

PROBING EARLY UNIVERSE COSMOLOGIES WITH
SPIDER & PLANCK HFI

JÓN EMIL GUÐMUNDSSON

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHYSICS

ADVISER: PROFESSOR WILLIAM C. JONES

SEPTEMBER 2014

© Copyright by Jón Emil Guðmundsson, 2014.

All rights reserved.

Abstract

We describe the design and characterization of the SPIDER balloon-borne experiment which will launch from Antarctica in December 2014. The experiment is designed to measure the polarization of the Cosmic Microwave Background (CMB) with unparalleled instantaneous sensitivity, and in doing so, constrain early universe models. The experiment is ready to deploy. We will emphasize: 1) The cryogenic architecture of the SPIDER flight cryostat. 2) The design and characterization of a capillary assembly which provides a continuous flow of superfluid helium to a 1.8 K temperature stage required to operate adsorption refrigerators cooling each focal plane. 3) The design and build of a Fourier transform spectrometer used to characterize the spectral response of the SPIDER detectors. 4) The optical characterization of the SPIDER telescopes and simulations characterizing susceptibility to polarized sidelobes contamination.

We also describe the analysis of the spatial (beam) response of the High Frequency Instrument (HFI) onboard the *Planck* satellite. This characterization work is required for high fidelity cosmological analysis of all-sky maps in six frequency bands spanning 100 to 857 GHz. The beam reconstruction error and bias are constrained using time-domain simulations that include the most significant non-idealities that affect the analysis. Using these simulations, we also verify the consistency of the beam product used for cosmological analysis of the 2014 data release. As part of the beam reconstruction, we characterize the flux of the five outer planets in the six HFI frequency bands. We also verify the absolute photometric calibration of the experiment by comparing planet flux estimates with existing models. Finally, we use planet flux measurements to show that the absolute calibrations of the *WMAP* and *Planck* experiments are consistent at the half percent level.

Ágrip

Í þessu doktorsverkefni fjöllum við um hönnun og kvörðun á SPIDER loftbelgs-tilrauninni sem skotið verður upp frá Suðurskautslandinu í desember 2014. Tilrauninni er ætlað að mæla skautun örbylgjukliðsins með meiri næmni en nokkur önnur núverandi tilraun. Með þessum mælingum mun SPIDER skorða líkön er snúa að frumbernsku alheimsins. Hún er nú tilbúin til notkunnar. Hér munum við leggja áherslu á: 1) Hönnun kuldahalds sem hýsir SPIDER sjónaukana. 2) Hönnun og kvörðun á ofurþráðum sem veita stöðugt flæði af ofurflæðandi helíum yfir í 1.8 K hitasvið sem að knýr ásogskælibúnað sem kælir ljósflögur sérhvers sjónauka. 3) Hönnun og smíði litrófsgreinis sem notaður er til þess að kvarða litrófssvörum SPIDER sjónaukanna. 4) Kvörðun á ljósgeislum SPIDER sjónaukanna og hermanir sem að skorða viðtak vegna skautaðra hliðargeira.

Í þessari ritgerð munum við einnig lýsa rannsóknum á ljósgeislum High Frequency Instrument um borð í *Planck* gervitunglinu. Þessi kvörðun er nauðsynleg fyrir nákvæma heimsfræðilega greiningu á himinkortum sem spanna sex tíðnisvið frá 100 til 857 GHz. Óvissa og bjagi við kvörðun ljósgeislanna er áætluð með hermunum á tímaröðum sem innihalda helstu óvissu- og bjögunarþætti. Með þessum hermunum sýnum við fram á samkvæmni áætlaðra ljósgeisla sem að notaðir eru fyrir heimsfræðiniðurstöður sem að birtast í lok ársins 2014. Við staðfestum einnig ljósafskvörðun mælitækjanna með því að bera saman niðurstöður á mældu ljósafli ytri reikistjarnanna við spágildi líkanna. Við notum mælingar á ljósafli þessara reikistjarna til þess að sýna fram á 0.5% samkvæmni í ljósafskvörðun *WMAP* og *Planck* gervitunglanna.

