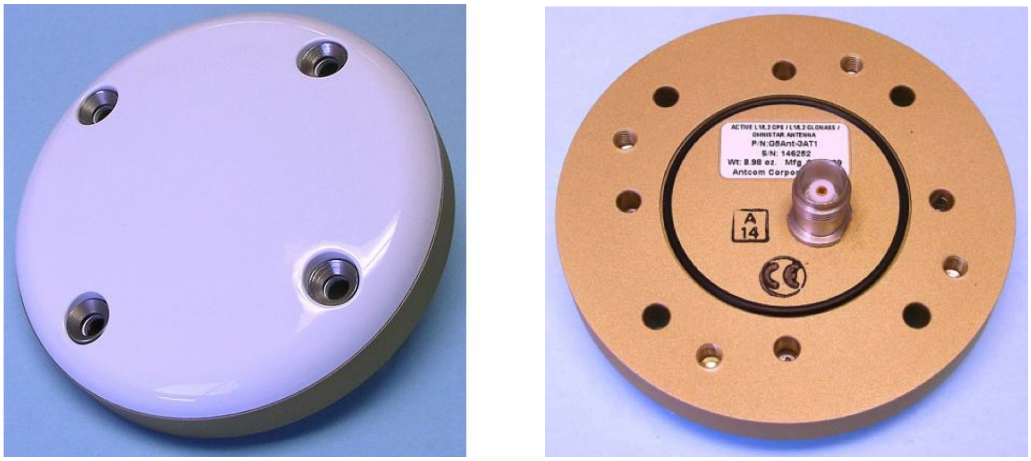

Figure 3.7: Antcom G5Ant-3AT1 GPS antenna.

baselines. The expected accuracy with this antenna spacing is $< 0.1^\circ$ in heading, and $< 0.2^\circ$ in pitch and roll. Figure 3.3 shows the locations of the SPIDER GPS antennas.

As in BLASTPol, the COM1 serial output of the GPS is set up to report the message blocks described in Section 2.8.4. For the 2014 flight, the GPS receiver is mounted on the gondola floor, on top of the serial hub, and powered directly from the ACS without a DC-DC. One pair of SIP commands is allocated to switching GPS power on or off.

Sun Sensor

SPIDER uses an array of six pinhole sun sensors, each of the same design described in Section 2.8.4. The PSS's are mounted at 25° elevation using two small aluminum "feet", which are glued to the square PSS case. The six PSS's are then glued onto an aluminum honeycomb mounting plate, spaced in azimuth at 36° from each other, in order to overlap their 40° azimuthal coverage by 2° degrees. The entire array covers an azimuthal range of 240°. The mounting plate is bolted to a 3 foot long aluminum boom which is attached to the bottom rear port joint of the sunshields. The plate is mounted on the boom at an azimuth of 200° offset from boresight, placing the centre of the sun sensor array in the middle of SPIDER's expected range of sun exposure. The sun sensor is powered through a DC-DC, which is mounted inside a box attached to the gondola floor, to keep it out of the sun. The DC-DC provides the +/- 12V required by the sun sensor readout circuit (as shown in Figure 2.16).

Magnetometer

The SPIDER magnetometer is a Honeywell HMR2300 three-axis digital magnetometer. It uses three magneto-resistive sensors (HMC 1001) to measure the orthogonal X,Y and Z components

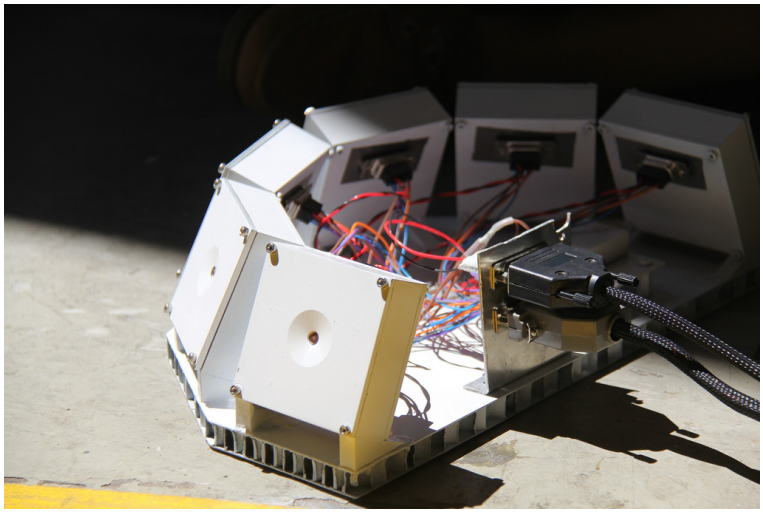

Figure 3.8: PSS array.

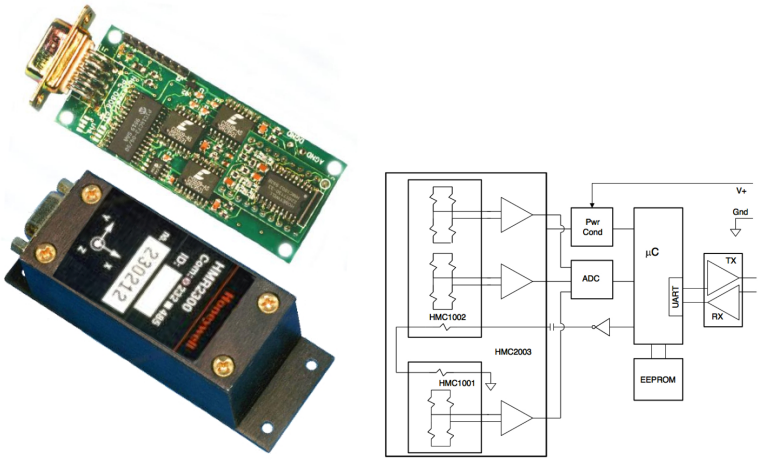

Figure 3.9: SPIDER magnetometer.

of the magnetic field. Unlike the analog magnetometers flown on BLASTPol, the HMR2300 has an internal analog-to-digital converter, and outputs in serial RS-485. A RS-485 to RS-232 serial converter is used to communicate with the flight computers over the RS-232 serial hub. The magnetometer is set up to output messages in ASCII format, with a continuous output rate of 20Hz. Configuration settings are stored by an onboard EEPROM. The magnetometer is mounted on the joint of the bottom front sunshield tube, which is not a structural part of the sunshields and is only used during sunshield assembly. The 2.4-meter carbon fiber tube is cantilevered off the bottom front port sunshield joint (see Figure 3.2), and stabilized by tying it to the gondola with KEVLAR string.

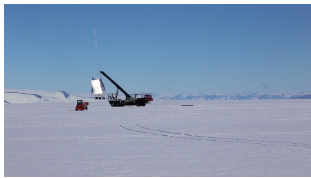

Clinometers and Gondola Balancing

SPIDER has two clinometers, both Model 904-TH from Applied Geomechanics. One is mounted in the pivot motor control box, which is attached to the bottom face of the pivot, and measures the pitch and roll of the pivot. The second clinometer is mounted inside the ACS crate and measures the pitch and roll of the gondola floor.

During ground operations, the readings from these clinometers are used to balance the gondola. In order to minimize pitch and roll pendulations of the gondola while scanning, the pivot and outer frame are made to be as level as possible. Because the star cameras are only able to track azimuthal motions of the gondola, and not pitch and roll motions, the latter types of motions must be minimized in order to prevent stars in the images from streaking across the CCD. In order to balance the gondola, first the outer frame is balanced using weights (e.g. sandbags, lead shot containers, lead bricks). Then the pivot is leveled by adjusting one or more of the three turnbuckles attaching the gondola ropes to the gondola. This changes the balance of the outer frame, so the two steps are repeated until the balance converges on an acceptable value, usually a few tenths of a degree in pitch and roll.

3.2.4 Star Cameras

SPIDER uses three star cameras in order to reconstruct its pointing after the flight. Two are mounted on the outer frame, and one on the cryostat. The outer frame cameras sit on a rotating platform at the front of the gondola floor. These rotating star cameras stay fixed on the sky as the gondola scans sinusoidally back and forth in azimuth, with a peak azimuthal velocity of $6^\circ/\mathrm{s}$. The rotating cameras take pictures every ~ 4 seconds. SPIDER also has one boresight star camera, mounted near the top of the cryostat, which takes pictures at every scan turnaround (every ~ 30s), when the gondola is stationary in azimuth, and before the inner frame is stepped in elevation. The bore sight camera is named "The Good" and the rotating star cameras are named "The Bad" and "The Ugly".

Components

Each star camera consists of multiple components, selected to satisfy the pointing requirements and to operate reliably during the flight. All three star camera units are identical and completely interchangeable, and differ only in how they are mounted to the gondola, and how often they are commanded by the flight computers to take images.

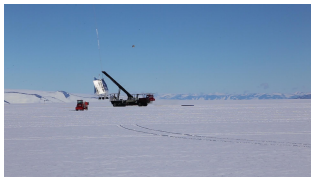

Figure 3.10: SBIG CCD camera (left) and Sigma lens (right) used in SPIDER star cameras.

CCD Camera

The cameras are ST-1603 ME Santa Barbara Instrument Group (SBIG) CCD cameras with KAF-1603ME detectors. The specifications are listed in Table 3.1. The cameras were chosen for their high well depth, allowing for longer exposures without saturating. Because the sky is not moving across their field of view during an exposure, the exposures can be longer than the BLASTPol star cameras. During the BLASTPol 2012 flight, the SPIDER star camera took 200ms exposures. The CCD's have 9x9 μm pixels with an array size of 1530x1020 pixels. They have a peak quantum efficiency of $\sim 80\%$ at 650nm, and saturate at $\sim 100,000$ e-. Communication with the camera is over USB 2.0, with a full frame download time of $\sim$ 2s. The linux based SBIG Universal Driver Library is used to send commands to the camera.

camera	SBIG ST-1603ME
CCD	Kodak KAF-1603ME
Pixel Size	$9 \times 9\mu$m
Pixel Array	1530×1020 pixels
Read Noise	18e- RMS
Well Depth	$\sim 100,000$e-
Digital Output	16 bits

Table 3.1: SPIDER Star Camera.

Optics

The camera is coupled with a Sigma 120-300mm f/2.8 zoom lens. A 200 mm focal length gives the camera a $4^\circ \times 2.7^\circ$ field of view and a pixel scale of $\sim 9.3''$/pixel. In order to reduce the background level due to sky brightness, a Schott RG600 filter is used to cut off wavelengths shorter than 600nm.

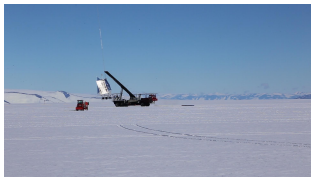

Figure 3.11: Birger electronic lens interface.

During ascent, large temperature fluctuations cause the focus of the lens to change. Therefore a mechanism is needed to be able to focus the camera once the balloon has reached float altitude. This is done using a Canon EF-232 electronic lens interface made by Birger Engineering that mimics the interface of a Canon camera, allowing the Sigma lens to be controlled over USB. The lens interface can control the focus position as well as the aperture position, and also respond to queries about the current focus and aperture positions. The focus position is included in the header of each image taken by the camera. The lens interface can also report the current focal length of the lens, which nominally is set to 200mm, but could change if the lens is moved unintentionally during handling. Knowledge of the focal length is helpful in determining the pixel scale of the image, which is needed in order to properly determine a pointing solution.

Baffles

In order to block stray light, each star camera is equipped with a baffle mounted in front of the lens. The baffle is designed to be as light as possible while also being extremely rigid. It consists of multiple knife-edge 0.20in aluminum disks, laser cut to the right inner and outer diameter. The spacing of the disks is designed to eliminate direct illumination from sources $> 10^\circ$ from the optical axis. Figure 3.12 shows how the spacing of the disks was determined. The outer disk is placed 23.5in away from the base, giving a 10° margin over the 7.91° field of view of the baffle. Then the next disk is placed where A-B intersects the baffle field of view. And so on for the rest of the disks. This prevents the camera from directly seeing any part of the baffle wall. The disks are held in place by a carbon fiber truss, made of Easton Lightspeed size 500 arrow shafts cut to the desired length. The shafts are slotted through holes in the disks and glued in place with ScotchWeld adhesive epoxy. The baffle is constructed using an aluminum jig that holds the baffle discs in place at the correct spacing while the glue cures

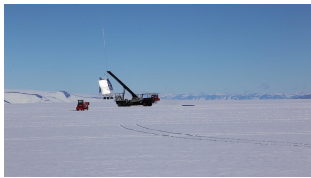

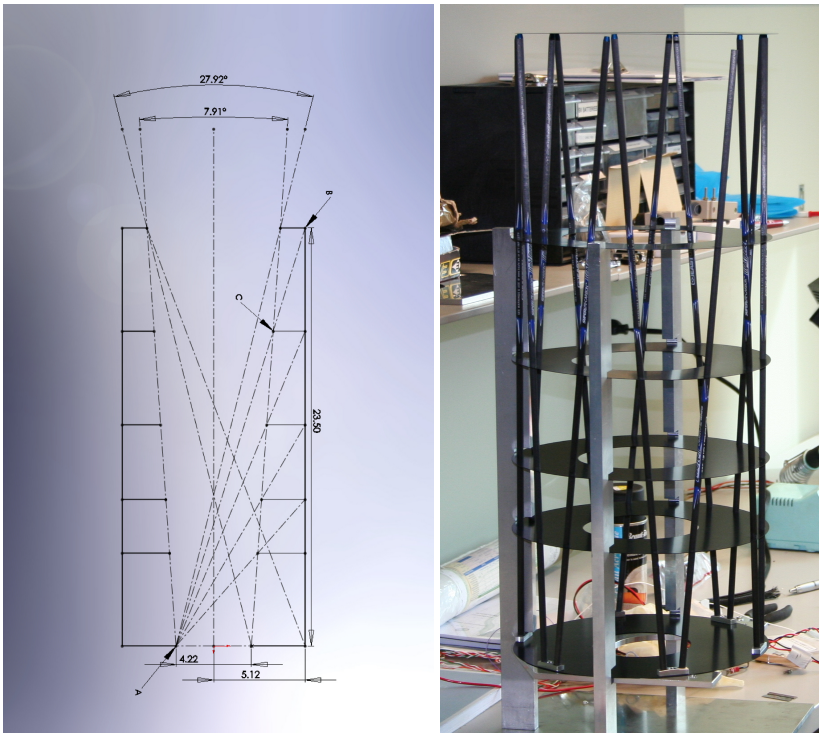

Figure 3.12: Left: Baffle sketch (units in inches), Right: Baffle assembly.

(Figure 3.12). The ends of the rods sit in circular depressions in small aluminum pieces. The middle of each piece has a threaded center hole to attach the baffle to a baffle holder, which is a U-shaped hexcel plate mounted in front of the lens (Figure 3.13). Before flight the baffle is wrapped in aluminized mylar, with the inside surface spray painted black. The space between the lens and the first baffle disk is covered with black electrical tape.

Computer and Software

Each camera has its own PC-104 single board computer, booting from a compact flash card. A SSD is mounted as a data drive to store images during the flight. The computer runs software that interfaces over USB with both the camera and the lens, and that communicates with the SPIDER flight computers over ethernet. The software performs the same functions described in Section 2.8.2 (downloading images from the CCD, processing them in order to identify stars). Unlike for BLASTPol, SPIDER does not require pointing solutions from the star cameas, and star identification is only useful to ensure the proper functioning of the star cameras. In order to save computer resources, a command was written that could turn off the Pyramid search once the cameras were shown to be working properly.

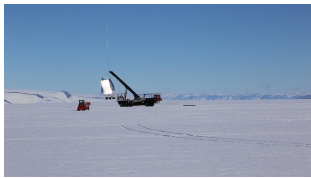

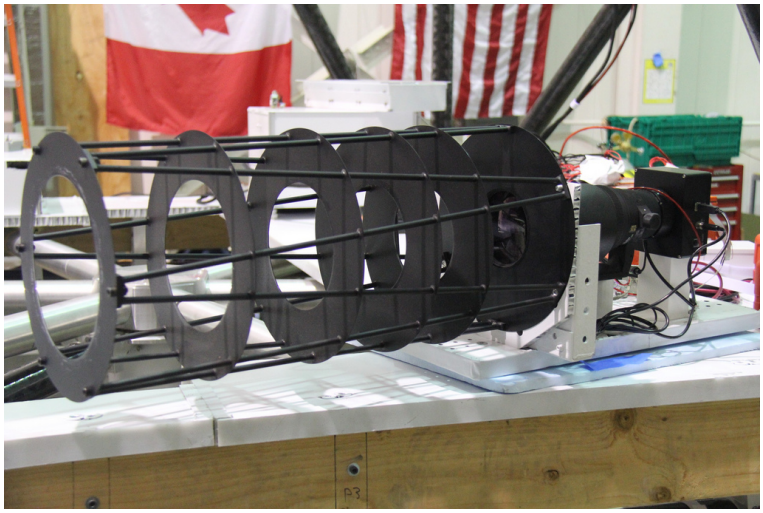

Figure 3.13: Baffle mounted in front of the star camera.

Video

Video output from the bore sight star camera computer is connected to a VGA converter which outputs the video in NTSC format. Unlike in BLASTPol, the NTSC signal is sent directly over coaxial cable to the CSBF LOS video transmitter, and no relay box is required to switch between cameras. Instead, the boresight star camera computers runs a video display program that displays the most current image taken by one of the three star cameras. The two rotating star cameras send their most recent image to the boresight star camera over the local gondola network. Which of the three star camera's image is displayed is determined by a command sent to the bore sight star camera. Since the video output is the boresight star camera computer's video output, it is not possible to see the video output of the two rotating star cameras, only the images that they send to the boresight camera. So for example, there is no way to watch one of the rotating star cameras booting up in LOS video stream.

Rotating Star Cameras

The rotating cameras are mounted to an aluminum platform on an Aerotech ADRT 150-135 direct drive rotating stage (Figure 3.14). A 16-bit encoder in the motor measures the position to an accuracy of $4''$. The motor is controlled using a Technosoft IDM240-5EI Intelligent Servo Drive. The motor speed is commanded over RS-232 serial by the flight computers. While the star cameras are exposing, the stage rotates at a speed to counter the azimuthal rotation of the gondola, as reported by the gyros. The full with of a typical SPIDER scan is $90°$, however the opening at the front of the gondola frame through which the cameras can see is $\sim 40°$. Therefore after each exposure the rotary stage returns the cameras to a forward facing position.

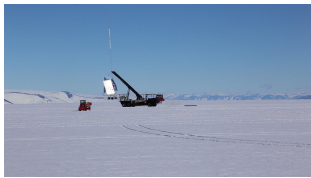

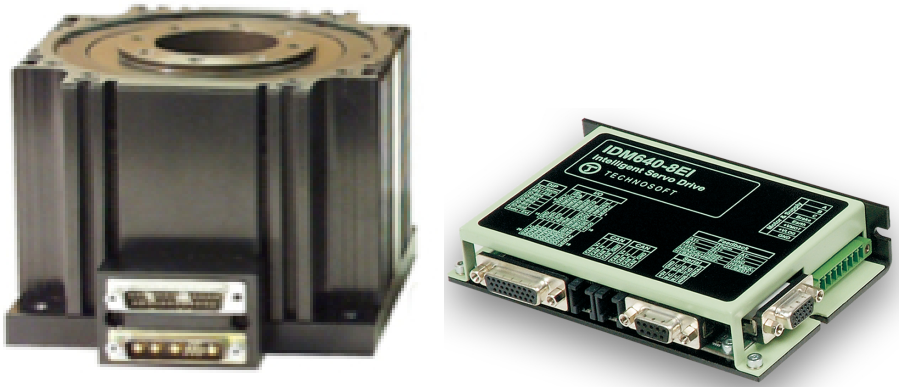

Figure 3.14: Motor and drive used by the rotating star cameras.

Figure 3.15: Mounts for the rotating star cameras.

Mechanical limit switches prevent the cameras from rotating too far in either direction where they might collide with the cryostat.

Each camera (together with its lens) is mounted onto a triangular aluminum base, as shown in Figure 3.15. A Finite Element Analysis was performed in SolidWorks on the design of the mounts in order to ensure that given the stress they experience, they would not deflect by more than $4''$. All of the electronics (computers, DC-DCs) associated with the star cameras are mounted to the walls of the triangular bases, as well as to the platform. During flight, the electronics, cameras, and lenses are covered in blue foam in order to prevent them from freezing. The cameras are offset from each other by 30° in azimuth, and 20° in elevation. One camera is mounted at 30° elevation and the other at 50° elevation. Mounting the cameras at different angles allows the pointing solution from one camera to constrain the roll solution of the other camera.

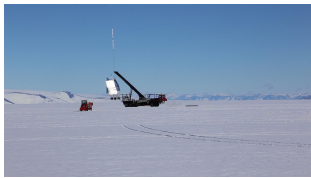

Chapter 4

Flight Performance

Figure 4.1: You've got to ask yourself one question: 'Do I feel lucky?' Well, do ya, punk?

4.1 Flights

The flight tracks of both BLASTPol Antarctic flights are shown in Figure 4.2. The 2010 flight lasted 9.5 days, and the 2012 flight lasted 12.5 days. Both flights took the telescope to an mean altitude of $\sim$ 38.5km, with daily variations of $\sim$ 2km peak to peak. Both flights terminated near the South Pole Traverse line, and the telescope was able to be recovered weeks after termination.

A number of problems occurred during the 2010 BLASTPol flight. There was difficulty with the reaction wheel motor, which led to poor azimuth control in-flight. Recovery of the telescope discovered partial damage to an infrared blocking filter, which was believed to have been melted by exposure to the sun on ascent. This degraded the beam shape of the telescope, and produced a high amount of instrumental polarization. Post-flight analysis also revealed correlated noise in all detectors, believed to be partially caused by interference from TDRSS

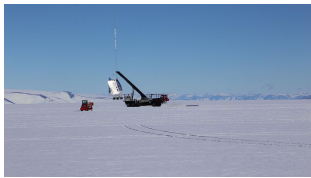

antennas. Due to all these issues, analysis of the data from the 2010 flight is not presented in Chapters 5 and 6. However, performance of various subsystems during that flight is included in this chapter. Table 4.1 lists the targets observed during the 2012 flight, and the amount of observing time for each target.

Target	Type	Distance (pc)	Size (degrees^2)	Time observed (hours)
Vela C	GMC	~ 700	3.9	43.91
Vela C Ref.	GMC	~ 700	14.0	10.80
Vela Wide	GMC	~ 700	23.1	89.49
Carina Nebula	Calibrator	~ 2300	2.7	4.19
CG 12	Dark Cloud	~ 550	0.1	1.72
G331.5-0.1	Calibrator	~ 7500	3.6	4.35
Lupus I	Dark Cloud	~ 155	1.7	15.38
IRAS 08470-4243	Calibrator	~ 700	1.0	4.94
Puppis	Nearby Cloud	~ 1000	1.4	13.18
VY CMa	Calibrator	~ 1200	0.4	4.97
Saturn	Calibrator	$\sim$ 4e-5	1.0	4.17

Table 4.1: Targets observed during BLASTPol 2012 flight.

4.2 Telemetry

During the 2012 flight, it became apparent after launch from looking at the detector timestreams that the detectors were sensitive to RF interference from CSBF telemetry antennas, specifically those transmitting data via the TDRSS satellite network. It was decided that the TDRSS antennas would be switched off, and data would only be downlinked over the Iridium network in dialup mode. The dialup connection is frequently lost (every couple minutes) and has to be re-established, and sometimes the link is down for several minutes. On January 1, 2013, one of the two SIP flight computers failed. After that it was decided that re-dialing the Iridium connection every few minutes would no longer be done, thus removing the only telemetry link to the payload. For the last few days of the flight, the TDRSS high gain antenna was turned on every 2 hours for 5-10 minutes in order to check telescope operation and send any necessary commands.

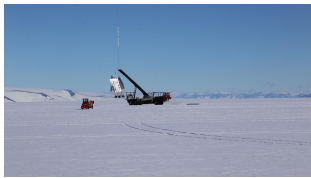

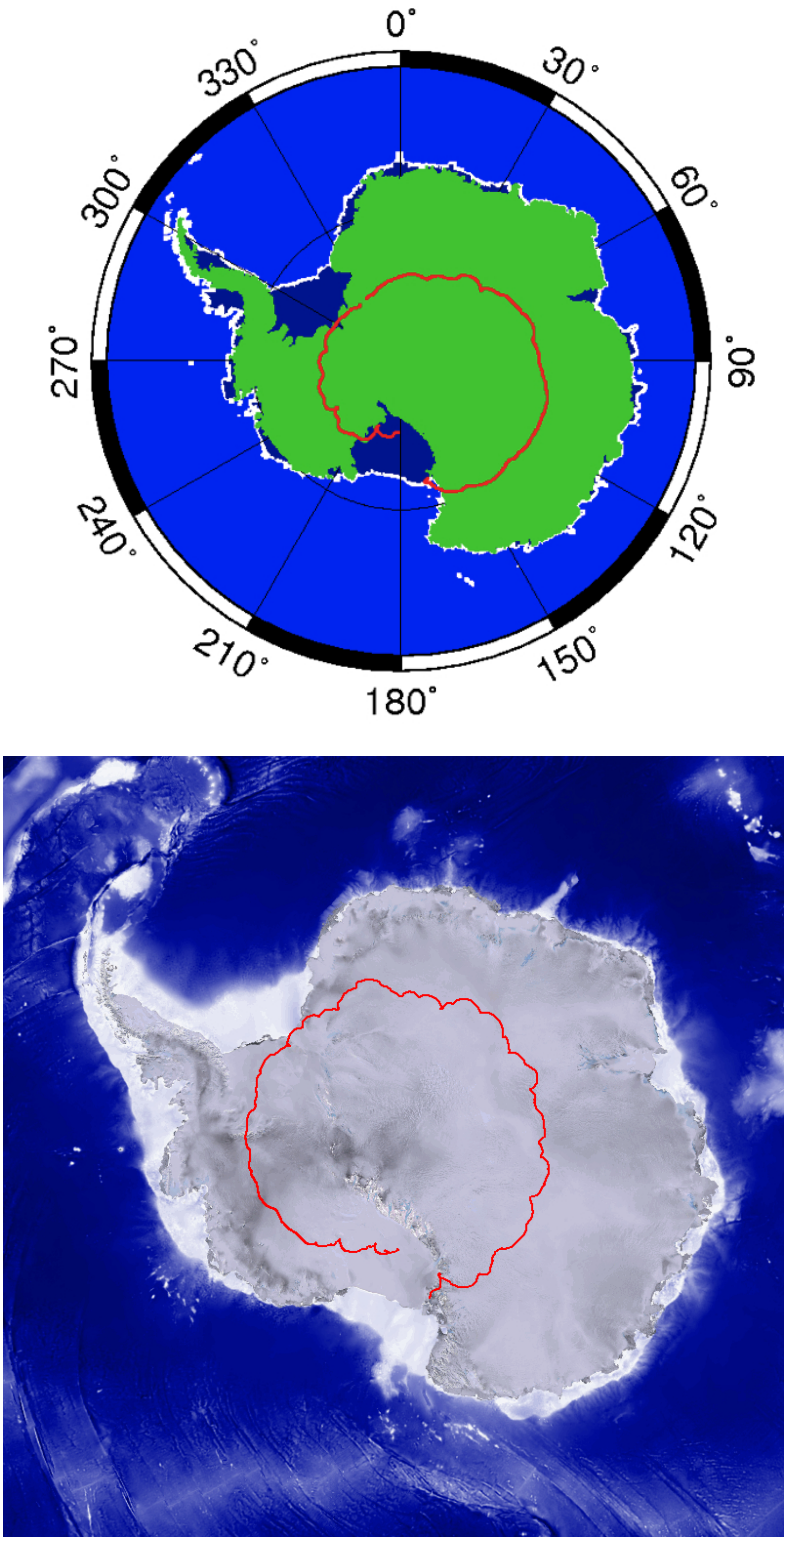

Figure 4.2: BLASTPol flight tracks provided by CSBF. The top panel shows the 2010 flight, and the bottom panel shows the 2012 flight.

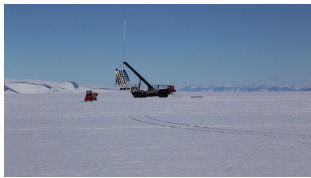

4.3 Pointing

4.3.1 Coarse sensors

After the flight, the performance of each coarse pointing sensor is evaluated by analyzing the difference between the raw azimuth provided by the sensor and the reconstructed star camera-based azimuth. Figure 4.4 shows histograms of the differences between the coarse sensors and the star camera for the BLASTPol 2012 flight. For each sensor, a Gaussian is fit to the histogram, and the full-width at half maximum of the fit is taken as the accuracy of the sensor. The accuracy of the combined in-flight pointing solution was $\sim$0.05° rms for the 2012 flight. Individual sensors are described below.

Magnetometer

As expected, the analog magnetometer provided the least accurate in-flight azimuth. The rms accuracy is seen to vary over the flight, due to the changing proximity of the payload to the south magnetic pole. The rms accuracy of the magnetometer over the first half BLASTPol 2012 flight is $\sim$6°. Despite its low accuracy, the magnetometer remains the most reliable sensor, with no failures of any kind occurring in both BLASTPol flights, as well as in previous BLAST flights. Therefore it is recommended that one should always be flown.

GPS

Both GPS units performed as expected given the size of their antenna arrays. The performance of the Septentrio GPS on the 2010 BLASTPol flight was sometimes affected by the inner frame baffle of the telescope. After being shielded in Mylar, the baffle became a large reflective surface, resulting in large GPS error any time the inner frame was pointed above 40° in elevation. The multipath effect is clearly visible when looking at the dependence of the GPS accuracy on the elevation of the inner frame (Figure 4.3). When not being blocked by the inner frame baffle, the attitude was accurate to $\sim$0.14° rms. The azimuth measured by the ADU5 in the BLASTPol 2012 flight differed from the star camera azimuth by $\sim$0.33° rms (Figure 4.4), which is expected for a 1m antenna array.

Pinhole Sun Sensor

During the 2012 flight, one of the sun sensors (PSS1) malfunctioned during the flight, and was not used. The other three sun sensors performed well and were the best available azimuth sensors during the second half of the flight when the ISC was dead.

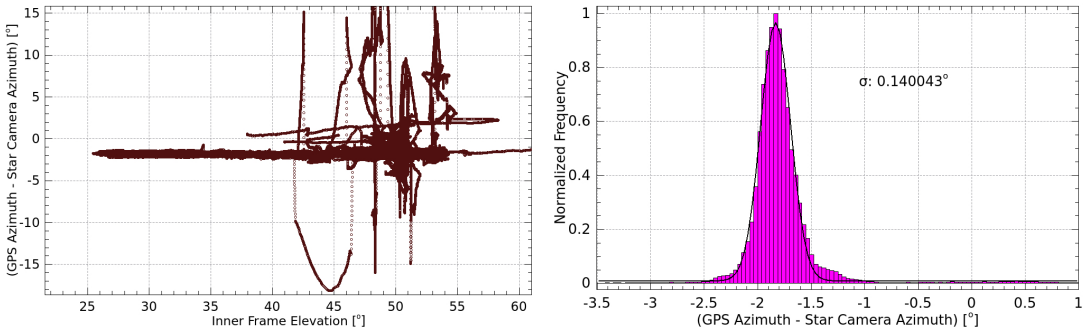

Figure 4.3: Left: GPS Azimuth error versus inner frame elevation angle. Right: GPS azimuth accuracy.

The accuracy of the pinhole sun sensors can be improved post-flight by adjusting various parameters that affect the azimuth calculation for each sensor. These parameters are defined in-flight according to design specifications and pre-flight calibration tests, but their actual value can be more accurately determined post-flight by doing a least-squares fit of the sun sensor azimuth to the star camera post-flight azimuth solution. Four parameters are left free in the fit. Table 4.2 shows the values found in the fit, and the values used during flight. The first is the distance, D, between the pinhole and the sensor, which is designed to be 10mm. The best fit solution found deviations from this by < 1mm. This value is adjusted for flight based on ground testing as well as adjustments made in-flight. Tabel 2.1 shows that the pre-flight and in-flight calibrations resulted in values that differed from the best fit post-flight value by < 0.5mm. The second parameter is the elevation mounting angle of the sensors, α, which is designed to be 25° (and has been measured to be 24.3° for PSS2). The best-fit post flight solutions show deviations of $\sim 3 - 4^\circ$ from the nominal values. The third parameter is the rotation of the sensor, ψ, which is assumed to be zero for the square sun sensors. PSS2 was a circular sun sensor built for BLASTPol 2010 flight, and is rotated at 11.0°. No attempt is made to measure the actual value pre-flight. The post-flight fitting shows minor deviations from zero of $\sim 1^\circ$. The fourth parameter is the azimuth mounting angle offset between the telescope boresight and the sensor. These are nominally the azimuths listed in Table 2.1. The best-fit values are almost identical to the in-flight values, with deviations of only a few tenths of a degree. After post-flight calibration, the accuracy of the pinhole sun sensors was found to be $\sim$0.08° rms.

sensor	PSS2		PSS3		PSS4	
	in-flight	post-flight	in-flight	post-flight	in-flight	post-flight
D (mm)	9.90	9.64	10.10	9.64	9.79	10.02
ψ (°)	11.0	11.0	0.0	1.5	0.0	0.9
β (∘)	84.376	84.459	121.866	121.937	156.786	157.468
α (°)	24.3	21.3	25.0	21.5	25.0	28.5

Table 4.2: Post-flight fitting of parameters used in PSS azimuth calculation, compared to the values of those parameters used in-flight.

4.3.2 Star Cameras

ISC and OSC

During the 2010 BLASTPol flight, the ISC and OSC had to be rebooted on average once a day. The cause of this is thought to be problems with the program running on the star camera computers, *netisc*, that would cause it to stall without restarting. During the 2012 integration in Texas, the ISC and OSC were left running for hours/days in order to find and fix any remaining issues with *netisc*. Six days into the 2010 flight, the ISC suddenly stopped reporting data to the flight computers. During the rest of the 9-day flight, the ISC was power cycled several times but never appeared to recover. The ISC hard drive was found to be dead when retrieved after the flight, but attempts to recover the drive and access the data were unsuccessful.

During the 2012 flight, both the ISC and OSC images showed spurious bright pixels in the corners of the images that became visible in the downlink video once the telescope reached float altitude. Approximately 12 hours after launch, during which the payload was still in line-of-sight, the OSC star camera stopped reporting data. A system reboot was done, and the video downlink displayed an error message indicating that the star camera computer could not find a boot device. Several more reboots were attempted, and they all failed to boot from the hard drive. Approximately 6.5 days into the flight, the ISC stopped reporting data. Post flight recovery of the ISC hard drive also indicates that it failed in a similar manner to the OSC. The failure is believed to be due to attributable to the brand of hard drive (Intel320) used by both the ISC and OSC, which has a known issue with the drive controller which might have been triggered by a cosmic ray. All other drives on the flight were of a different model, and none failed, including the SPIDER star camera drive and the two flight computer drives. A separate experiment, SuperTIGER, that was in the stratosphere above Antarctica at the same time as BLASTPol also had all six of their Intel320 hard drives fail.

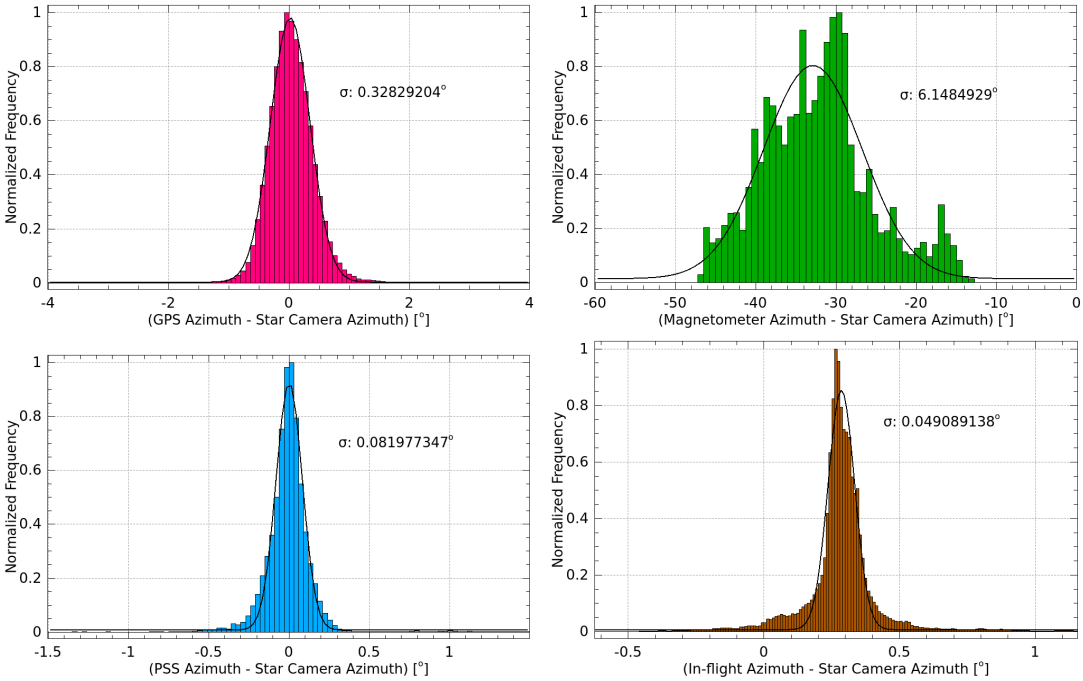

Figure 4.4: Accuracy of coarse sensors and combined in-flight pointing solution. Top Left: ADU5 GPS. Top Right: Magnetometer. Bottom Left: Pinhole Sun Sensor. Bottom Right: In-flight combined azimuth.

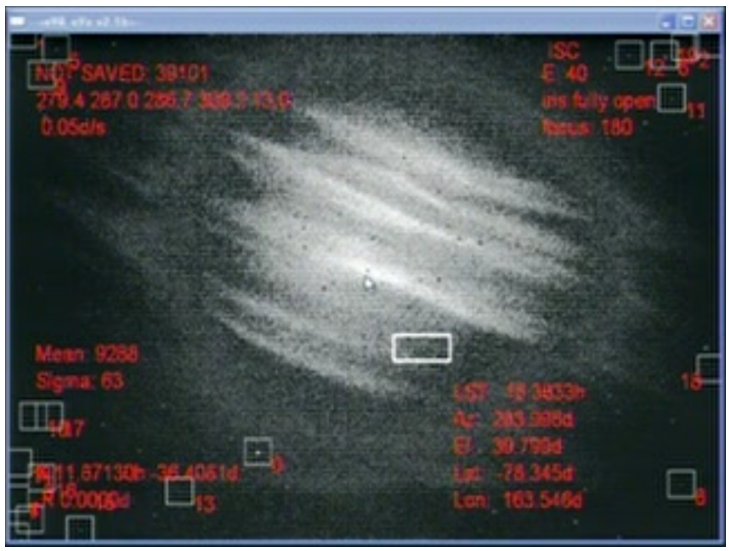

Figure 4.5: Corner blobs and streaks

SPIDER Star camera

The first test of the SPIDER star camera on a stratospheric balloon flight was on the 2010 BLASTPol flight. Although the camera reported identifying "blobs" throughout the flight, it was determined after the flight that none of the blobs were in fact stars. The blobs seen by the camera were most likely due to cosmic rays. The autofocus routine was run twice at the start of the flight, and once again on day six of the flight, however the routine was confused by the cosmic rays and the camera remained out of focus. Since the camera was not designed to calculate pointing solutions in flight, this failure went undetected. The line-of-sight video downlink could not be used to diagnose the failure due to improper settings on the star camera computer display that resulted in the images not being shown fullscreen.

The failures of the camera's first test flight led to several improvements in the system. The autofocus routine now checks that the blob being used to calculate the best focus position is in fact a star. It determines this by ensuring that the blob does not move more than 5 pixels between the three images taken at the best focus position and the images taken before and after. The display software was also improved, with full screen images now displayed with blobs labeled in order of highest to lowest flux. The video baluns used in 2012 also drastically reduced the noise in the downlink images.

Figure 4.6 shows the mean brightness of the camera image in ADU (for a 16-bit converter, saturation is at 65535) during ascent of the balloon. The exposure time is 100ms, except around the 5th hour of ascent in 2012 when various other exposure times were tried. It s unknown why the CCD remained saturated up to a higher altitude in 2012, but at the float altitude of ~ 40km it appears to be at the same level (~ 12000 ADU) in both flights.

Upon reaching float, the autofocus routine was done, but a focus position was not able to be found due to the presence of non-stellar features in the field of view (see Section 4.3.2). Focusing was then done manually, and it was determined that the best focus position was beyond infinity, so the focus was set to infinity. The cause of the focus being beyond infinity was due to elongation of the focal length by the presence of a 600nm fliter. Since the filter was mounted between the camera and the lens adapter, it increased the space between the camera and lens as the lens adapter could not be screwed in all the way. A modification made for the SPIDER flight was to place the filter in the hollow space inside the lens adapter, allowing the adapter to be fully threaded to the camera. Although the maximum focus position could not be obtained for the 2012 flight, this did not prevent stars from being identified (see section on SBSC pointing solution).