Recognition

This thesis was typeset using the L^AT_EX markup language using T_EXShop as a front-end to a T_EX Live distribution. To get things started, I used an unofficial Ph.D. dissertation template for Princeton University created by Jeffrey Dwoskin. Thanks for sharing the document! Many thanks to Professor John Collins, Pennsylvania State University, for making his “Latexmk” script available in the public domain, see <http://www.ctan.org/pkg/latexmk/>. Many figures in this document use colors from www.ColorBrewer.org, provided by Professor Cynthia A. Brewer, Geography, also at Penn State. Analysis was primarily performed using Python [1], Matlab [2], and associated libraries and packages. Figures were made using tools provided in Matlab and Matplotlib [3]. Some of the results presented here used the HEALPix [4], PolSpice [5], and CAMB [6] packages.

Acknowledgements

First and foremost I would like to thank my advisor, Bill Jones, for his unwavering support and encouragement. His bright spirit and sharp insight illuminate any situation. It has been a privilege.

I have been blessed with many great instructors throughout my life. To name a few: Gunnar Ásgeirsson and Jón Pétur Zimsen at Réttarholtsskóli. Birgir Guðjónsson, Hrjóbjartur Örn Guðmundsson, and Davíð Þorsteinsson at Menntaskólinn í Reykjavík. Ari Ólafsson, Einar H. Guðmundsson, Lárus Thorlacius, and Viðar Guðmundsson at the University of Iceland. While at Princeton, Aurelien Fraisse, Brendan Crill, Cynthia Chiang, and Zigmund Kermish have all been informal mentors, glad to debug my code and answer any questions I might have. You guys set a great example! I am thankful to Norman Jarosik for his constant guidance. Suzanne, Lyman, and the entire Gravity Group provide a wonderful learning environment. In particular, I would like to thank Jim Peebles for pleasant conversations about Λ CDM and for agreeing, along with Bruce Draine, to be on the thesis committee with five day notice.

The administrative and technical staff at the Princeton Physics Department is surely beyond compare. There are so many people to thank, just to name a few: Claudin Champagne, William Dix, Jessica Heslin, Geoffrey Gettelfinger, Darryl Johnson, Jim Kukon, Angela Lewis, Ted Lewis, and Steven Lowe. It has been a pleasure to work with all of you!

Thanks to the SPIDER collaboration – such a great group of hard working people – for the countless delightful moments. The diverse range of talents and perspectives has helped us develop and mature. Ohio was nice, Texas was wonderful, I’m sure Antarctica will be something else.

As a member of the *Planck* HFI core team, I have had the privilege of working with a large number of incredibly talented and friendly individuals, all of whom have been tireless in assisting me grasp the various *Planck*-related concepts. I would especially like to thank Brendan Crill, Duncan Hanson, Eric Hivon, Gaël Roudier, and Jean-Loup Puget.

Many thanks to Aurelien, Bruce, Cynthia, Dóra, and Norm for reading over the thesis and providing insightful comments. Anne, you are a champ for reading the entire thing and fixing my broken grammar. I owe you infinite pancakes and/or french toast.

My fellow grad students have been wonderful. Sasha Rahlin, thank you for your constant assistance – you have taught me so much. Edward Young and Anne Gambrel, you guys were a true blessing for the group. It's been a blast working with you.

Andrew, David, Doug, Jack, and all the others. Thank you for introducing us to pancakes, pecan pie, bok choy, american football, power hour, and Smash Brothers! You guys are truly wonderful friends. Football (the English kind) enlivened the days in Princeton. Gonzalo, Thanh, GC Opal, and the rest, may we play again! Geir and Valla, the first years in Princeton were adorned with your warmth and hospitality. It would not have been the same without you. Also, Lilian Grosz, our American grandmother who took us under her wing when we arrived, we are forever grateful. You are a hoot!

Dóra, my now wife, thank you for taking a leap of faith and moving to Princeton and for allowing our relationship to flourish while inuring my perpetual cryogenic monitoring. Your care and support made this possible. Finally, I have to thank my family. Especially, my mother Valgerður, my father Guðmundur, and my sister Álfheiður Erla. I owe you everything.

Princeton, August 26, 2014

Þessi ritgerð er tileinkuð afa mínum, Jóni R. Einarssyni.

Contents

Abstract	iii
Ágrip	iv
Recognition	v
Acknowledgements	vi
List of Tables	xiii
List of Figures	xiv
1 Introduction	1
1.1 Modern Cosmology Theory – μ Review	8
1.2 Observational Cosmology	23
1.3 Thesis Work and Content	27
2 SPIDER: A Balloon-Borne CMB Polarization Experiment	29
2.1 Instrument Overview	29
2.1.1 Foregrounds and Noise	30
2.1.2 Recent Developments	37
2.2 Collaboration	39
2.3 Ballooning	40
2.4 Cryogenic Architecture	44
2.4.1 Radiation Shields, Heat Exchangers, and MLI	47
2.4.2 Flexures and Plumbing	48

2.4.3	Cryogenic Operations	51
2.4.4	Cryogenic Qualification	55
2.4.5	Thermal Model	57
2.4.6	Helium Leak	59
2.5	Capillaries	61
2.5.1	Brief Literature Review	62
2.5.2	Design	64
2.5.3	Experimental Results	66
2.5.4	Performance Characterization	72
2.5.5	Launch Configuration	74
2.6	The Telescopes	76
2.7	Half-Wave Plate	79
2.8	Detector Architecture	84
2.9	Fourier Transform Spectroscopy	89
2.9.1	FTS Design	92
2.10	Outer Frame	99
2.11	Current Status	103
3	Calibration of the <i>Planck</i> High Frequency Instrument	104
3.1	Instrument Overview	105
3.2	Scan Strategy and Pointing	108
3.3	Absolute Photometric Calibration	110
3.3.1	Beam Definitions	113
3.3.2	Planet Observations	115
3.4	Beam Reconstruction	116
3.4.1	The Gauss-Hermite Function Basis on \mathbb{R}^2	118
3.5	Time-Domain Simulations	121
3.5.1	Methodology	121

3.5.2	Quantitative Description of Simulation Framework	123
3.5.3	Non-Idealities	125
3.5.4	Error Budget	135
3.5.5	Principal Component Analysis	137
3.6	Data Derived Beams	139
3.6.1	Pipeline Comparison	140
3.6.2	Consistency Tests	143
3.6.3	Dimpling Lobes and Ruze Envelopes	146
3.6.4	Bolometer Non-Linearity and Saturation	149
3.7	Validation of Absolute Calibration using Planets	150
3.7.1	Planet Brightness Measurements	151
3.7.2	Mars	153
3.7.3	Jupiter	155
3.7.4	Saturn	157
3.7.5	Uranus and Neptune	159
3.7.6	Summary	161
3.7.7	Tabulated Results	162
4	Calibrations and Simulations of the SPIDER Optical System	163
4.1	Baffles and Far Sidelobes	163
4.1.1	Baffle Design and Sidelobe Phenomenology	164
4.1.2	Spurious Polarization from Baffle Diffraction	166
4.2	Beam Maps	169
4.2.1	Stacking Beams	171
4.2.2	Optical Ghosting	177
5	Conclusions	181

A Thermal Modeling	184
A.1 Conduction	184
A.2 Radiation	186
B Fourier Transform Spectroscopy	188
B.1 Wire Spacing of Polarizing Beamsplitters	188
B.2 Propagation of Electric Fields in an FTS	190
B.3 Collimating mirror	194
B.4 G-code for collimating mirror	195
C Gauss-Hermite Math	198
C.1 Preliminaries	198
C.2 Gauss-Hermite Decomposition	199
C.3 Window Functions	202
C.4 Ruze Envelope	206
D Photometric Calibration Using	
Planet Flux Estimates	208
D.1 Describing Flux	208
D.2 Estimating Flux from Peak Signal	210
D.2.1 Estimating Photometric Parameters	214
D.2.2 Defining Brightness Temperature – Subtle Differences	215
D.3 Color Corrections	217
D.4 Measured Planet Fluxes	219
Bibliography	225