When the ISC died on day 6 of the 2012 BLASTPol flight, an attempt was made to improve the performance of the SPIDER star camera. Because the telemetry link was insufficient to be

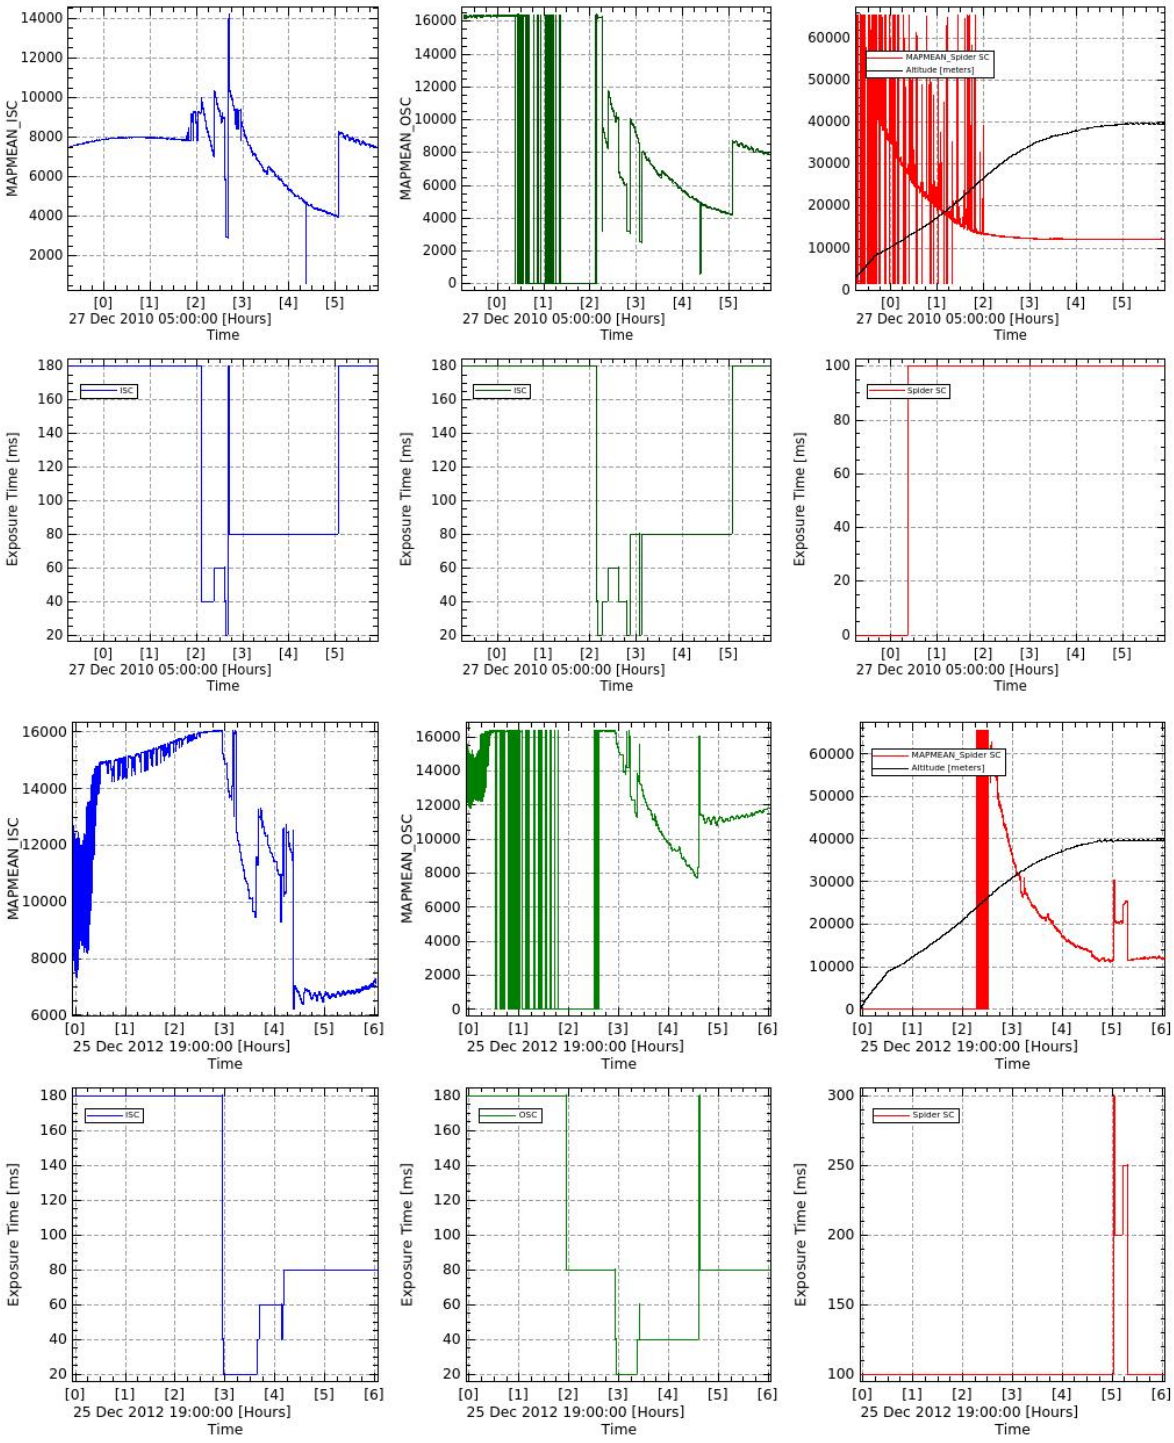

Figure 4.6: Mapmean and exposure times of the star camera images during ascent. Also plotted in the Spider SC plot is altitude. Top: BLASTPol 2010 flight; Bottom: BLASTPol 2012 flight.

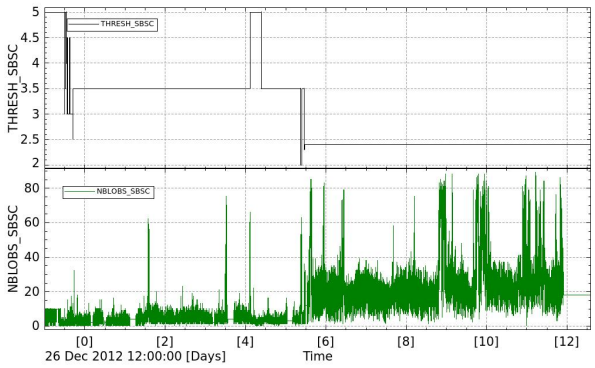

Figure 4.7: Top: threshold for identification of blobs; Bottom: Number of blobs identified.

able to continuously diagnose the star camera performance, it was decided that the threshold for blob detection would be lowered. While information is stored for the 15 brightest blobs, before the ISC died the SPIDER star camera was often detecting less than 15 blobs. While not all blobs detected are stars, lowering the detection threshold increased the number of blobs detected to be always greater than 15, which increased the probability of stars being among the 15 blobs. In order to diagnose whether the camera was detecting blobs on the sky or artifacts in the CCD, such as the cosmic rays detected in 2010, the focus position was moved back and forth and a decrease in the number of blobs was observed. Since the appearance of cosmic ray-induced blobs would not be affected by focus position, it was determined that the blobs being detected were stars.

During the 2012 flight, the SPIDER star camera ceased operating on several occasions, ie. the frame number stopped increasing, which is an indication that the star camera program is not running. This did not happen at regularly spaced time intervals, but typically happened after the frame number reached ~ 8000. On these occasions the star camera was power cycled, and resumed normal operation. It was determined during SPIDER integration that the cause of this failure was a memory leak in the star camera software. This was fixed before the SPIDER flight.

Mesospheric Clouds

During line-of-sight, the images transmitted over the video downlink from the star cameras displayed numerous strange features (e.g. streaks, ripples). These were initially thought to be caused by light leakage from faulty star camera baffles. However a separate experiment, EBEX (the E and B Experiment), which was also in the stratosphere above Antarctica at the same time as BLASTPol, reported seeing similar features in their star camera images. The presence of these features interfered with the SPIDER star camera's ability to identify stars,

since large portions of the features would be identified as blobs instead (Figure 4.8, top).

During line-of-sight, only the spatially filtered images are transmitted to the ground, but after the flight, if the star camera hard drives are recovered, the unfiltered raw images can be examined. Because of the hard drive failures of the ISC and OSC, all stored images were lost. The SPIDER star camera hard drive was recovered and all stored image files were intact. Figure 4.5 shows a line-of-sight image from a BLASTPol and Figure 4.8 shows line-of-sight images from the SPIDER star camera, as well as two stored images recovered from the SPIDER star camera hard drive. The structures seen in these images are consistent with other images of polar mesospheric clouds, or noctilucent clouds. These clouds are composed of a diffuse layer of water ice crystals in the mesosphere. They are more commonly observed at high latitudes (above 70° to 75°), and increase in occurance during the polar summer, when the temperature in the mesosphere is coldest. Mesospheric clouds are not well understood, but it is believed that their occurrence is increasing, possibly due to climate change, or an increase in space traffic. They were not observed in the star camera images during the BLASTPol 2010 flight, or during the 2006 Antarctic BLAST flight. The clouds are observed in a small fraction ($< 10\%$) of the raw SPIDER star camera images. A post-flight pointing solution was still able to be constructed (Section 5.3). However if the brightness and frequency of these clouds increases, then a solution will have to be devised to ensure that stars can still be identified. The brightness of the clouds can be tracked using daily images from the AIM (Aeronomy of Ice in the Mesosphere) satellite [1].

[1] http://lasp.colorado.edu/aim/browse-images.php

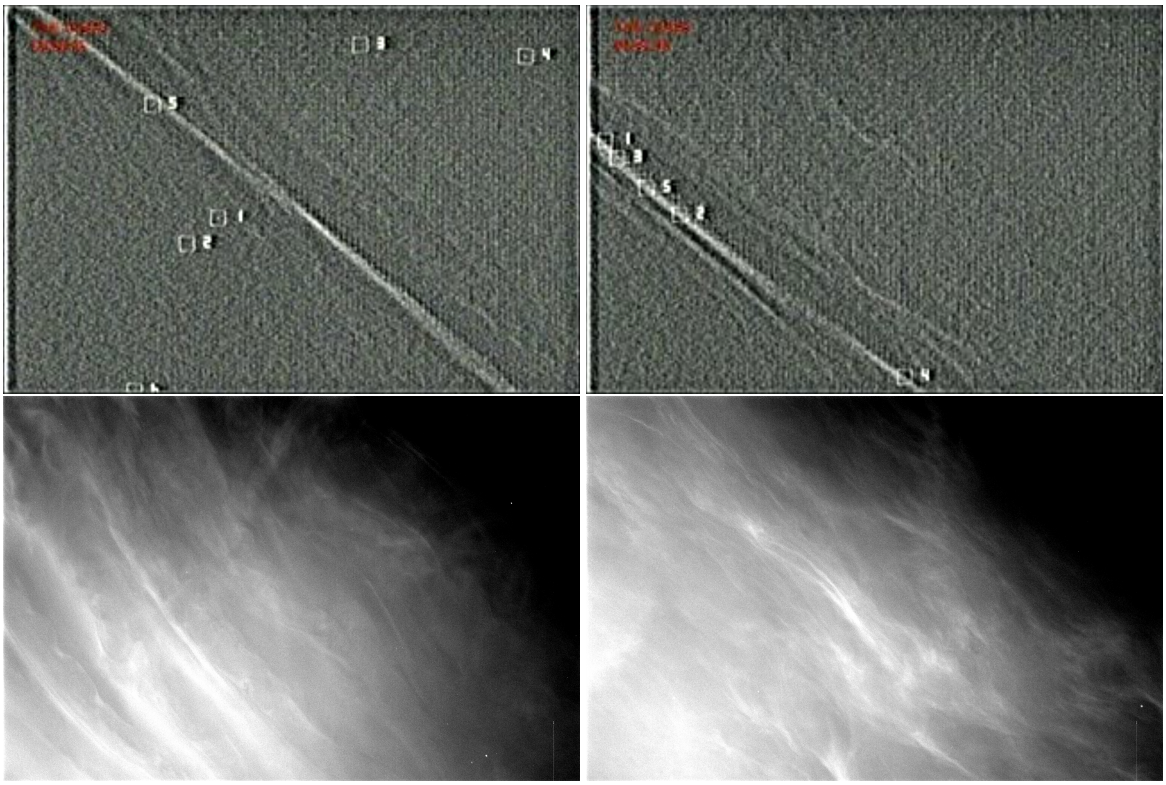

Figure 4.8: Screen captures of images transmitted over line-of-sight from the SPIDER star camera (top left and top right). Raw images recovered from the SPIDER star camera hard drive (bottom left and bottom right). The features in the images are thought to be mesospheric clouds.

Chapter 5

Analysis

Figure 5.1: See, in this world there's two kinds of people, my friend: those with loaded guns, and those who dig. You dig.

5.1 Timestream Processing

The reduction pipeline that is applied to the raw flight data is described in detail in Wiebe (2007) and Fissel (2013). The preliminary steps include merging of the datasets from the two flight computers, cosmic ray despiking and removal of data glitches, flagging of data where there was interference from CSBF transmitters, deconvolution of the detector transfer function, and fitting of detector response to calibration pulses.

5.2 Pointing Solution

The post-flight reconstructed pointing solution is meant to be based solely on data from the boresight star cameras and the gyroscopes. The methodology for determining the pointing

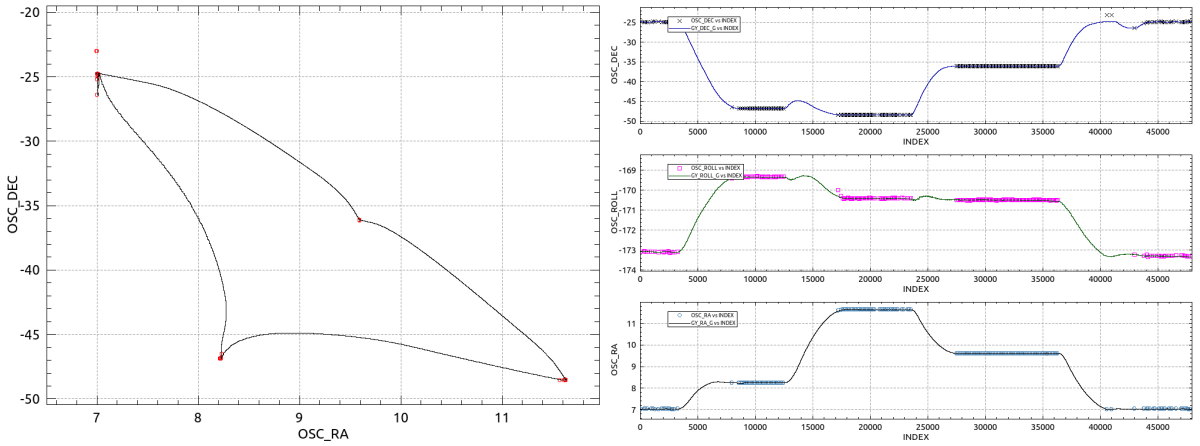

Figure 5.2: Slews. The solid line is the gyro-integrated pointing solution. The points are OSC solutions.

solution was developed for the BLAST 2006 flight and is described in detail in Pascale (2007).

5.2.1 Gyroscope Rotation

In order to use the gyroscope data to integrate between star camera solutions, the rotation from the gyros to the star camera reference frame must be calculated. This is done using data from large slews performed during the flight. At the end of the 2010 flight, four large slews were done for the purpose of the gyro rotation calculation (see Figure 5.2). At the end of each slew, enough time is given for the star cameras to get pointing solutions before beginning the next slew. Starting from the RA,Dec star camera solution at the beginning of the first slew, a pointing solution at the end of the first slew is calculated based solely on integration of gyroscope rate data. This gyroscope-based pointing solution differs from the solution that the star cameras report at the end of the first slew. A least-squares fit is done using Powell's method (Powell 1964) to calculate the gyroscope gains, rotations, and offsets that minimize this error over all four slews. Figure 5.2 shows the residuals in the fit: after the 3rd slew, which lasts 80s, the gyroscope-integrated solution is off by $\sim 0.48'$ in DEC, $\sim 2.4'$ in ROLL and $\sim 0.18'$ in RA. The residual in roll is the largest, and well above the gyroscope angle random walk noise of $4.0''\sqrt{s}$ or $\sim 0.60'$ in 80s. This likely reflects the inaccuracy of the star camera roll solution, which is $\sim 0.10^{\circ}$.

In the 2012 flight, both the ISC and OSC died before large slews could be performed for this purpose. Therefore the calculation was done using slews that were executed during the course of regular science scans.

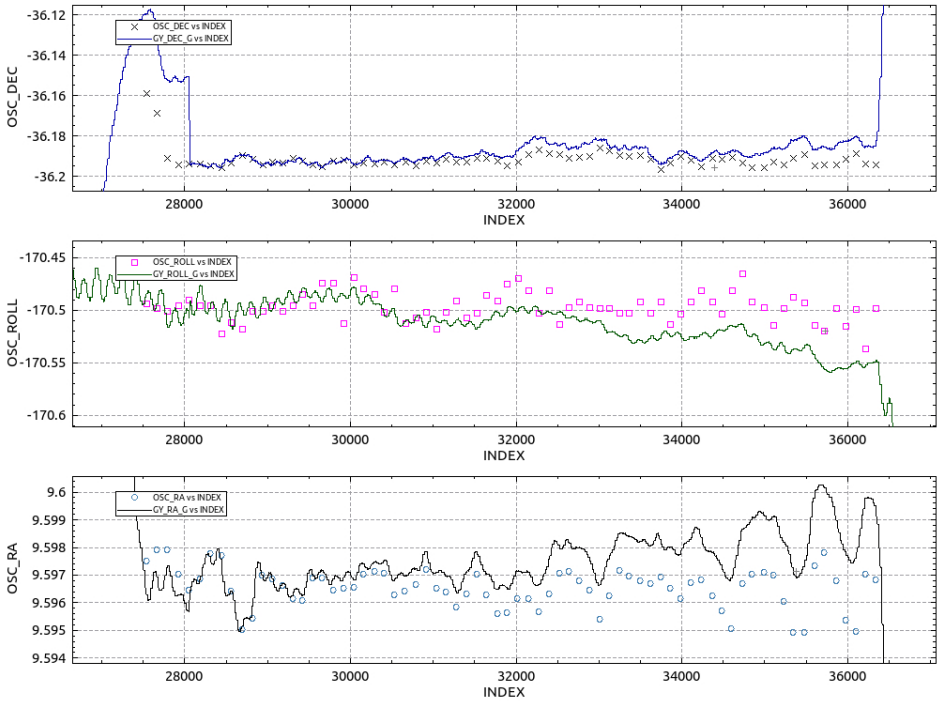

Figure 5.3: Slews

5.2.2 Star Camera Synchronization

After each image, the star cameras send data packets to the flight computers containing all relevant information about that image. Because of the time it takes to process the image and send the packet, the time at which the packet information appears in the flight data is not coincident with the time at which the image was taken. Each star camera solution must be used at the time the image was taken. Therefore star camera fields must be produced which are time-synchronized with the rest of the flight data. These synchronized fields are then used in the construction of the post-flight pointing solution.

Each camera has an electrical trigger pulse that can be used to pinpoint the exposure time, and produce a set of time-synchronized star camera fields. The flight data also contains the exposure time of each image, as reported by the star camera in microseconds. In the rare cases where a trigger pulse did not register, this information can be used to determine the time at which the image was taken. However this requires the cpu clocks of the flight computer and star cameras to be synchronized. Figure 5.4 shows star camera fields that are in the flight data, and the same fields after time-synchronization.

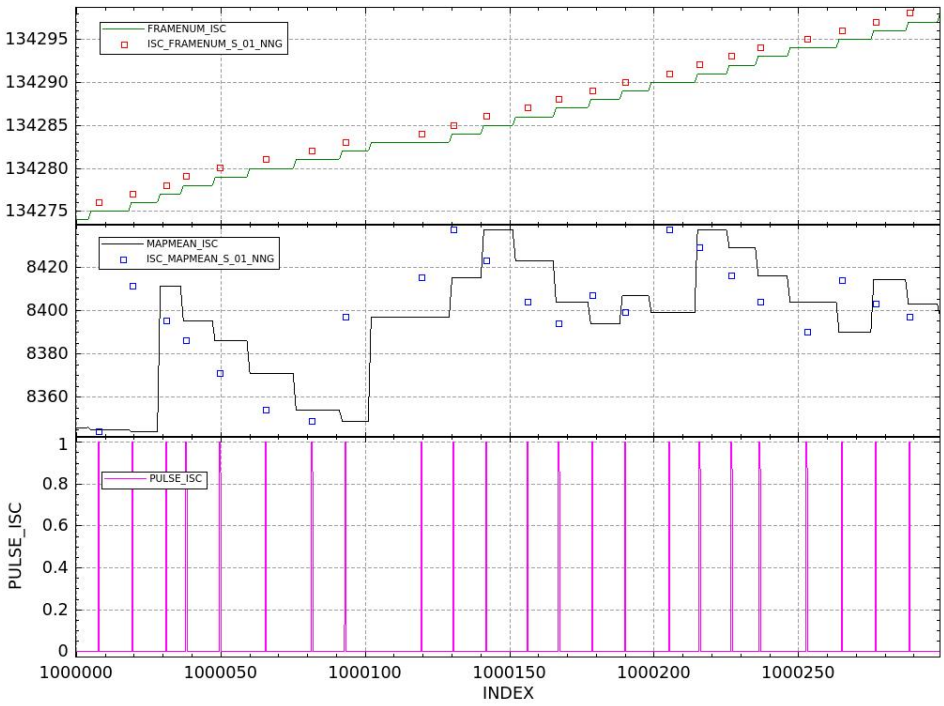

Figure 5.4: Top and middle: in-flight fields (lines) and post-flight synchronized fields (points), Bottom: Trigger pulse. PULSE=1 when the image is taken.

5.2.3 Post-flight Star Camera Solutions

The pyramid algorithm (Section 2.8.2) is applied to calculate pointing solutions based on reported blob positions for the entire flight. Whereas in flight, a minimum of 3 stars must be matched before a pointing solution is accepted, post-flight this minimum can be reduced to 1. At first, the in-flight RA and DEC are input as a coarse pointing solution guess, but once a gyro-integrated pointing solution is calculated, this is used instead. As the accuracy of the gyro-integrated pointing solution increases, more 1- and 2-star solutions can be found. When observing Saturn, instead of generating a pyramid-based solution of stars in the image, the pointing solution for each frame was simply a "1-star" solution, based on the known coordinates of Saturn.

As described in Section 4.3.2, the BLASTPol boresight star camera images included several spurious bright pixels which were identified as blobs, and therefore included in the flight data blob list. During post-flight analysis, these spurious blobs led to pyramid calculating solutions for these frames that were clearly wrong. This effect was reduced by identifying the recurring bad blobs, and removing them from the blob list (see Figures 5.5 and 5.6). The solutions generated by pyramid improved when these bad blobs were removed.