List of Tables

2.1	Primary characteristics of the SPIDER experiment	32
2.2	Flight cryostat main specifications	45
2.3	Cryogenic tasks required for cooldown	51
2.4	Cryogenic score sheet	53
2.5	Thermal properties of ^4He and N_2	58
2.6	Measured properties of various capillary assemblies	70
2.7	Spectral properties of SPIDER's telescopes	91
3.1	Summary of the <i>Planck</i> HFI	105
3.2	Summary of <i>Planck</i> planet observations	116
3.3	The mean value of DX11 scanning beam parameters	142
3.4	Beam window function consistency test	146
4.1	Telescope averaged beam properties	173
D.1	The band average color corrections	218
D.2	Flux table 1	220
D.3	Flux table 4	221
D.4	Flux table 3	222
D.5	Flux table 4	223
D.6	Flux table 5	224

List of Figures

1.1	The spectrum of the CMB	4
1.2	Results from the 2dF Galaxy Redshift Survey	7
1.3	The CMB temperature power spectrum	16
1.4	<i>E</i> - and <i>B</i> -mode polarization patterns	17
1.5	Published <i>EE</i> and <i>BB</i> polarization power spectra	21
1.6	The CMB as observed by <i>COBE</i> , <i>WMAP</i> , and <i>Planck</i>	24
2.1	The SPIDER payload	31
2.2	CMB foreground maps	35
2.3	Polarized foreground power spectra	39
2.4	Histogram of LDB flight durations	42
2.5	The average Antarctic LDB trajectory	43
2.6	A cross section of the flight cryostat	46
2.7	The view of the bottom of the main tank	47
2.8	CAD model of the vacuum vessel	49
2.9	CAD of main structural flexures and heat exchangers	50
2.10	Pumpdown profiles for Runs 13 and 15	54
2.11	A CAD model of the custom bolometer	56
2.12	Predictions from the thermal model	59
2.13	Adsorption isotherms of helium on activated charcoal	61
2.14	The capillary assembly	63

2.15	A cross section through one of the capsules	65
2.16	The superfluid tank and capillary assembly	68
2.17	Cryogenic data from the MT and SFT	71
2.18	Atmospheric pressure as a function of time from launch	75
2.19	The SPIDER telescope assembly	78
2.20	The SPIDER HWP	81
2.21	The properties of typical TES bolometers	85
2.22	The main components of a SPIDER detector	86
2.23	A collection of photos showing the SPIDER TES architecture	87
2.24	The SPIDER TES thermoelectrical circuit	88
2.25	The average spectra of all six focal planes	90
2.26	The path of light in a Martin-Puplett interferometer	92
2.27	A CAD model of the entire FTS assembly	94
2.28	The shape of the collimating mirror	96
2.29	The wire grid beamsplitter and the collimating mirror	97
2.30	The SPIDER FTS on a lab stool	99
2.31	A rendering of the Gondola assembly	100
2.32	Simulated SPIDER coverage map	101
3.1	A drawing of the <i>Planck</i> satellite	106
3.2	A depiction of <i>Planck</i> 's scan strategy	108
3.3	A collage of <i>Planck</i> HFI maps	117
3.4	The first few Gauss-Hermite functions	120
3.5	A flow diagram describing the simulation methodology	122
3.6	Simulated beam map, coverage map, and noise map	126
3.7	Noise only simulation results	128
3.8	Simulations of pointing error	129
3.9	The co- and cross-scan centroid offset	130

3.10	Effects of non-uniform coverage	133
3.11	Error budget	136
3.12	Simulations of pointing error	138
3.13	Comparison of beam products	141
3.14	<i>Planck</i> HFI beam profiles	144
3.15	B-spline and GH comparison	145
3.16	The sub-millimeter beams	147
3.17	The <i>Planck</i> reflectors	148
3.18	Mars, Jupiter, Saturn, and Uranus	152
3.19	Mars brightness temperature	154
3.20	Jupiter brightness temperature	156
3.21	Saturn brightness temperature	158
3.22	Uranus brightness temperature	159
3.23	Neptune brightness temperature	160
3.24	Measured planet flux compared to models	162
4.1	A section view rendering of the SPIDER flight cryostat	164
4.2	A cross section through the SPIDER telescope	165
4.3	Measured sidelobe profiles	167
4.4	Sidelobe smoothed temperature map	168
4.5	The B-mode power spectra generated by Galactic pickup	169
4.6	Beam maps of X6 and X3	170
4.7	SPIDER's azimuthally averaged beam profiles	172
4.8	SPIDER's beam window functions	174
4.9	SPIDER's beam window function ratios	175
4.10	Beam difference maps	176
4.11	A ghost map	177
4.12	Ghosts abound	178