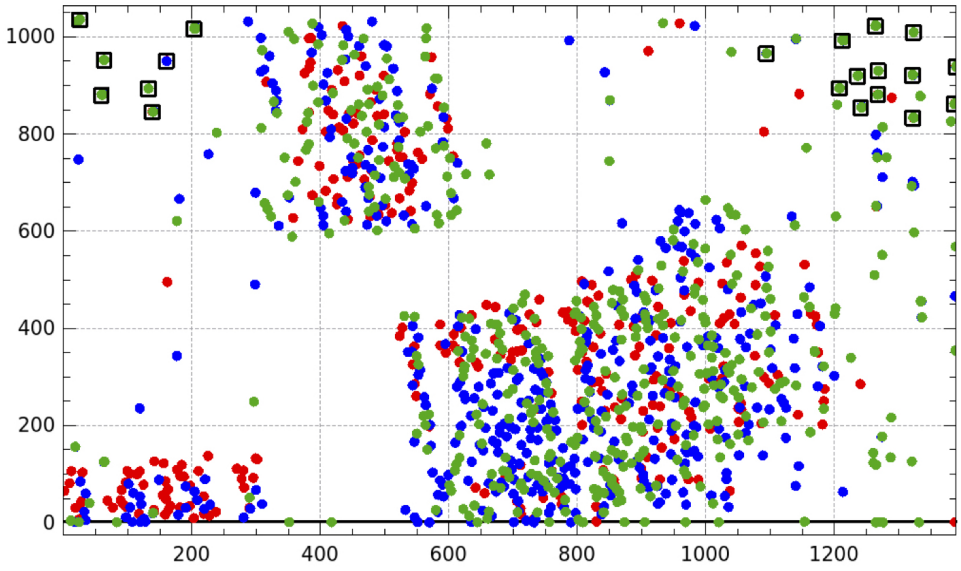

Figure 5.5: Blob positions in x, y CCD coordinates. Red, blue and green are the 1st, 2nd and 3rd brightest blob, respectively. Blobs with boxes around them, and blobs below $y = 5$ were removed from the blob list.

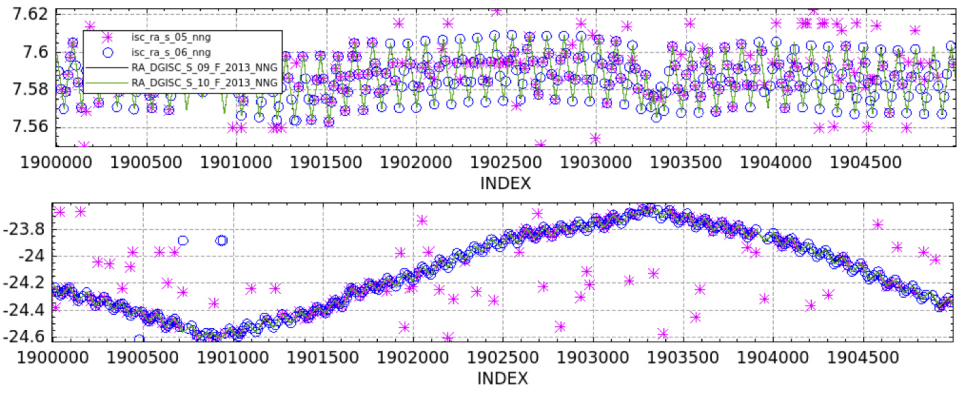

Figure 5.6: Two versions of the post-flight pointing solution. Top: RA(hrs), Bottom: DEC(deg). Points are pyramid solutions (versions 05 and 06), and lines are Kalman-integrated solutions. In version 05, the bad blobs were included, and in the version 06 they were removed.

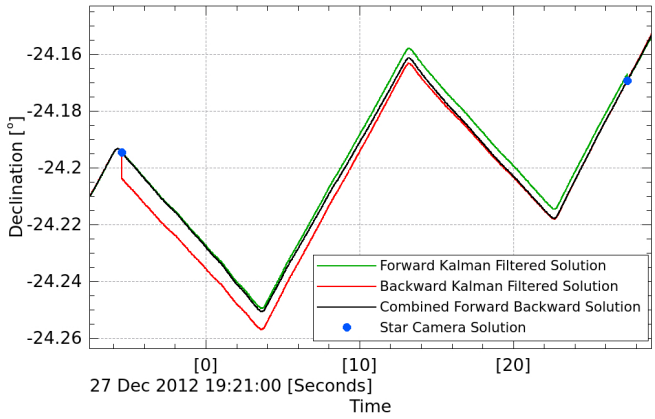

Figure 5.7: Kalman Filtered pointing solution.

5.2.4 Kalman Integration

The rotated gyroscope timestreams (Section 5.2.1) are used to interpolate between each star camera solution obtained as described in Section 5.2.3. The pointing solution algorithm employs an Extended Kalman Filter method used by the Wilkinson Microwave Anisotropy Probe (WMAP) as described in Harman (2005). The attitude of the telescope at each point in time is represented as a quaternion, which can be written as a set of 4 numbers $q = (w, x, y, z)$ where w is real and x, y, z are imaginary. In this representation, $w = cos(\theta/2)$, where θ is the amount of rotation about the axis defined by (x, y, z). A quaternion can also be written as a scalar and a vector $q = (s, \mathbf{v})$, with $s = w$ and $\mathbf{v} = (x, y, z)$. At each time step, a quaternion multiplication is performed to rotate the attitude of the telescope by the rotation quaternion given by the gyroscopes

$$q_{n+1} = \begin{pmatrix} 1 \\ \frac{1}{2} w_n \Delta T \end{pmatrix} q_n + \begin{pmatrix} 0 \\ \frac{1}{2} \mathbf{b}_n \Delta T \end{pmatrix} q_n + \begin{pmatrix} 0 \\ \frac{1}{2} \mathbf{p}_{w_n} \Delta T \end{pmatrix} q_n \tag{5.1}$$

where $\mathbf{b}_n$ is the gyro bias and $\mathbf{p}_{w_n}$ is the noise in the gyro rate measurement.

Starting with the initial star camera measurement, the attitude is propagated forward with the gyroscopes by quaternion multiplication at each discrete time step. As this propagated solution moves forward, its error grows according to the gyroscope angle random walk error. When a new star camera solution is available, the Kalman Filter calculates a weighted combination of the predicted attitude, given by the gyros, with the measured attitude, given by the star camera solution. The predicted attitude is weighted by the error in the propagated solution, and the measured attitude is weighted by the error in the star camera solution.

The filter is run both forward and backward in time, and the forward and backward solutions are then combined, weighted by their respective errors. This smooths the solution over the periods in which there are no star camera solutions (Figure 5.7).

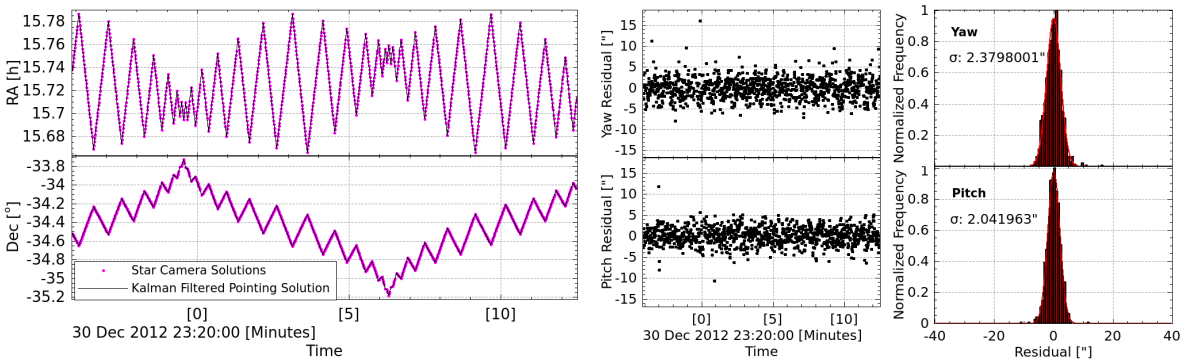

Figure 5.8: Left: Points are star camera solutions and solid line is the final Kalman filtered pointing solution. Right: Residuals in the pointing solution in pitch and yaw.

5.2.5 Results

Figure 5.8 shows the discrete star camera solutions and the Kalman filtered pointing solution in RA and Dec for ∼15 minutes of the BLASTPol 2012 flight. The residuals in yaw (=RA×cos(Dec)) and pitch (=Dec) are also shown. These are the difference between the Kalman filtered solution and the star camera solutions. The sum in quadrature of the rms residual in pitch and yaw is the error in the pointing solution, which is ∼3″.

5.3 SPIDER star camera

The same pointing solution pipeline used on the ISC was used on the SPIDER star camera, with slight modifications. The SPIDER star camera was not mounted along the boresight of the telescope, but was fixed to the outer frame of the gondola. Therefore before using the in-flight pointing solution as an initial guess to input into the post-flight star camera solution finder, a constant azimuth offset of 70.28° was added, and the elevation was fixed at 23.7°. Figure 5.9 shows the number of stars identified in each star camera solution during the entire 2012 flight. During the second half of the flight there is an increase in the frequency of solutions and the number of stars identified, probably due to the threshhold for blob identification being lowered, leading to more stars being included in the blob list. Most of the solutions are 3-star solutions, and the lack of 1- and 2-star solutions is due to a minimum of 3 being imposed. A large fraction of solutions identified 4 or 5 stars, with a maximum of 13 stars.

Since the gyroscope data had been rotated into the ISC reference frame, the post-flight SPIDER star camera solutions were then also rotated into the ISC reference frame, and the Kalman integration was performed on them as was done on the ISC. The two solutions were compared during a time when the ISC was still working, as shown in Figure 5.10.

Figure 5.11 shows that the difference between the SPIDER star camera pointing solution

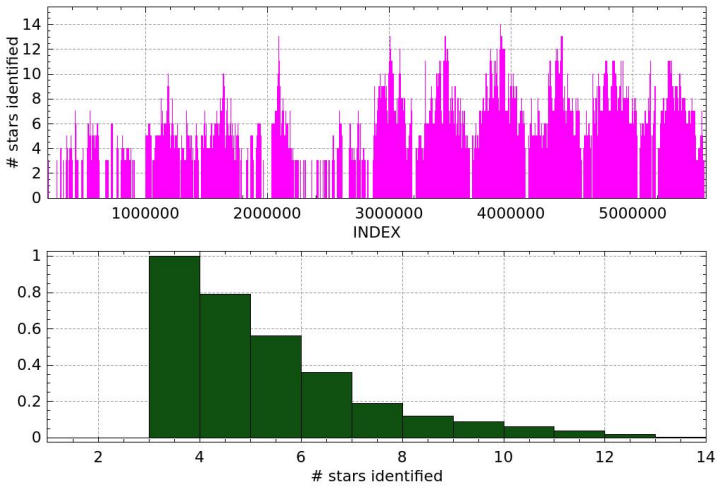

Figure 5.9: Top: Number of stars identified in each star camera solution over the entire 2012 flight. Bottom: Normalized histogram of the number of stars identified.

and the ISC pointing solution is $\sim 40''$ rms. Figure 5.12 shows a preliminary bowtie map of Vela C made using the roll star camera pointing solution. Future improvements to the pointing solution can be made by constraining the pointing offset using Herschel maps.

5.4 Beam Analysis

5.4.1 Beam shape

Characterization of the beam shape was done using observations of Mickey and Saturn. The analysis was primarily done using Mickey because the observation of Saturn did not cover all of the bolometers in the detector arrays, due to difficulties in pointing at Saturn in-flight. The beam was found to have an elongated shape with multiple lobes at 250μm and 350μm. Models of the beam shape were made by fitting a sum of multiple elliptical gaussians to the lobes. Although the beam shape was observed to vary across the array, and vary in time during the flight, the final beam model combined data from all bolometers in an array, and over all Mickey observations during the flight. The beam models for the three wavelengths are shown in Figure 5.13. Also shown are the flux-weighted centroids, calculated from the flux enclosed within the 27.5% contour level.

The beam shape was observed to vary across the detector array, with different detectors having more power in different lobes of a double-lobed beam. Also observed was variation in time of beam properties such as the ratio of power in the each of the lobes, the rotation angle of the beam. the separation of the lobes, and the FWHMs of the beams.

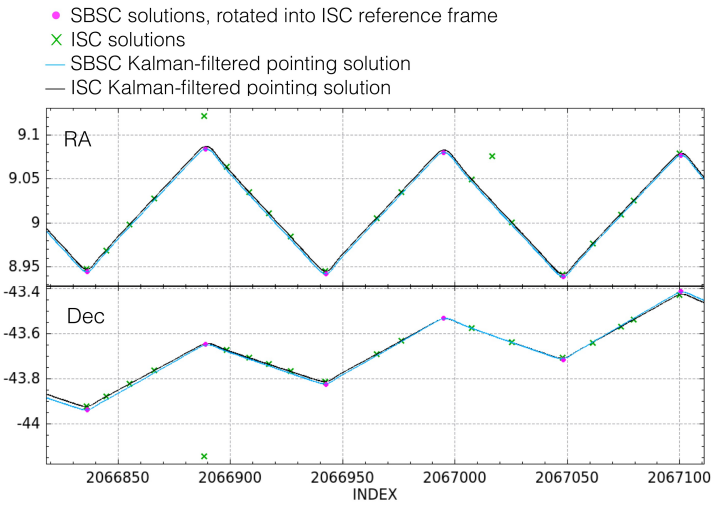

Figure 5.10: Comparison of roll star camera integrated solution to that of the ISC.

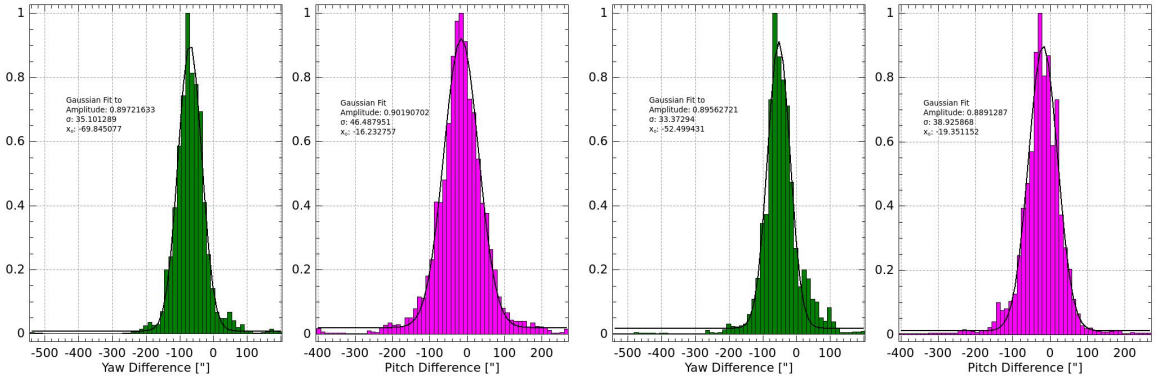

Figure 5.11: Difference between the integrated pointing solutions of the ISC and roll star camera. ΔYaw=ΔRA*Cos(DEC) and ΔPitch=ΔDEC. Left two panels and right two panels are from different frame ranges.

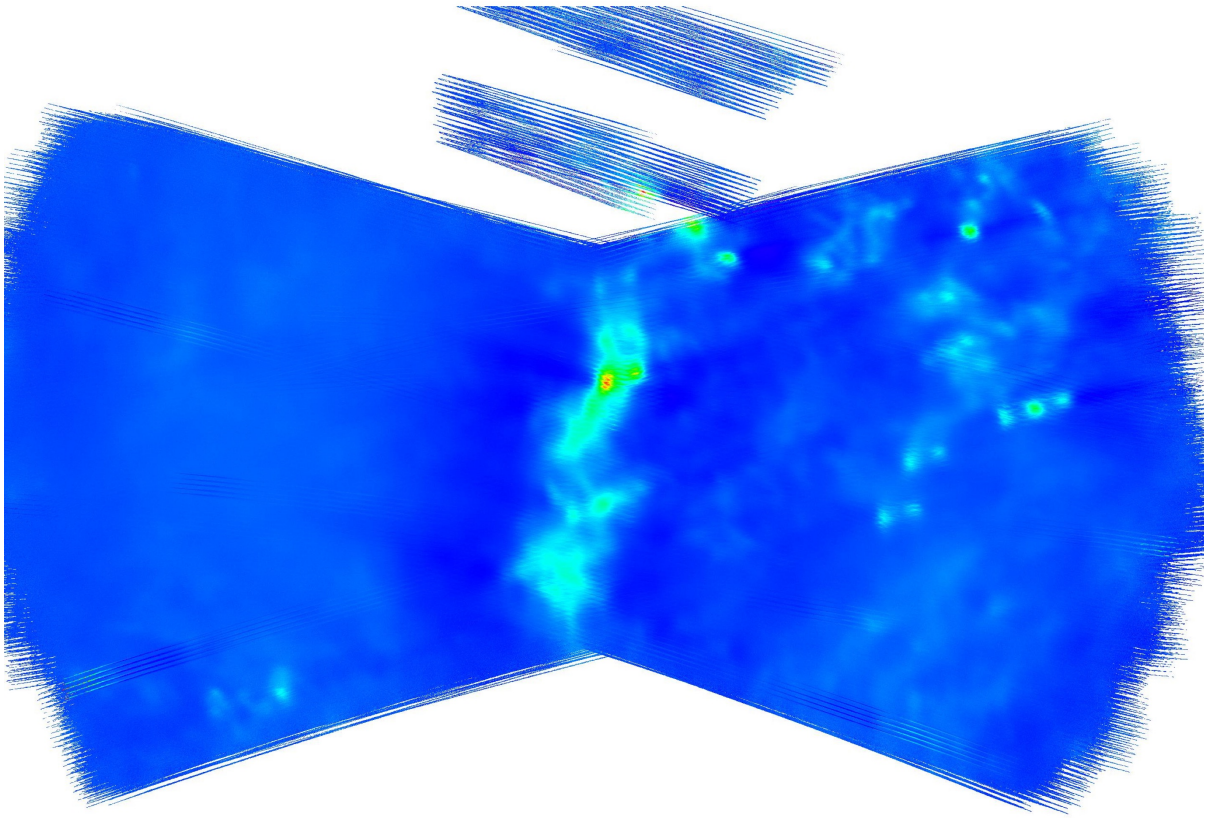

Figure 5.12: Preliminary bowtie map of Vela C, made using the roll star camera pointing solution.

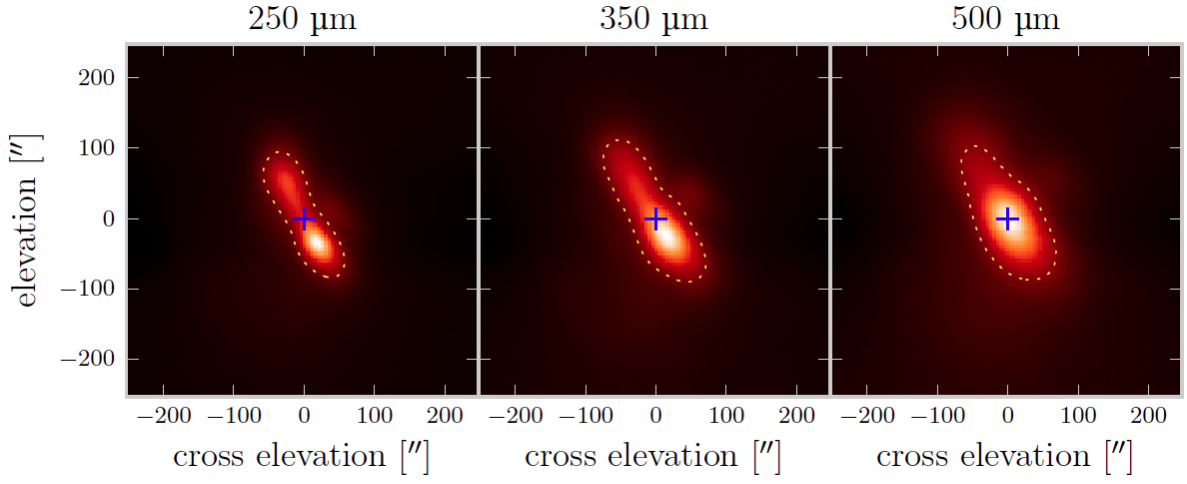

Figure 5.13: Templates of BLASTPol beams at the three wavelengths. Blue crosses are the flux-weighted centroids, and yellow dashed lines indicate the 27.5% contour level.

5.5 Pointing Offsets

In order to make maps from the bolometer timestreams, the star camera pointing solution must first be rotated into each bolometer's reference frame. There are two types of pointing offsets that must be determined and used in the mapmaking process:

1. the offsets between the star camera and each of the three bolometer arrays
2. the offsets between the central bolometer of each array and the other bolometers in that array

The first type of offsets slowly vary in time over the course of the flight, due to deflections in the gondola inner frame causing changes in the alignment between the star cameras and the telescope boresight. These offsets are calculated using 250μm observations of the bright compact source IRAS 08470-4243 ("Mickey"), as described in detail in Shariff (2015). For the analysis of the 2012 flight, a separate offset was calculated for each scan of a science target (where a scan is as described in Section 2.7).

The offsets were calculated by comparing BLASTPol maps at 350μm with simulated Herschel SPIRE maps. These simulated maps were obtained using *simsky*, a software package developed during BLAST 2006 analysis, and described in detail in Fissel (2013). For each science target, a $0.33^\circ \times 0.5^\circ$ region was extracted from the BLASTPol map and the same region was extracted from the simulated Herschel maps. After re-normalizing the maps, a brute-force search was then done to find pointing offset and gain that minimized the residual between the two maps. The search was done with a resolution of $2''$, over a range in pointing offsets of -60 to +60 pixels, which corresponds to -10 to +10 arcminutes for maps with $5''$ pixels. The gain was allowed to vary from 0.7 to 1.0, with a resolution of 0.1.

Figure 5.14 shows the residuals obtained when the best fit gain and pointing offset is applied and the BLASTPol maps are subtracted from the simulated Herschel maps. The offsets were calculated separately for each of the 3 wavelengths. Figure 5.15 shows the results in all three bands for every scan in the first half of the flight, when there was still a working boresight star camera. The difference in results between the three bands is $< 10''$ during Vela ("Axehead") scans. This is smaller than the variation in beam shape described in the previous section, which was up to $20''$ at 250μm. For the purpose of mapmaking the 350μm offsets were used.

Offsets calculated using observations of Mickey also appear to follow a similar trend (Figure 5.16). The results show a roughly diurnal drift in the offsets of $\sim 20''$, although this drift was not found to correlate well with the temperature of the primary mirror, the elevation of the inner frame, or parallactic angle. Each offset was assigned to the frame number in the

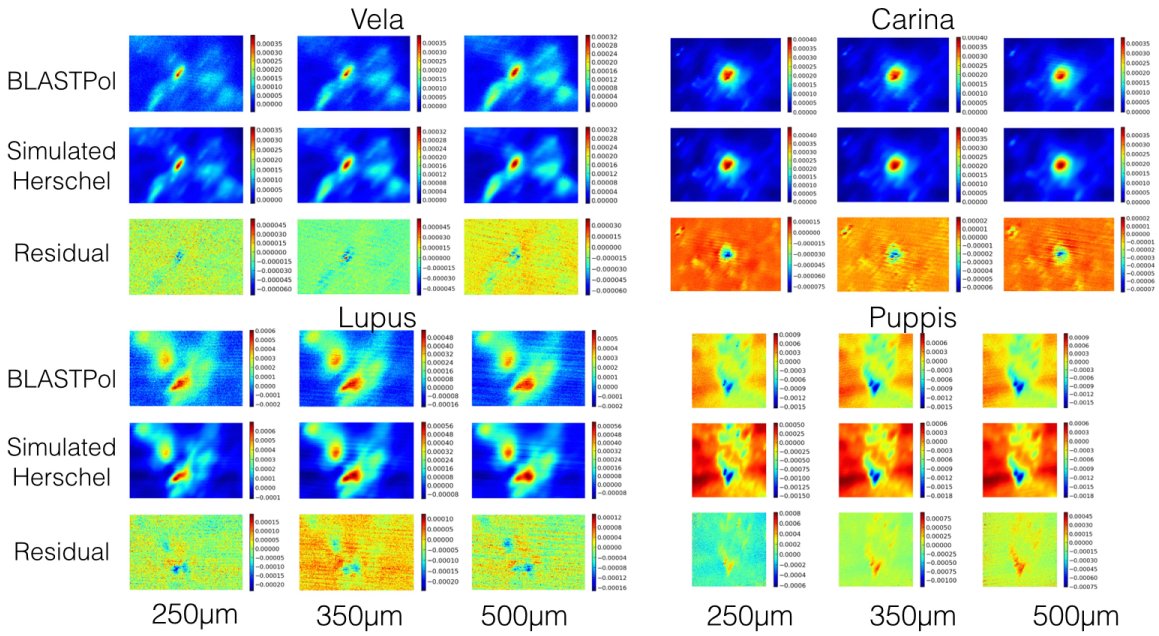

Figure 5.14: Method used for estimating pointing offsets. Results for a single scan of each science target are shown. The top panels show the BLASTPol map, middle panels show simulated Herschel maps of the same region, and bottom panels are the residuals for the best fit pointing offset.

middle of the scan used to measure it, and applied over a continuous frame range until the next offset measurement.

The second type of offsets are determined using single bolometer maps of Mickey, and these offsets do not vary during the flight. Using the time varying pointing offsets described above, single-bolometer maps of Mickey were made. The difference between the flux-weighted centroid of the beam and the known location of Mickey was used as that bolometer's pointing offset.

5.6 Further Tests

5.6.1 Flatfielding

In order to produce a uniform response across the arrays, flatfielding coefficients are applied. Flatfielding is done by comparing single-bolometer maps of Mickey and adjusting the gains until all bolometers see the same total power in each map.

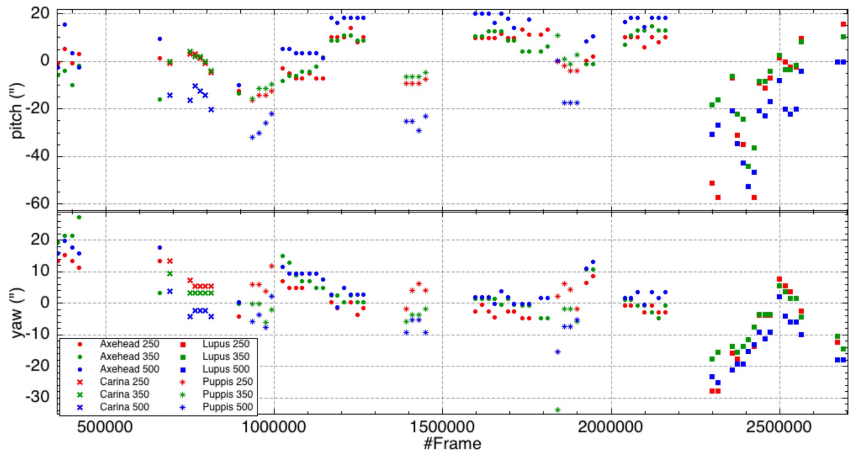

Figure 5.15: Pointing offsets (in arcseconds) obtained over the first half of the BLASTPol 2012 flight, for every scan and in all 3 wavelengths.

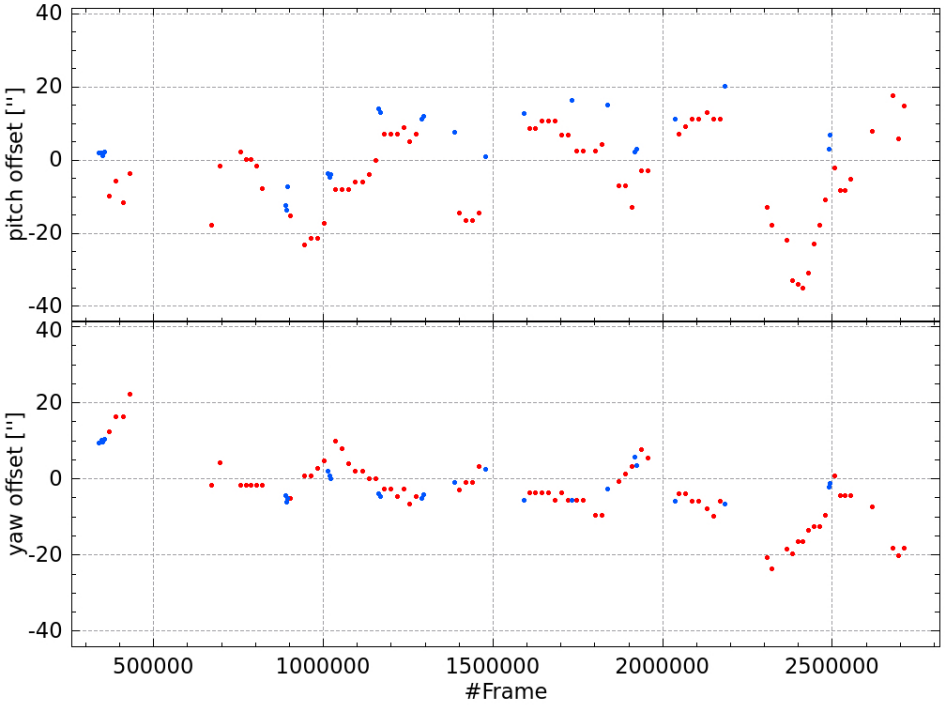

Figure 5.16: Pointing offsets (in arcseconds) used for mapmaking. Red points were obtained using comparison with simulated Herschel maps, blue points using the pointing calibrator 'Mickey'.

5.6.2 Instrumental Polarization

IP is characterized after the flight using bright sources in Vela C, observed with high signal to noise at different parallactic angles. While source polarization rotates relative to the telescope reference frame, instrumental polarization stays fixed. The array-averaged IP values were found to be similar to those measured pre-flight ($< 1\%$ in all three bands). The IP values measured are then subtracted out when making maps. The error in the measured IP values was taken to be $\sim 0.5\%$, which is the level at which simulations were able to reproduce the values. After subtracting the IP, comparison with SPARO observations of G331.5-0.1 (Li et al. 2006) found an agreement to 0.5% in the peak percentage polarization of the source.

5.7 Flux Calibration

Absolute flux calibration factors were obtained by comparing BLASTPol maps of Vela-C to Herschel SPIRE maps of the same region. The Herschel maps were convolved with a BLASTPol beam template, averaged over parallactic angles, and the BLASTPol maps were convolved with the Herschel beam. A $0.8^\circ \times 0.8^\circ$ section of the Vela-C maps was extracted from the BLASTPol maps, and the same square region was extracted from the Herschel maps (Figure 5.17). The power spectrum $(P(k))$ of the BLASTPol and Herschel maps was then computed. The flux calibration factor was then taken to be $(P(k)_{Herschel}/P(k)_{BLASTPol})$ at k=0.1 arcminutes^{-1} (Figure 5.18). The results are presented in Table 5.1.

Wavelength	250μm	350μm	500μm
Calibration factor	4603316.014	992501.898	278888.833

Table 5.1: Absolute flux calibration factors for BLASTPol 2012, based on Herschel maps of Vela-C.

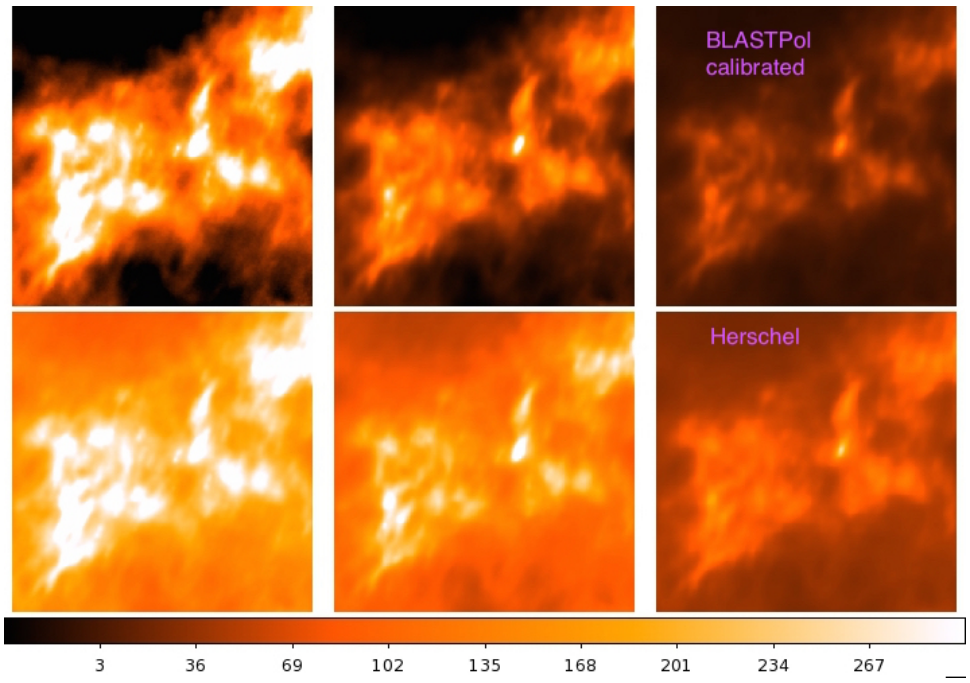

Figure 5.17: Top: BLASTPol calibrated maps. Bottom: Herschel maps convolved with BLASTPol beam. Left: 250μm, Centre: 350μm, Right: 500μm. Color bar units in MJy/Sr.

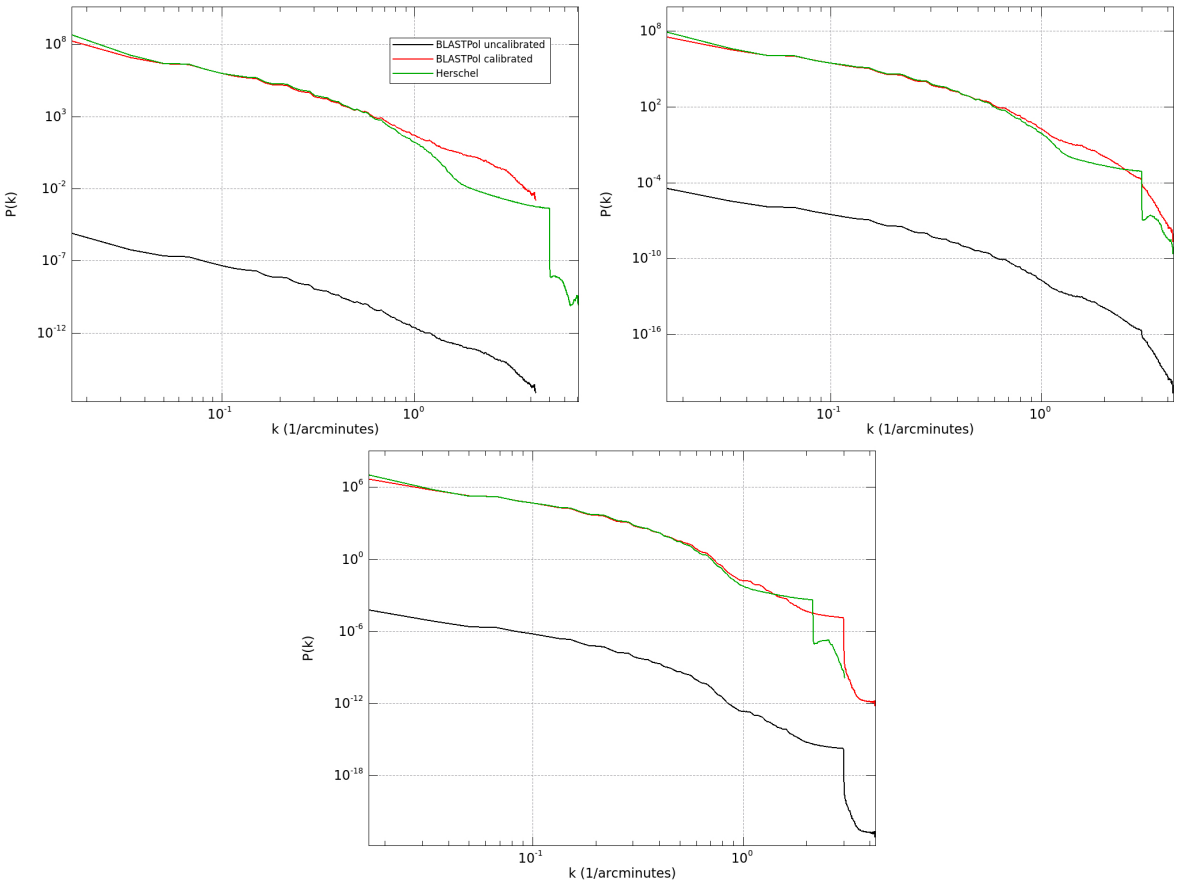

Figure 5.18: Power spectra of Herschel maps (green), BLASTPol uncalibrated maps (black), and BLASTPol calibrated maps (red). Top left: 250μm, Top right: 350μm, Bottom: 500μm.

Chapter 6

Science Results

Figure 6.1: Everybody's got a right to be a sucker once.

6.1 Vela C Maps

Maps are made from the time-ordered data of bolometer signals, $\mathbf{d}$, and the pointing matrix $\mathbf{P}$. These are related to the sky signal, $\mathbf{s}$, by the following equation:

$$\mathbf{d} = \mathbf{P}\mathbf{s} + \mathbf{n} \tag{6.1}$$

where $\mathbf{n}$ is the noise in the time-ordered data. For the 2012 data analysis BLASTPol used the software package TOAST (Kisner 2015) to solve Equation 6.1 for $\mathbf{s}$. TOAST produces maps of I, Q and U, as well as their associated covariances. The I, Q and U maps are shown in Figures 6.2, 6.3 and 6.4, with $10''$ pixel scale.

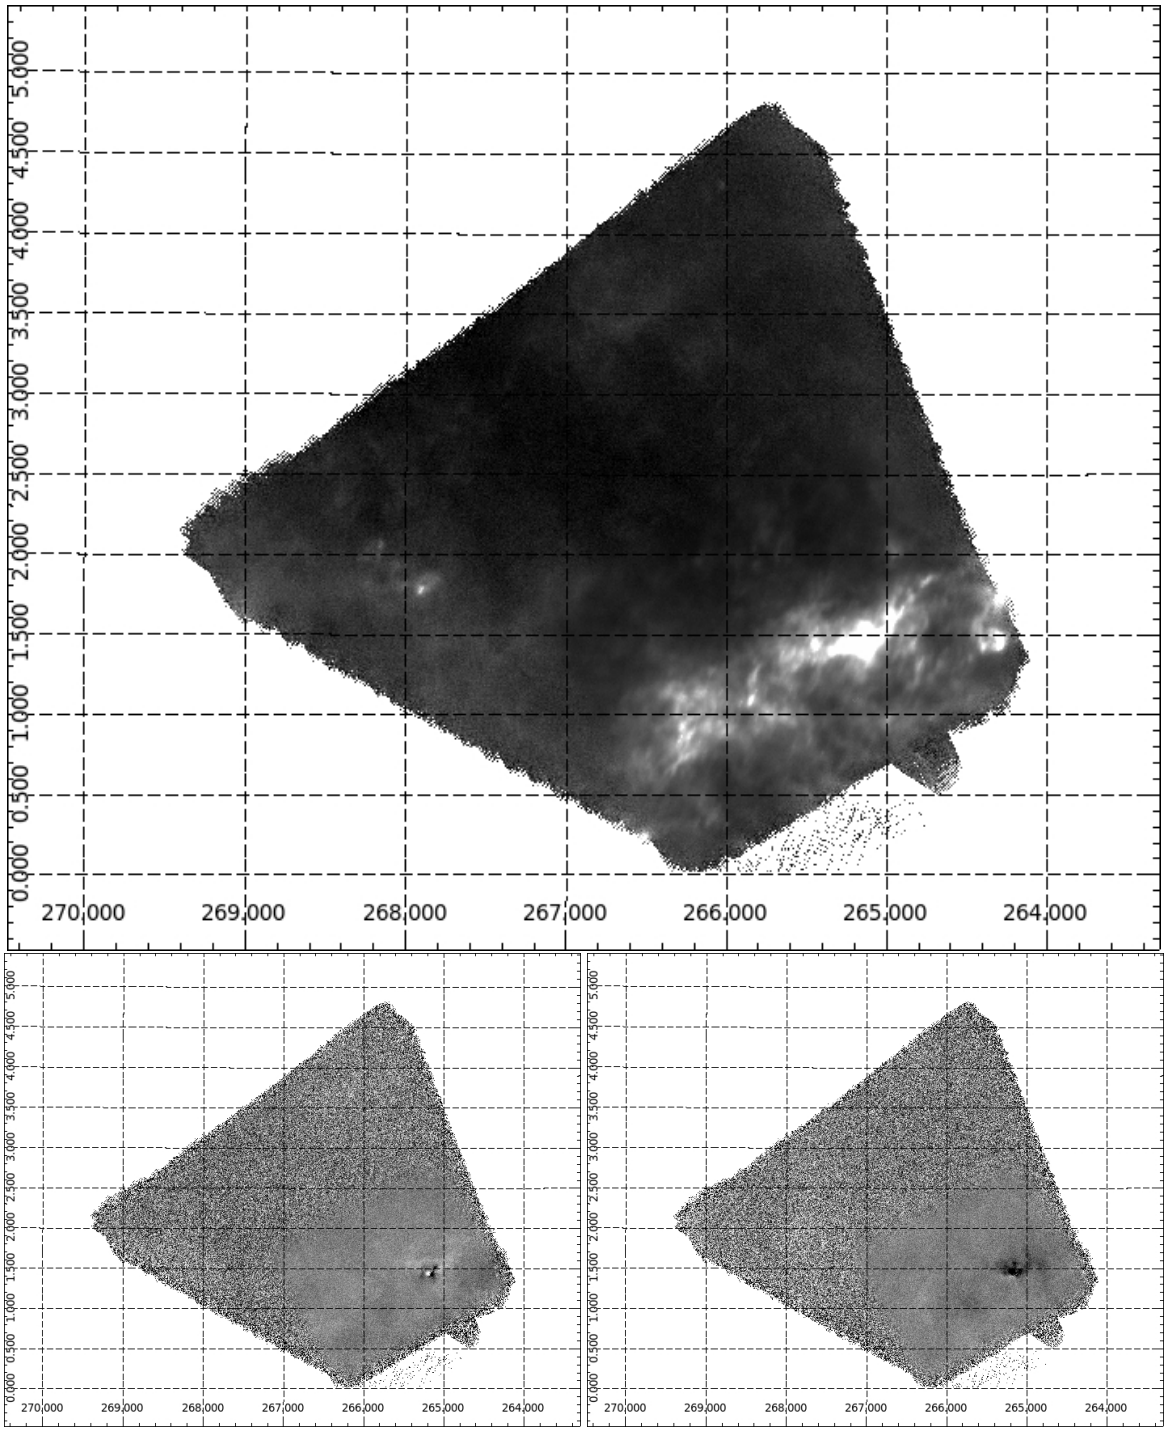

Figure 6.2: Vela C maps at 250μm in galactic coordinates. Top: I; Bottom left: Q; Bottom right: U.

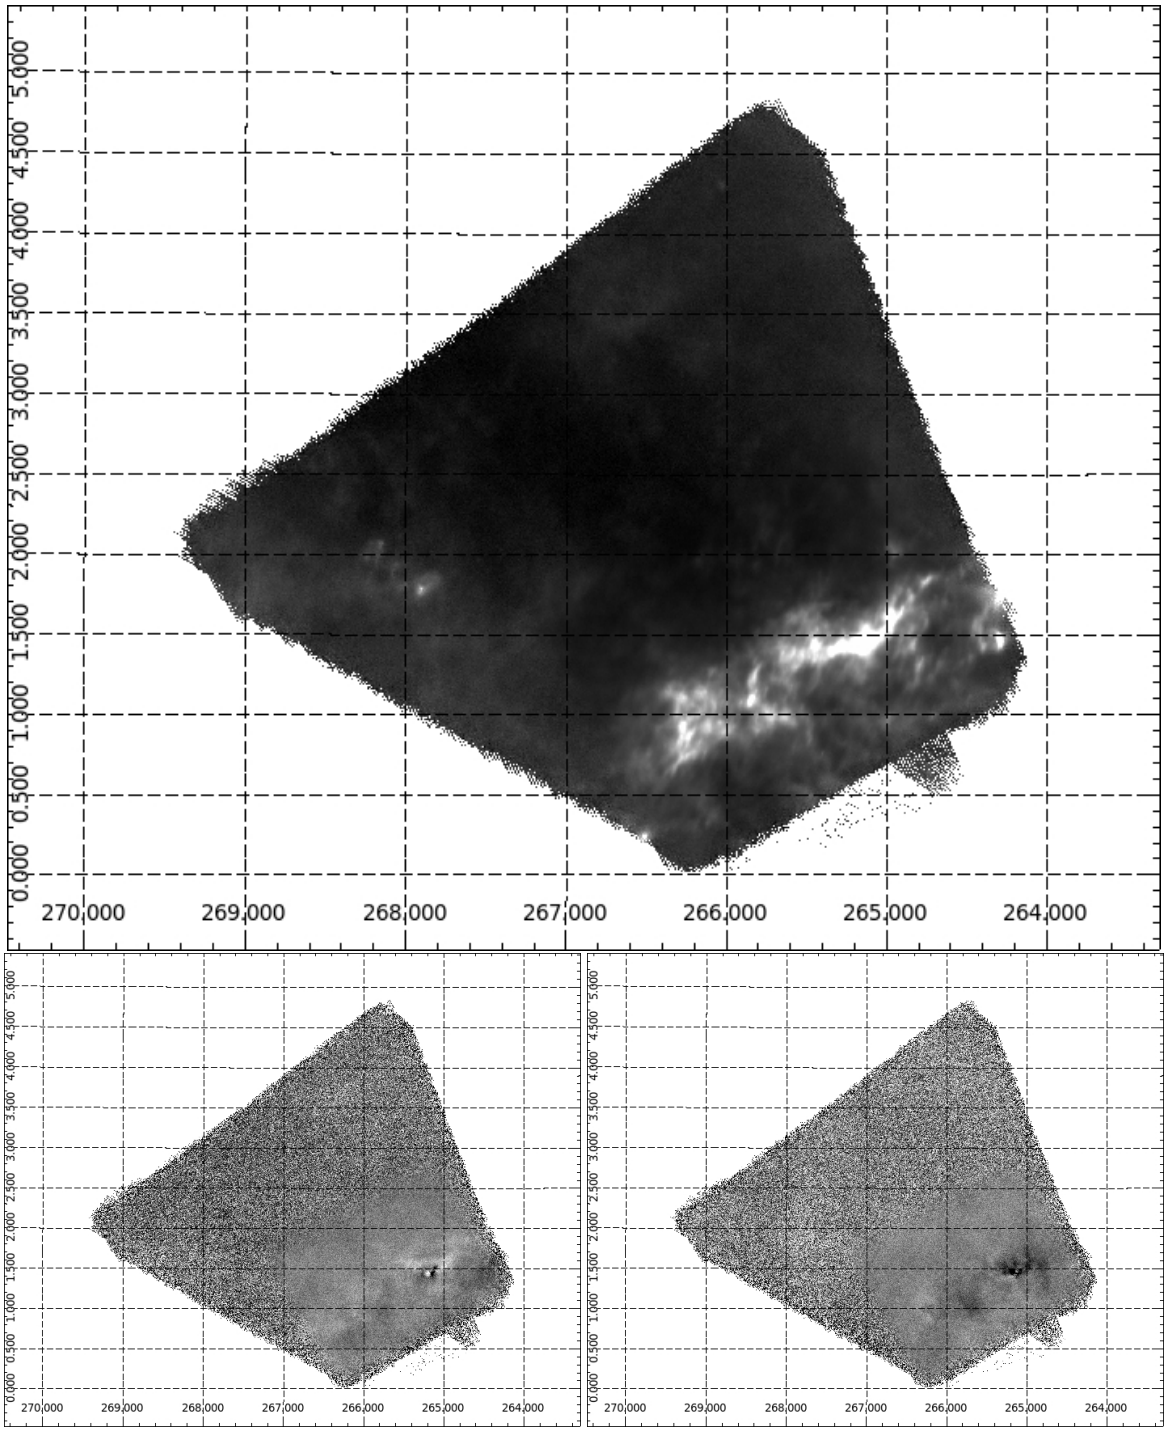

Figure 6.3: Vela C maps at 350μm in galactic coordinates. Top: I; Bottom left: Q; Bottom right: U.

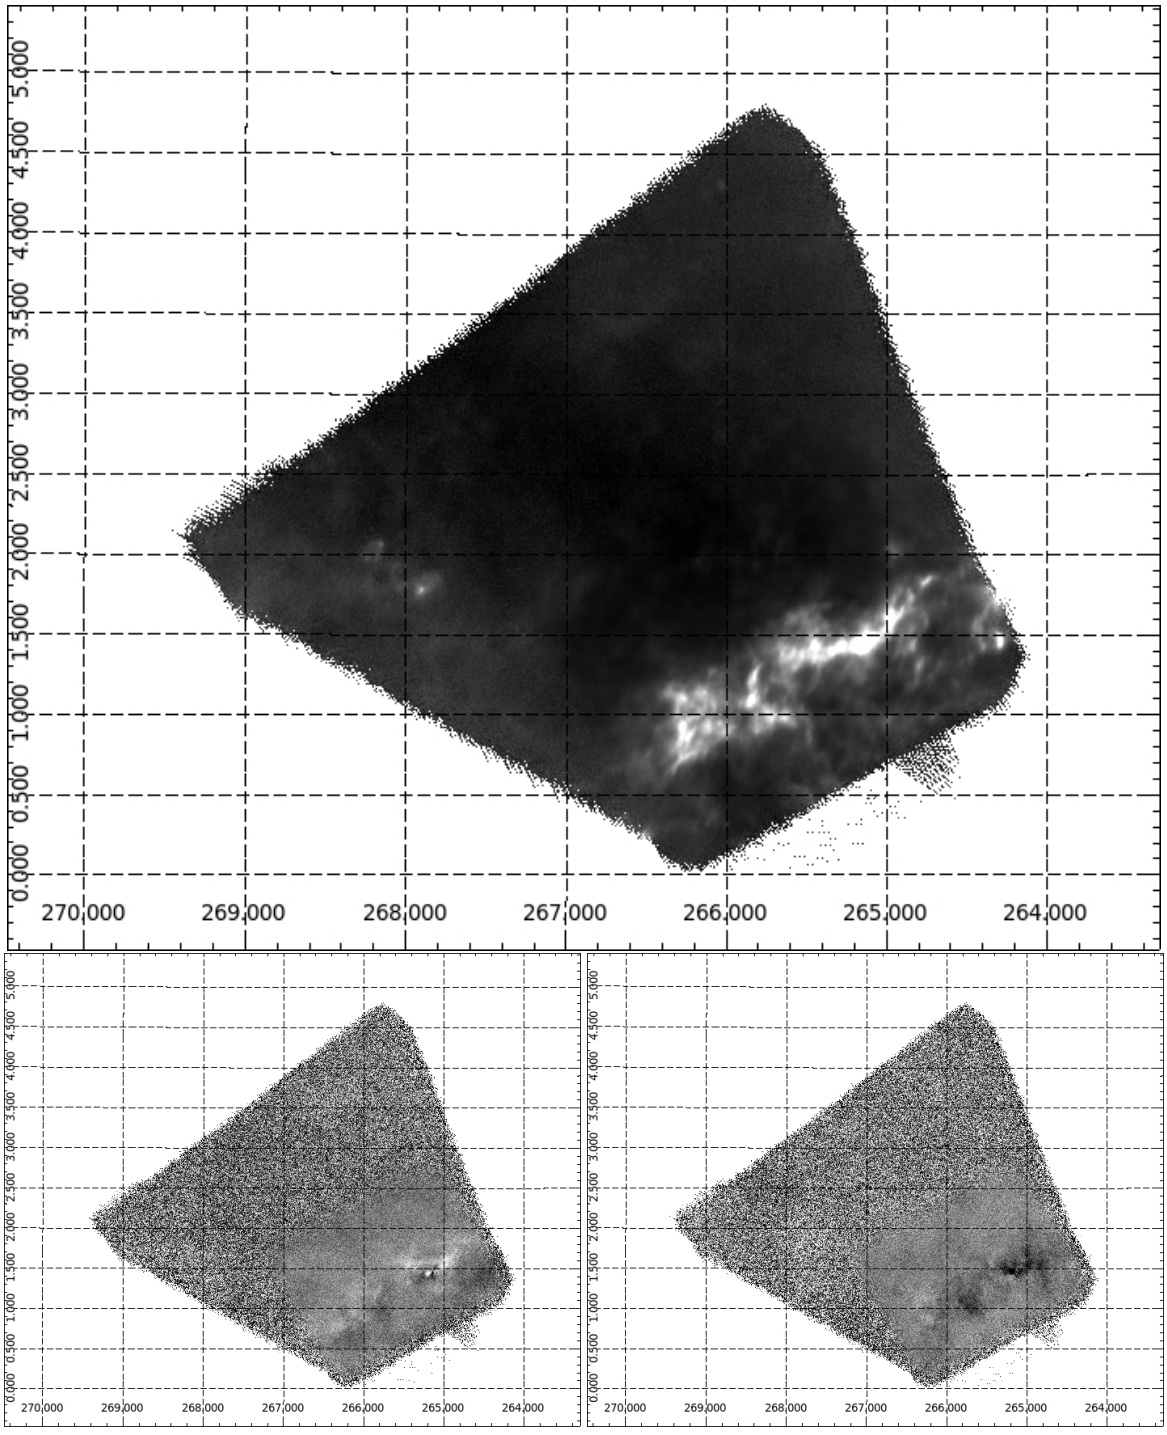

Figure 6.4: Vela C maps at 500μm in galactic coordinates. Top: I; Bottom left: Q; Bottom right: U.

6.1.1 Jackknife Tests

Systematic errors are characterized by splitting the data into independent halves, making maps from each half, and examining the agreement between the two maps. Four ways of splitting the data were used:

- left half of the detector array versus right half
- top half of the array versus bottom half
- first half of the flight versus second half
- one set of data chunks versus the other

where a "chunk" is one up and down raster scan at one half-wave plate position. The flight was divided sequentially into alternating chunks (A,B,A,B...) and maps made from all A chunks were compared to maps made from all B chunks. In each case, data from the two halves were differenced in three quantities: polarization percentage (p), total polarized flux (P), and polarization angle (ϕ). Data was rejected where the difference in p was greater than $1/3$ of the measured value of p, and similarly for P. Data was also rejected where the difference in ϕ was greater than 20°.

6.1.2 Foreground Subtraction

Two methods were used to subtract the galactic dust that is not part of the Vela C molecular cloud. The first used a region near the cloud shown in Figure 6.5 called the "near reference region" to represent the dust off the cloud. The mean value of I, Q and U in this region was then subtracted from the map. The second method fit a plane to two regions above and below the cloud (also shown in Figure 6.5), and subtracted from each pixel in the map the value of the planar fit (for I, Q or U) in that pixel. The only data points included in the analysis were those where the polarization fractions obtained using both methods agreed to less than $1/3$ of their average value, and the angles obtained agreed to less than 10°. A region defining the cloud is also shown in Figure 6.5, which is limited to be inside the regions used in the planar fit. Only data inside the cloud region was used. The results shown in the next section apply the planar fit subtraction method, and using the other subtraction method did not significantly alter the results.

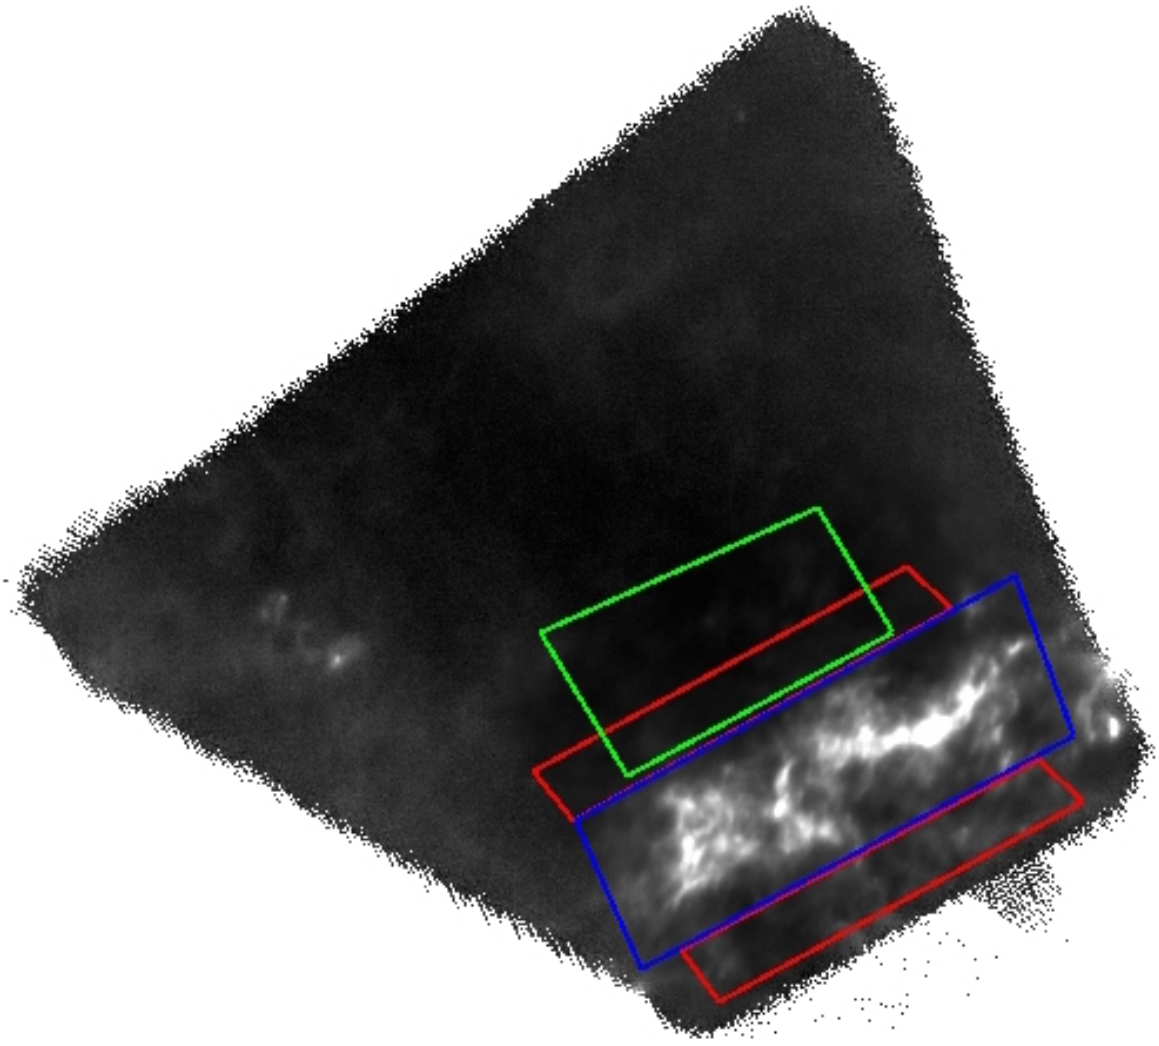

Figure 6.5: Vela C reference regions. Near reference region is green, regions used for the planar fit are red, and blue shows the region defining the cloud.

6.2 Polarization Spectrum in the Vela C Cloud

6.2.1 Data Processing

Maps were smoothed to $3'$ and sampled every $1.5'$. Polarization fraction, p, is calculated from the I, Q, and U maps as $p = \sqrt{Q^2 + U^2}/I$. The polarization angle. ϕ, with respect to the Q, U reference frame, is calculated as $\phi = \frac{1}{2}\tan^{-1}\left(\frac{U}{Q}\right)$. The variance in p is

$$\sigma_p^2 = \left(\frac{q^2}{p^2}\sigma_q^2 + \frac{u^2}{p^2}\sigma_u^2 + 2\frac{qu}{p^2}\mathrm{Cov}(q,u)\right) \tag{6.2}$$

where

$$\sigma_q^2 = q^2\left(\frac{\sigma_I^2}{I^2} + \frac{\sigma_Q^2}{Q^2} - 2\frac{\mathrm{Cov}(I,Q)}{IQ}\right) \tag{6.3}$$

$$\sigma_u^2 = u^2\left(\frac{\sigma_I^2}{I^2} + \frac{\sigma_U^2}{U^2} - 2\frac{\mathrm{Cov}(I,U)}{IU}\right) \tag{6.4}$$

$$\mathrm{Cov}(q,u) = qu\left(\frac{\sigma_I^2}{I^2} + \frac{\mathrm{Cov}(Q,U)}{QU} - \frac{\mathrm{Cov}(I,Q)}{IQ} - \frac{\mathrm{Cov}(I,U)}{IU}\right) \tag{6.5}$$

The polarization measurements are corrected for positive polarization bias using the formula: $p_c \approx (p_m^2 - \sigma_p^2)^{1/2}$, where p_c is the corrected polarization fraction, p_m is the measured polarization fraction, and σ_p^2 is the variance in p_m. Only measurements where $p_m \geq 3\sigma_p$ are used. Ratios of the polarization fraction at two wavelengths are only calculated for sight lines where 3σ polarization measurements exist at both wavelengths, and the measured polarization angles at both wavelengths agree to $<10^\circ$. Requiring the polarization angles to agree reduces the probability that the different wavelengths are measuring different points along the line of sight.

6.2.2 Polarization Ratios

A total of 485 measurements were obtained in the Vela C cloud. Figure 6.6 shows maps of the polarization ratios $p_{250\mu m}/p_{350\mu m}$ and $p_{500\mu m}/p_{350\mu m}$. One obvious feature in the maps is the change in polarization ratio around the massive star-forming region RCW 36 (to the north-east of the centre of the map). Both an increase in $p_{250\mu m}/p_{350\mu m}$ and a decrease in $p_{500\mu m}/p_{350\mu m}$ are seen towards RCW 36. The value of $p_{250\mu m}/p_{350\mu m}$ also increases strongly in one corner of the south part of the cloud. However no strong feature is seen there in the $p_{500\mu m}/p_{350\mu m}$ map.

Overall there appears to be a wider distribution of values in $p_{250\mu m}/p_{350\mu m}$ across the cloud, while $p_{500\mu m}/p_{350\mu m}$ is close to 1.0 across the cloud. Figure 6.7 shows histograms of the two polarization ratios, with median values of $p_{250\mu m}/p_{350\mu m} = 1.11$ and $p_{500\mu m}/p_{350\mu m} = 1.07$.

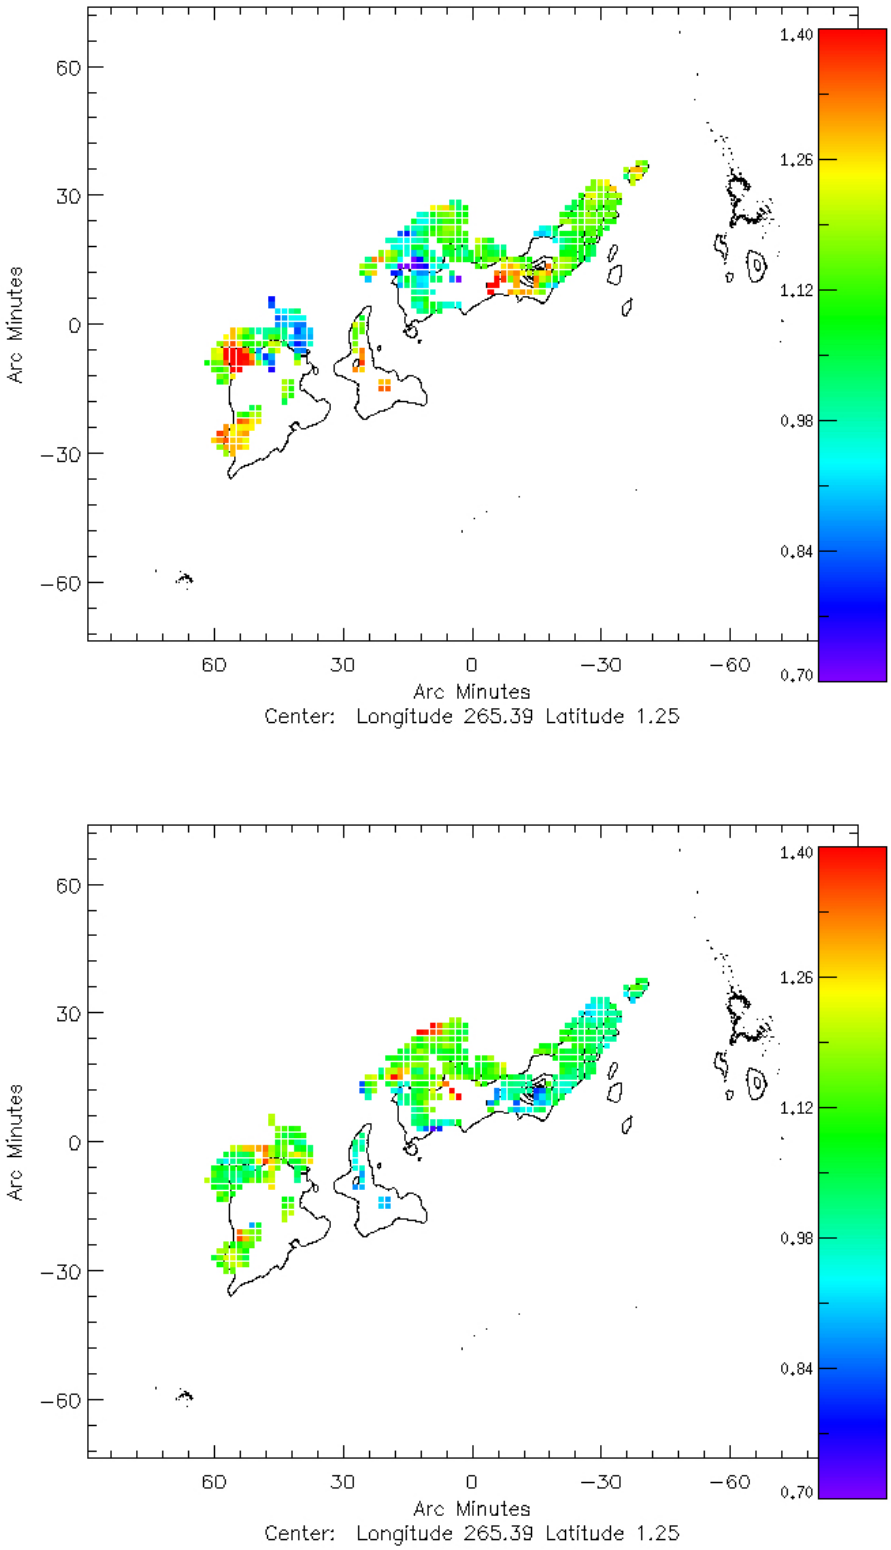

Figure 6.6: Top: $p_{250\mu m}/p_{350\mu m}$; Bottom: $p_{500\mu m}/p_{350\mu m}$.. Contours are 500μm intensity.

Medians are used because the distributions are non-normal, so they are not well characterized by a simple mean and standard deviation. A method similar to that in Vaillancourt & Matthews (2012) is used to describe the uncertainty in the histogram's median and distribution. In order to determine whether the distribution in the histogram is real, or simply a result of the uncertainties in the measurements of p, a reduced-χ^2 statistic is defined:

$$\chi_r^2 = \frac{1}{N-1}\sum_{i=1}^{N}\frac{(x_i - x_m)^2}{\sigma_i^2} \tag{6.6}$$

where x_m is the median value of the samples x_i, σ_i is the measurement uncertainty on the quantity x_i. If the quantity is a polarization ratio $(p1/p2)$, then the measurement uncertainty is:

$$\sigma_{p1/p2} = \frac{p1}{p2}\sqrt{\left(\frac{\sigma_{p1}}{p1}\right)^2 + \left(\frac{\sigma_{p2}}{p2}\right)^2} \tag{6.7}$$

If most data points are consistent with the median polarization ratio within their uncertainties then χ_r^2 should be around 1. For $p_{250\mu m}/p_{350\mu m}$ $\chi_r^2 = 263$, and for $p_{500\mu m}/p_{350\mu m}$ $\chi_r^2 = 605$. Since both values are much greater than 1, the widths in the distributions are taken to be intrinsic and not solely a result of measurement uncertainties.

The median absolute deviation (MAD) is used to quantify the error in the median value of the distribution. The MAD is the median value of the residuals:

$$\text{MAD} \equiv \text{median}(|x - x_m|) \tag{6.8}$$

where x_m is the median value of the measurements x. For $p_{250\mu m}/p_{350\mu m}$ MAD=0.09, and for $p_{500\mu m}/p_{350\mu m}$ MAD=0.07.

Figure 6.8 shows the resulting polarization spectrum measured by BLASTPol, with error bars equal to the MAD. The figure also shows measurements of the polarization spectra of other molecular clouds reported in several previous works (Vaillancourt & Matthews 2012; Zeng et al. 2013). All but one of the previous measurements have indicated a minimum in the spectrum around 350μm. The spectrum measured by Zeng et al. (2013) does not show a minimum, and instead continues decreasing towards longer wavelengths. The BLASTPol data produces a flatter spectrum than previous measurements, other than Zeng et al. (2013), but with a minimum at 350μm.

The existence of a minimum in the spectrum implies the existence of multiple populations of dust grains with different temperatures and spectral indices (Hildebrand et al. 1999; Vaillancourt 2002). These populations experience different alignment efficiencies, leading to multiple polarized components and a wavelength-dependence in the observed polarization percentage.

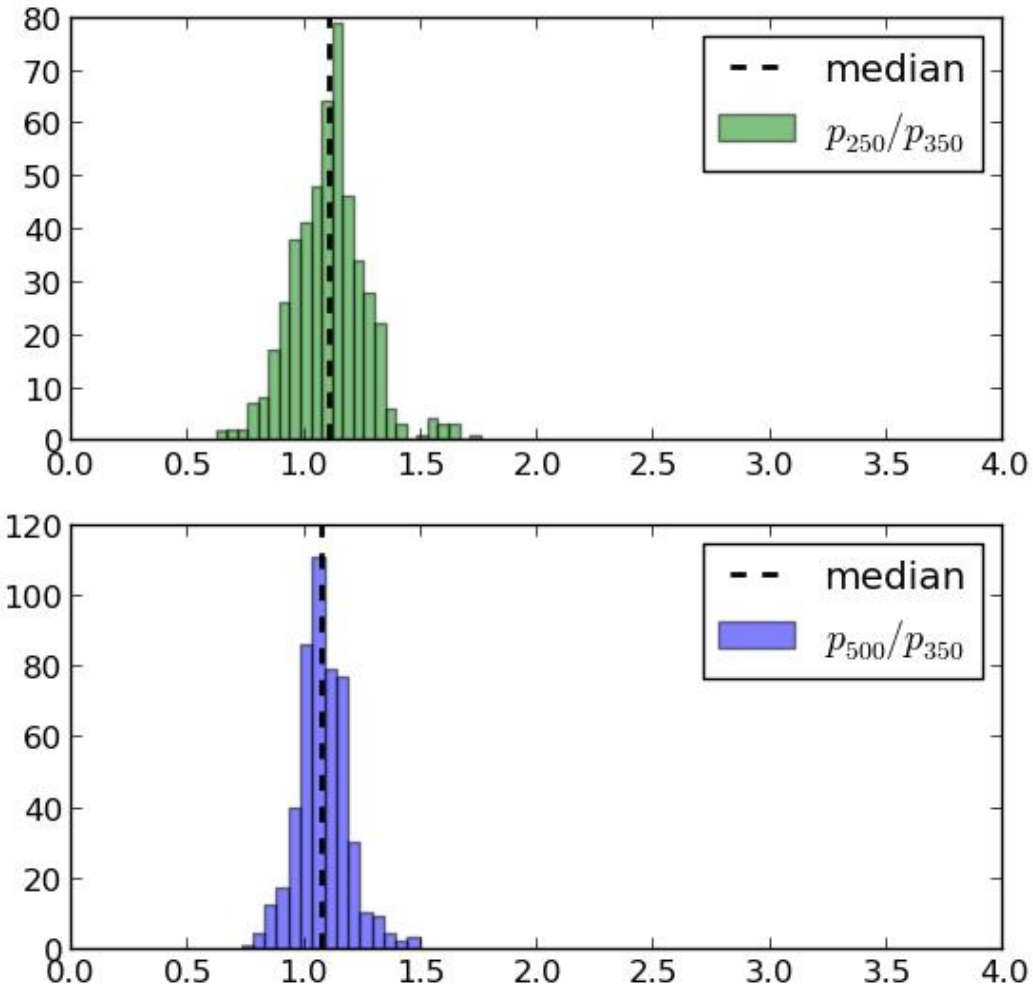

Figure 6.7: Histograms of polarization ratios. Dashed lines indicate medians.

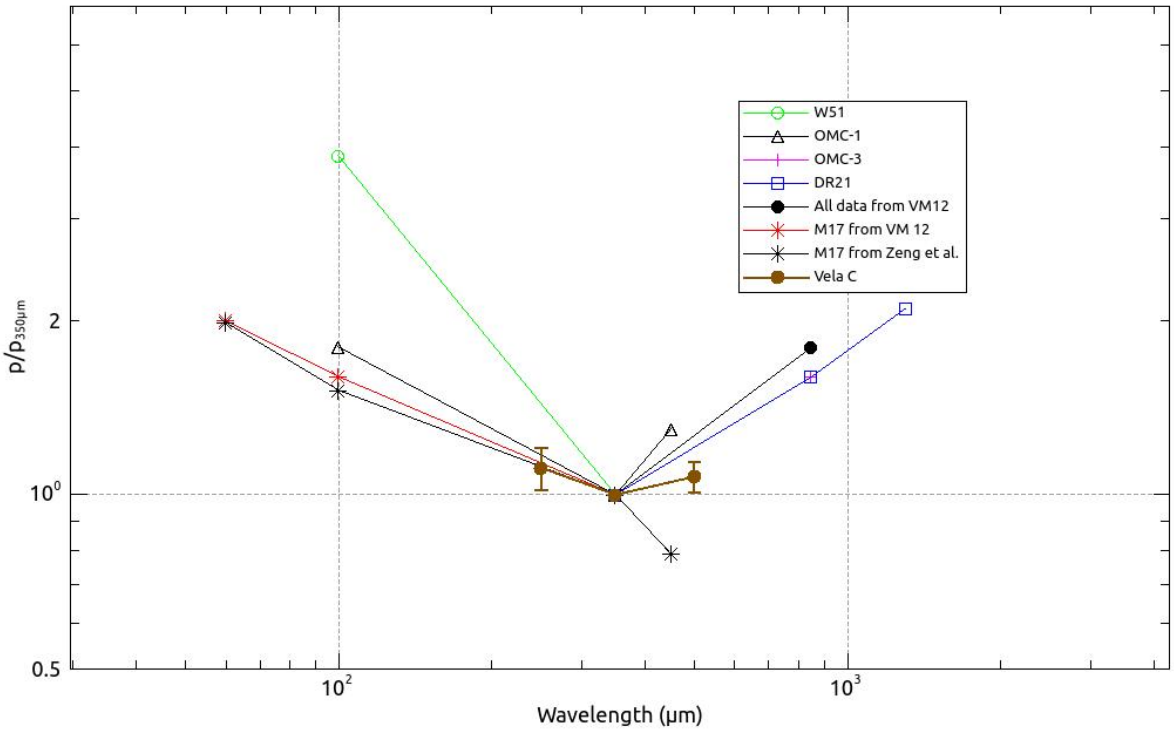

Figure 6.8: Polarization spectrum.

The BLASTPol polarization spectrum cannot be directly compared to any existing models, since there are no models that apply to bright dense molecular clouds like Vela C. Draine & Fraisse (2009) and Bethell et al. (2007) both make predictions of a steadily rising submillimeter polarization spectrum. However the Draine & Fraisse (2009) models apply to diffuse dust, and were intended to be used for high galactic latitude observations of the CMB. The Bethell et al. (2007) model of a molecular cloud assumes that the dust is exposed to a uniform interstellar radiation field, with no radiation coming from inside the cloud itself.

While the median polarization spectrum for Vela C as a whole seems to agree with previous measurements about the location of a minimum at 350μm, the maps in Figure 6.6 suggest coherent structure in the variation of these ratios across the cloud. RAT theory predicts that these variations are due to the dust being exposed to different radiative environments, with the asymmetry of the radiation field being an important factor in the process of grain alignment. Zeng et al. (2013) and Vaillancourt et al. (2008) measured trends in polarization ratios versus intensity ratio (as a proxy for column density) and flux density, respectively. Zeng et al. (2013) suggested that the short wavelength rise in the spectrum occurs near internal radiative sources in the cloud. Section 6.2.3 looks at the correlation between the polarization spectrum and the environmental conditions in Vela C.

6.2.3 Effect of Radiative Environment

In order to study the environmental conditions in the Vela C cloud, maps of number density (in cm^{-2}) and temperature (in K) were made by fitting SEDs to Herschel SPIRE maps using the method described in Hill et al. (2011) (Figure 6.9). The BLASTPol data was then divided into two categories, "protostellar" and "quiescent" (see Figure 6.10). The protostellar category includes data looking towards sightlines near RCW 36 where dust grains are much warmer, being highly irradiated by the ongoing star formation in that region. The quiescent category includes data from the cooler regions of Vela C, where no active star formation is occurring, and the dust is mostly externally heated by the interstellar radiation field. The number of data points in each category is $N_{quiescent} = 421$ and $N_{protostellar} = 59$.

Figure 6.11 shows the polarization ratios $p_{250\mu m}/p_{350\mu m}$ and $p_{500\mu m}/p_{350\mu m}$ versus number density, for data in both categories. Values greater than 1.0 in both polarization ratios lead to a more "V"-shaped polarization spectrum, while values close to 1.0 produce a flatter spectrum.

For both categories, the polarization spectrum shows a slight trend with N, where $p_{250\mu m}/p_{350\mu m}$ increases and $p_{500\mu m}/p_{350\mu m}$ decreases towards higher N. Figure 6.12 shows the polarization ratios versus temperature for the two categories. When plotted versus T there is more scatter in the data. However, for the quiescent data there is a trend of $p_{250\mu m}/p_{350\mu m}$ decreasing with T. For protostellar data, there does not appear to be any correlation wtih T. Figures 6.11 and 6.12 show that outside of RCW 36, the high $p_{250\mu m}/p_{350\mu m}$ regions seen in the top map in Figure 6.6 occur where the number density is the highest, and the temperature is low.

Overall the data in the protostellar category have a higher $p_{250\mu m}/p_{350\mu m}$, and a lower $p_{500\mu m}/p_{350\mu m}$. Figure 6.13 shows the shape of the polarization spectra obtained using the median polarization ratio in each category, compared to the shape of the spectrum for the cloud overall. The spectrum for the quiescent data appears to be similar to that seen for the cloud overall. The spectrum for the protostellar data has more of a decreasing slope, closer to that seen by Zeng et al. (2013). This is consistent with the suggestion in that paper that the presence of embedded sources in a cloud enhances the polarization at shorter wavelengths, shifting the minimum in the polarization spectrum towards longer wavelengths.

In order to separate the effects of temperature from the result shown in Figure 6.11, each category was divided into three temperature bins, and polarization ratios were plotted versus N for each bin. The bins (listed in Table 6.1) were chosen such that each bin contained an equal number of data points ($N = [19, 20, 20]$ for the three protostellar temperature bins, and $N = [141, 140, 140]$ for the three quiescent temperature bins). Figure 6.14 shows polarization ratios versus number density for the three quiescent temperature bins, and Figure 6.15 shows the result for protostellar sight lines. For the quiescent data, the second temperature bin shows

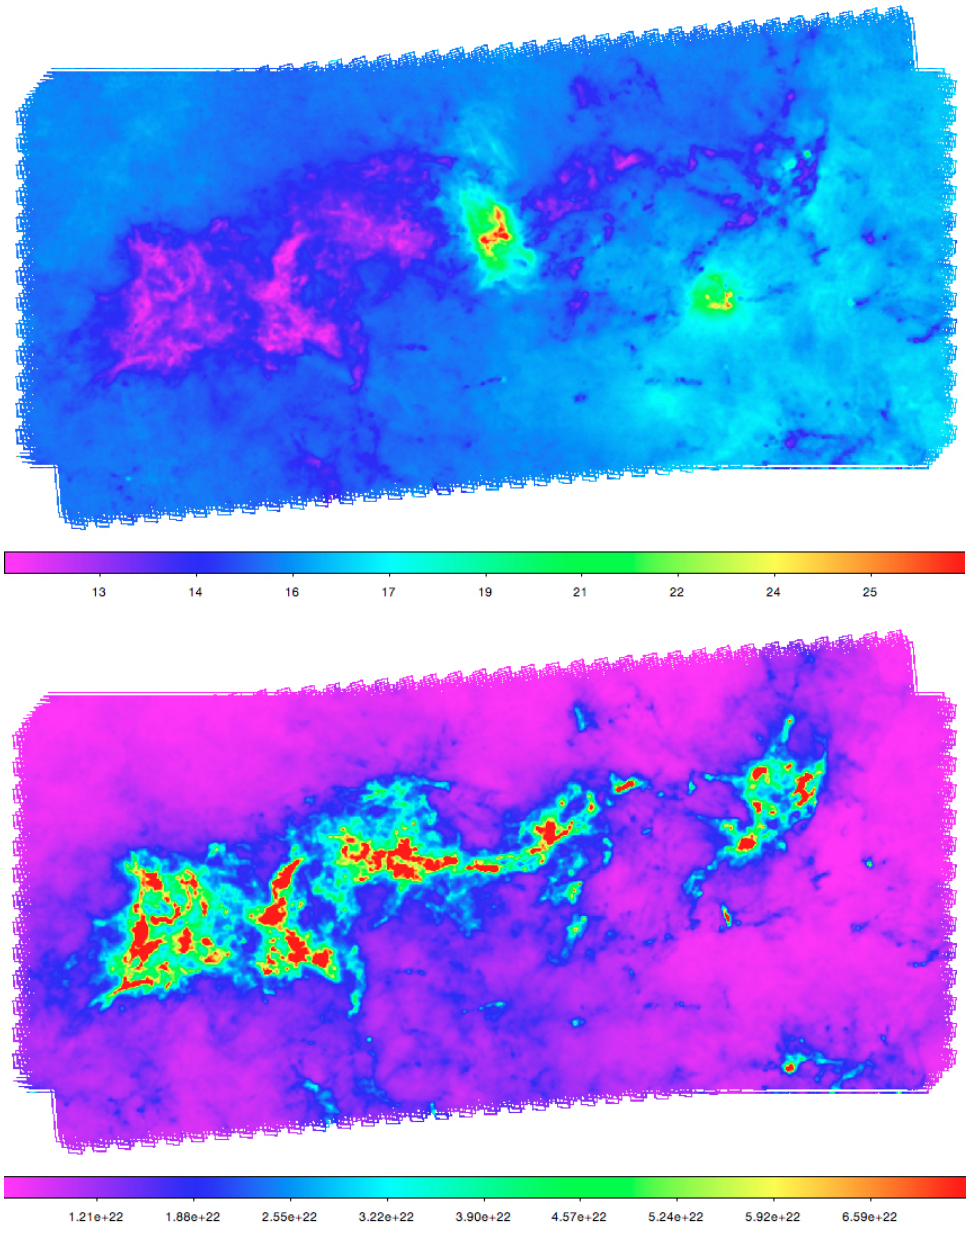

Figure 6.9: Top: temperature (K); Bottom: number density (cm^{-2}).

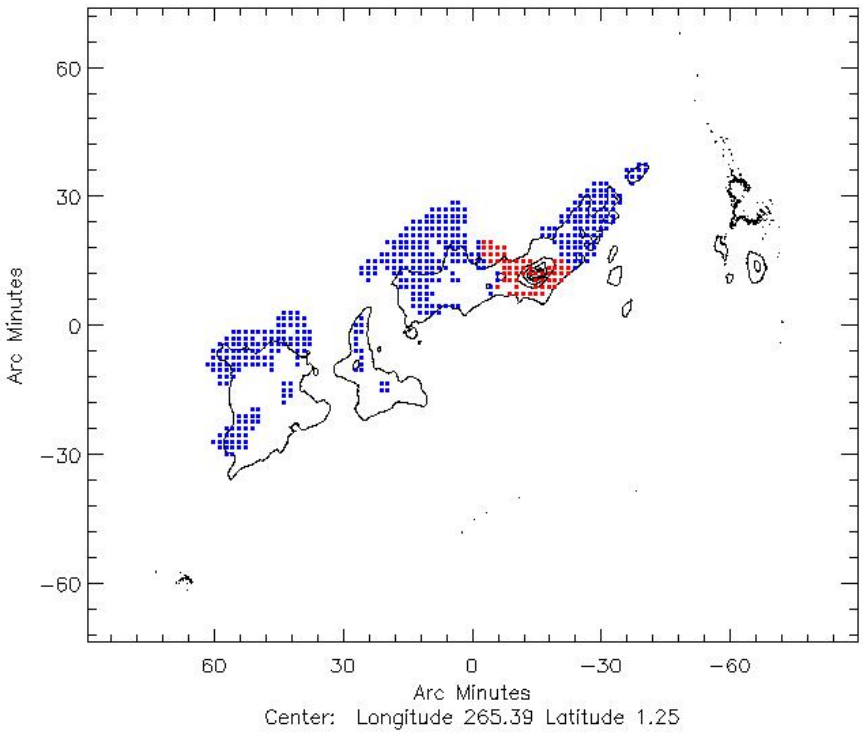

Figure 6.10: Vela C data divided into "protostellar" sightlines (red points) and "quiescent" sightlines (blue points). Contours are 500μm intensity.

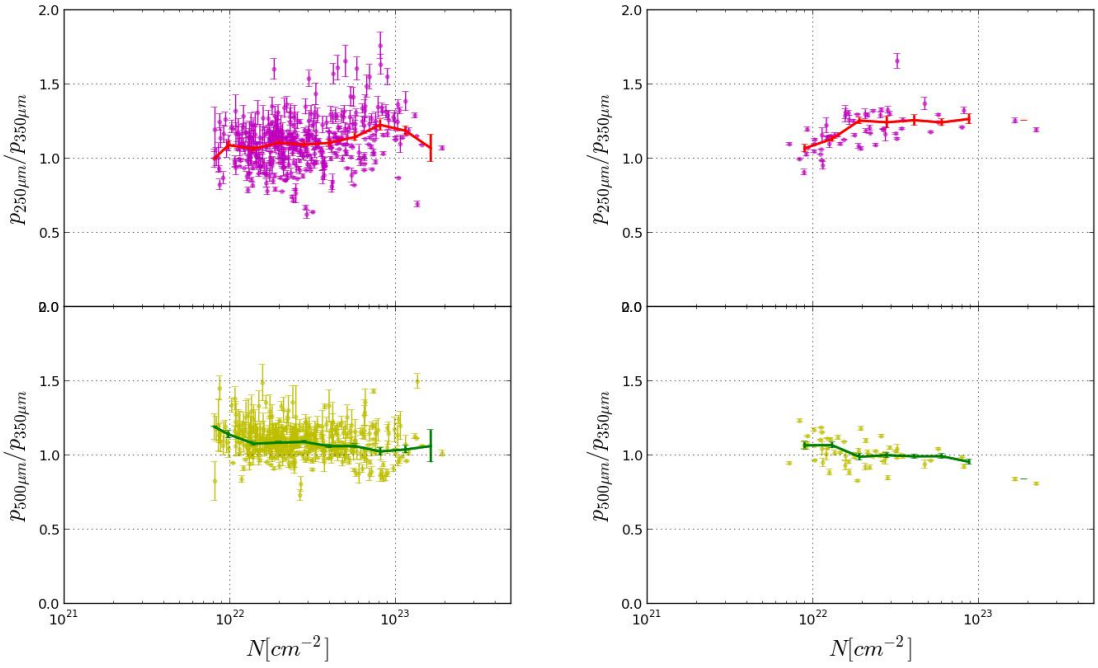

Figure 6.11: Polarization ratios versus number density, for quiescent sightlines (left) and protostellar sightlines (right). Solid lines represent running medians in logarithmic bins and error bars are the standard deviation of the mean in each bin.

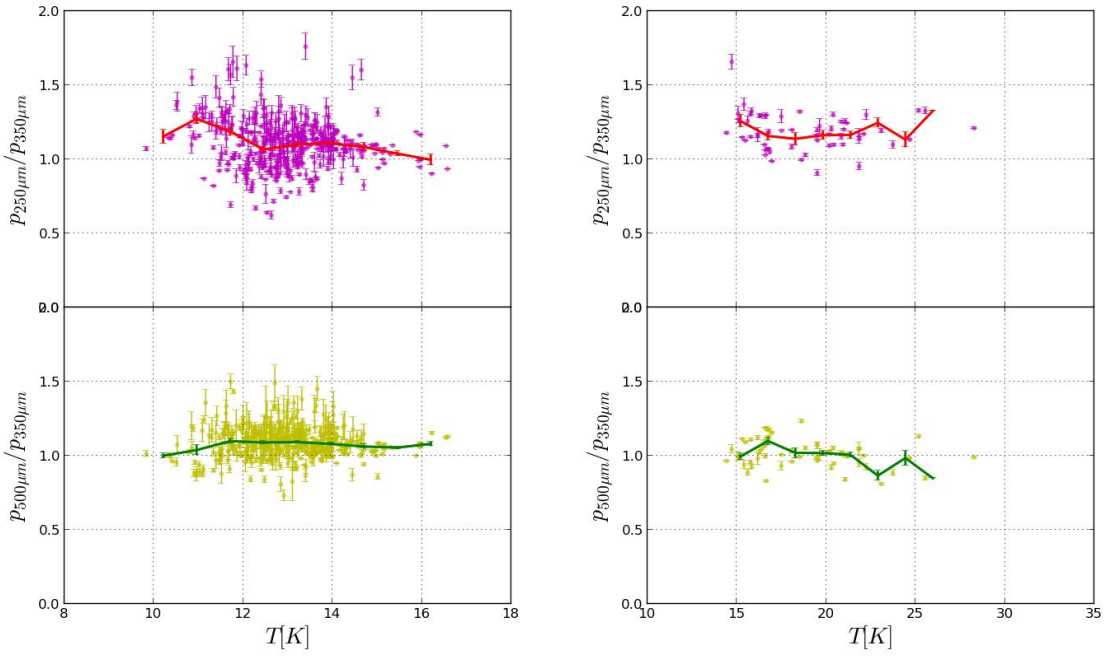

Figure 6.12: Polarization ratios versus temperature, for quiescent sightlines (left) and protostellar sightlines (right). Solid lines represent running medians in logarithmic bins and error bars are the standard deviation of the mean in each bin.

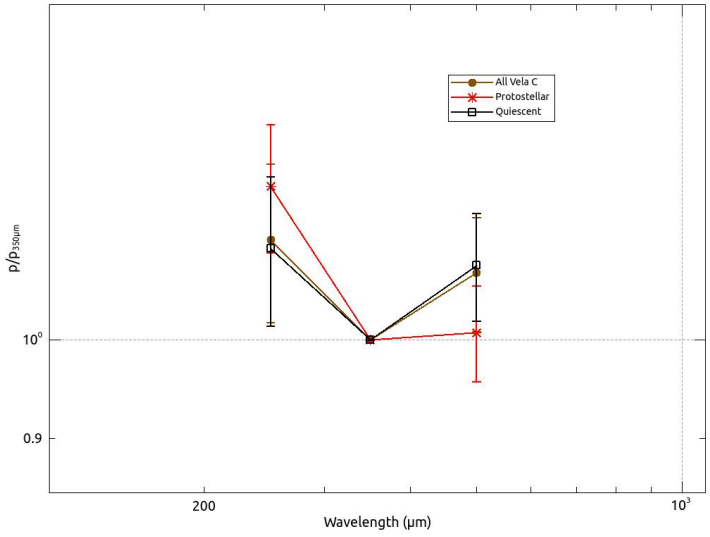

Figure 6.13: Polarization spectrum in protostellar and quiescent regions of Vela C.

Table 6.1: Temperature bins.

Bin	Quiescent	Protostellar
1	$9.8K < T < 12.5K$	$14.4K < T < 16.7K$
2	$12.5K < T < 13.4K$	$16.7K < T < 20.3K$
3	$13.4K < T < 16.6K$	$20.3K < T < 28.3K$

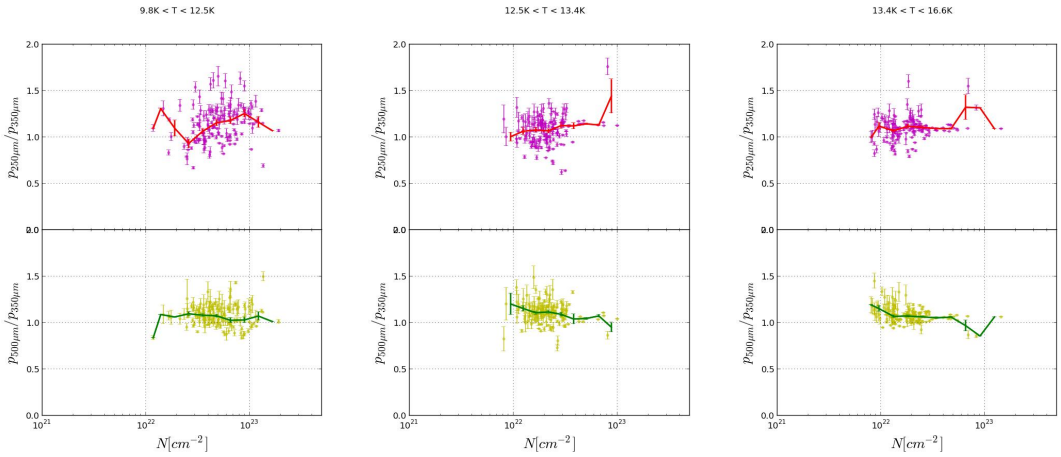

Figure 6.14: Quiescent sightlines. Polarization ratios versus number density, for three temperature bins. Solid lines represent running medians in logarithmic bins and error bars are the standard deviation of the mean in each bin.

the strongest trend with N, while the low and high T bins are flatter. For protostellar data, all three temperature bins show an increase in $p_{250\mu m}/p_{350\mu m}$ with N, while $p_{500\mu m}/p_{350\mu m}$ appears uncorrelated with N in all three bins.

The results of looking at polarization spectrum trends in different radiative environments suggests that the massive star-forming region RCW 36 affects the shape of the polarization spectrum, giving it more of a negative slope. In the quiescent parts of the cloud, $p_{250\mu m}/p_{350\mu m}$ appears to be correlated with high N and low T. RAT theory predicts that grain alignment will be lost in the densest regions of starless clouds. So it is unclear what would cause an enhancement in $p_{250\mu m}/p_{350\mu m}$. Further studies could look for the presence of embedded protostellar sources towards these high N, low T sightlines that might explain this result. Grain alignment by RATs also depends on the anisotropy of the radiation field. Therefore it might be worthwhile to look at the dependence of the polarization spectrum on the angle between the magnetic field and the radiation field (similar to Andersson & Potter (2010)).

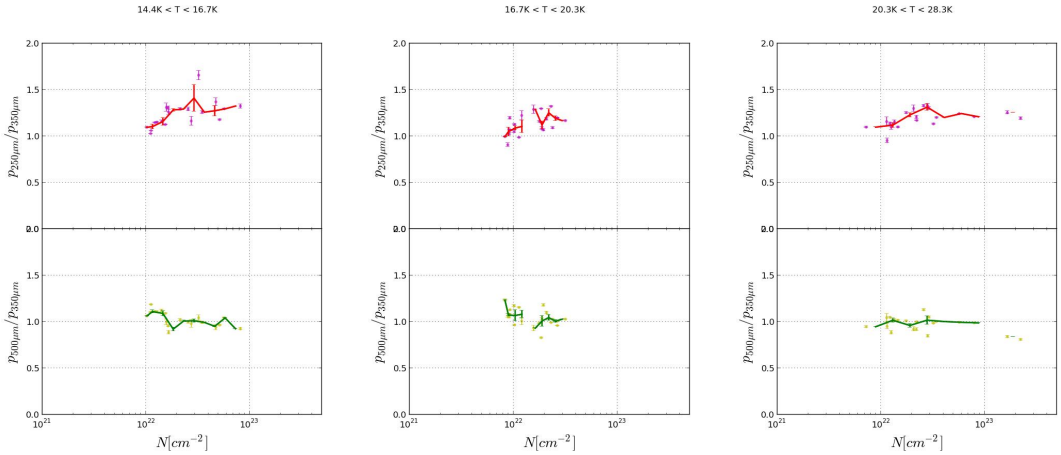

Figure 6.15: Protostellar sightlines. Polarization ratios versus number density, for three temperature bins. Solid lines represent running medians in logarithmic bins and error bars are the standard deviation of the mean in each bin.

Chapter 7

Summary

Figure 7.1: A man's got to know his limitations.

The balloon-borne telescope BLASTPol was designed to observe submillimeter polarization in order to trace the magnetic fields at high resolution over entire star-forming clouds. An attitude reconstruction built for BLASTPol was able to achieve the required pointing capability in-flight, and reconstruct the pointing after the flight. A similar system was designed for the balloon-borne CMB polarimeter, Spider. BLASTPol had two flights over Antarctica in 2010 and 2012, and Spider had its first flight in 2015. BLASTPol suffered from several hardware problems in the 2010 flight, but the 2012 obtained > 40 hours of observations of the Vela C molecular cloud.

BLASTPol's measurement of the submillimeter polarization spectrum in Vela C confirms the minimum at 350μm observed by several previous experiments. The shape of the spectrum appears to change near the massive star-forming region RCW 36, where the polarization at 250μm is increased, and at 500μm it is decreased. There is variation across the cloud in the shape of the spectrum, but this variation is not strongly correlated with temperature or num-

ber density. This result is not inconsistent with radiative grain alignment theory, since the radiative environment of the dust grains in Vela C is not simply represented by N and T. It may be necessary to take into account the angle between the magnetic field and the radiation field, which RAT theory predicts will influence grain alignment efficiency.

Analysis of the BLASTPol data is still ongoing, with other observed targets still to be studied (Carina Nebula, Puppis, Lupus). In Vela C, the diffuse parts of the map ($A_V < 4$) will be accessible once a better zero-point calibration of the map is done. Release of the *Planck* maps has provided the necessary data to do this calibration properly. Ongoing analysis of the cloud is also looking at the polarization signal towards embedded protostellar sources, detected by Herschel and other experiments. A prediction of radiative alignment of dust grains is that the polarization efficiency should be enhanced near embedded sources.

Bibliography

Ade, P. A. R., Aikin, R. W., Barkats, D., Benton, S. J., Bischoff, C. A., Bock, J. J., Brevik, J. A., Buder, I., Bullock, E., Dowell, C. D., Duband, L., Filippini, J. P., Fliescher, S., Golwala, S. R., Halpern, M., Hasselfield, M., Hildebrandt, S. R., Hilton, G. C., Hristov, V. V., Irwin, K. D., Karkare, K. S., Kaufman, J. P., Keating, B. G., Kernasovskiy, S. A., Kovac, J. M., Kuo, C. L., Leitch, E. M., Lueker, M., Mason, P., Netterfield, C. B., Nguyen, H. T., O'Brient, R., Ogburn, R. W., Orlando, A., Pryke, C., Reintsema, C. D., Richter, S., Schwarz, R., Sheehy, C. D., Staniszewski, Z. K., Sudiwala, R. V., Teply, G. P., Tolan, J. E., Turner, A. D., Vieregg, A. G., Wong, C. L., Yoon, K. W., & Bicep2 Collaboration. 2014, Physical Review Letters, 112, 241101

Andersson, B.-G. & Potter, S. B. 2010, Astrophysical Journal, 720, 1045

Bennett, C. L., Larson, D., Weiland, J. L., Jarosik, N., Hinshaw, G., Odegard, N., Smith, K. M., Hill, R. S., Gold, B., Halpern, M., Komatsu, E., Nolta, M. R., Page, L., Spergel, D. N., Wollack, E., Dunkley, J., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., & Wright, E. L. 2013, The Astrophysical Journal Supplement, 208, 20

Bethell, T. J., Chepurnov, A., Lazarian, A., & Kim, J. 2007, Astrophysical Journal, 663, 1055

BICEP2/Keck, Planck Collaborations, :, Ade, P. A. R., Aghanim, N., Ahmed, Z., Aikin, R. W., Alexander, K. D., Arnaud, M., Aumont, J., & et al. 2015, ArXiv e-prints

Chandrasekhar, S. & Fermi, E. 1953, Astrophysical Journal, 118, 113

Dolginov, A. Z. & Mitrofanov, I. G. 1976, Astrophysics and Space Science, 43, 291

Draine, B. T. 2003, Annual Review of Astronomy & Astrophysics, 41, 241

Draine, B. T. & Fraisse, A. A. 2009, Astrophysical Journal, 696, 1

Draine, B. T. & Weingartner, J. C. 1996, Astrophysical Journal, 470, 551

—. 1997, Astrophysical Journal, 480, 633

Filippini, J. P., Ade, P. A. R., Amiri, M., Benton, S. J., Bihary, R., Bock, J. J., Bond, J. R., Bonetti, J. A., Bryan, S. A., Burger, B., Chiang, H. C., Contaldi, C. R., Crill, B. P., Doré, O., Farhang, M., Fissel, L. M., Gandilo, N. N., Golwala, S. R., Gudmundsson, J. E., Halpern, M., Hasselfield, M., Hilton, G., Holmes, W., Hristov, V. V., Irwin, K. D., Jones, W. C., Kuo, C. L., MacTavish, C. J., Mason, P. V., Montroy, T. E., Morford, T. A., Netterfield, C. B., O'Dea, D. T., Rahlin, A. S., Reintsema, C. D., Ruhl, J. E., Runyan, M. C., Schenker, M. A., Shariff, J. A., Soler, J. D., Trangsrud, A., Tucker, C., Tucker, R. S., & Turner, A. D. 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7741, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series

Fissel, L. M. 2013, PhD thesis, University of Toronto

Girart, J. M., Beltrán, M. T., Zhang, Q., Rao, R., & Estalella, R. 2009, Science, 324, 1408

Harman, R. R. 2005

Hildebrand, R. H., Dotson, J. L., Dowell, C. D., Schleuning, D. A., & Vaillancourt, J. E. 1999, Astrophysical Journal, 516, 834

Hill, T., Motte, F., Didelon, P., Bontemps, S., Minier, V., Hennemann, M., Schneider, N., André, P., Men'shchikov, A., Anderson, L. D., Arzoumanian, D., Bernard, J.-P., di Francesco, J., Elia, D., Giannini, T., Griffin, M. J., Könyves, V., Kirk, J., Marston, A. P., Martin, P. G., Molinari, S., Nguyen Luong, Q., Peretto, N., Pezzuto, S., Roussel, H., Sauvage, M., Sousbie, T., Testi, L., Ward-Thompson, D., White, G. J., Wilson, C. D., & Zavagno, A. 2011, Astronomy & Astrophysics, 533, A94

Kisner, T. S. 2015, Time Ordered Astrophysics Scalable Tools (TOAST), [Online; accessed 21-April-2015]

Lazarian, A. & Hoang, T. 2007, Monthly Notics of the Royal Astronomical Society, 378, 910

Li, H., Griffin, G. S., Krejny, M., Novak, G., Loewenstein, R. F., Newcomb, M. G., Calisse, P. G., & Chuss, D. T. 2006, Astrophysical Journal, 648, 340

Mac Low, M.-M. & Klessen, R. S. 2004, Reviews of Modern Physics, 76, 125

MacTavish, C. J., Ade, P. A. R., Battistelli, E. S., Benton, S., Bihary, R., Bock, J. J., Bond, J. R., Brevik, J., Bryan, S., Contaldi, C. R., Crill, B. P., Doré, O., Fissel, L., Golwala, S. R., Halpern, M., Hilton, G., Holmes, W., Hristov, V. V., Irwin, K., Jones, W. C., Kuo, C. L., Lange, A. E., Lawrie, C., Martin, T. G., Mason, P., Montroy, T. E., Netterfield, C. B., Riley, D., Ruhl, J. E., Runyan, M., Trangsrud, A., Tucker, C., Turner, A., Viero, M., & Wiebe, D. 2008, Astrophysical Journal, 689, 655

Marsden, G. 2006, PhD thesis, University of British Columbia

Moncelsi, L., Ade, P. A. R., Angilè, F. E., Benton, S. J., Devlin, M. J., Fissel, L. M., Gandilo, N. N., Gundersen, J. O., Matthews, T. G., Netterfield, C. B., Novak, G., Nutter, D., Pascale, E., Poidevin, F., Savini, G., Scott, D., Soler, J. D., Spencer, L. D., Truch, M. D. P., Tucker, G. S., & Zhang, J. 2014, onthly Notices of the Royal Astronomical Society, 437, 2772

Mortari, D., Junkins, J. L., & Samaan, M. 2001, in Guidance and control 2001, 49–68

Ostriker, E. C., Stone, J. M., & Gammie, C. F. 2001, Astrophysical Journal, 546, 980

Padoan, P., Federrath, C., Chabrier, G., Evans, II, N. J., Johnstone, D., Jørgensen, J. K., McKee, C. F., & Nordlund, Å. 2014, Protostars and Planets VI, 77

Pascale, E. 2007, PhD thesis, University of Toronto

Pascale, E., Ade, P. A. R., Bock, J. J., Chapin, E. L., Chung, J., Devlin, M. J., Dicker, S., Griffin, M., Gundersen, J. O., Halpern, M., Hargrave, P. C., Hughes, D. H., Klein, J., MacTavish, C. J., Marsden, G., Martin, P. G., Martin, T. G., Mauskopf, P., Netterfield, C. B., Olmi, L., Patanchon, G., Rex, M., Scott, D., Semisch, C., Thomas, N., Truch, M. D. P., Tucker, C., Tucker, G. S., Viero, M. P., & Wiebe, D. V. 2008, Astrophysical Journal, 681, 400

Powell, M. J. D. 1964, The Computer Journal, 7, 155

Rex, M., Chapin, E., Devlin, M. J., Gundersen, J., Klein, J., Pascale, E., & Wiebe, D. 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 6269, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series

Shariff, J. A. 2015, PhD thesis, University of Toronto

Shariff, J. A., Ade, P. A. R., Amiri, M., Benton, S. J., Bock, J. J., Bond, J. R., Bryan, S. A., Chiang, H. C., Contaldi, C. R., Crill, B. P., Doré, O. P., Farhang, M., Filippini, J. P., Fissel, L. M., Fraisse, A. A., Gambrel, A. E., Gandilo, N. N., Golwala, S. R., Gudmundsson, J. E., Halpern, M., Hasselfield, M., Hilton, G. C., Holmes, W. A., Hristov, V. V., Irwin, K. D., Jones, W. C., Kermish, Z. D., Kuo, C. L., MacTavish, C. J., Mason, P. V., Megerian, K. G., Moncelsi, L., Morford, T. A., Nagy, J. M., Netterfield, C. B., O'Brient, R., Rahlin, A. S., Reintsema, C. D., Ruhl, J. E., Runyan, M. C., Soler, J. D., Trangsrud, A., Tucker, C. E., Tucker, R. S., Turner, A. D., Weber, A. C., Wiebe, D. V., & Young, E. Y. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9145, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 0

Soler, J. D., Hennebelle, P., Martin, P. G., Miville-Deschênes, M.-A., Netterfield, C. B., & Fissel, L. M. 2013, Astrophysical Journal, 774, 128

Vaillancourt, J., Andersson, B. G., & Lazarian, A. 2013, in Proceedings of The Life Cycle of Dust in the Universe: Observations, Theory, and Laboratory Experiments (LCDU2013). 18-22 November, 2013. Taipei, Taiwan. Editors: Anja Andersen (University of Copenhagen, Denmark), Maarten Baes (Universiteit Gent, Belgium), Haley Gomez (Cardiff University, UK), Ciska Kemper (Academia Sinica, Taiwan), Darach Watson (University of Copenhagen, Denmark). Online at ¡A href="http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=207"¿http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=207¡/A¿, id.3, 3

Vaillancourt, J. E. 2002, The Astrophysical Journal Supplement Series, 142, 53

Vaillancourt, J. E., Dowell, C. D., Hildebrand, R. H., Kirby, L., Krejny, M. M., Li, H.-b., Novak, G., Houde, M., Shinnaga, H., & Attard, M. 2008, Astrophysical Journal, 679, L25

Vaillancourt, J. E. & Matthews, B. C. 2012, The Astrophysical Journal Supplement, 201, 13

Wiebe, D. V. 2007, PhD thesis, University of Toronto

Zeng, L., Bennett, C. L., Chapman, N. L., Chuss, D. T., Jimenez-Serra, I., Novak, G., & Vaillancourt, J. E. 2013, The Astrophysical Journal, 773, 29