4.13 The location of ghost beams	179
B.1 Optimal wire grid dimensions as a function of frequency	189
B.2 Cross polar leakage as a function of frequency	190
B.3 Wire grid coordinate systems	191
B.4 The shape of the collimating mirror	196

Chapter 1

Introduction

Erst was the age, when nothing was;
sand was not, nor sea, nor biting surf;
Earth was not found, nor heaven above;
A yawning gap, and no grass.

– The Beguiling of Gylfi, from the Prose Edda

Cosmology is the study of the most extensive astrophysical phenomena. It embodies the brash notion that the universe and its development can be characterized with at least some scientific rigor. The ideas and questions involved most certainly predate the stone age, but historical evidence suggests spells of enlightenment at numerous times in human history. An explosion of 20th century observations spanning the entire electromagnetic spectrum has left us with a multi-faceted picture of the universe which can take a lifetime to process. Yet many argue that the properties of our universe can be captured by half a dozen parameters in a straightforward physical model. In this chapter I will briefly describe some of the more notable discoveries of modern cosmology with emphasis on experimental results. I will provide an overview of the physics of the Cosmic Microwave Background (CMB) and conclude by describing current experimental efforts in cosmology.

One could trace the beginning of modern cosmology to the development of large aperture refracting telescopes which allowed astronomers to resolve the shape of far away galaxies. For more than a century, we observed these phenomena through our telescopes unable to prove

that they were independent island universes. It was not until the early 20th century when, working at the Harvard College Observatory, Henrietta Leavitt discovered a relation between the luminosity and period of Cepheid variable stars. This established a galactic distance calibrator [7] which in conjunction with spectroscopy allowed Hubble, Humason, Slipher, and others to gather evidence for a linear relation between the redshift of nearby galaxies, caused by their recessional velocity, v , and distance, d ,

$$v = Hd. \quad (1.1)$$

This is Hubble's law, where H is the famous proportionality constant also named after Hubble. The first estimates of this value were rather crude, with Hubble's 1929 value, derived from two dozen systems, at $H = 500$ km/s/Mpc [8], compared to modern estimates approximating to $H = 70$ km/s/Mpc. This cosmic expansion was a remarkable concept, considering that only five years prior to Hubble's initial estimate, a New York Times article [9] revealed findings that Andromeda had to be a distant galaxy separate from our own. The first paragraph read:

“Washington, Nov. 22. – Confirmation of the view that the spiral nebulae, which appear in the heavens as whirling clouds, are in reality distant stellar systems, or ‘island universes,’ has been obtained by Dr. Edwin Hubbell [sic] of the Carnegie Institution’s Mount Wilson observatory, through investigations carried out with the observatory’s powerful telescopes.”

Like many concepts at the bleeding edge of science, the discovery was gradual, as other astronomers had reported similar results [10].

Two years prior to Hubble's famous 1929 paper, a Belgian physicist named Georges Lemaître had published an article in an obscure journal of his home country [11]. The article was not translated into English until 1931 [12], sparked by Hubble's nascent law. Using data gathered at Mount Wilson, and guided by Einstein's general theory of relativity, Lemaître devised a theory of the expanding universe. The proposed idea did not attract

much attention even after the article had been translated to English, possibly due to inertia in the scientific community and somewhat exacerbated by Einstein’s preference for a static universe. Lemaître’s initial work described a dynamically evolving universe, not necessarily one with a beginning. He eventually envisioned and described a “primeval atom” [13], formally addressing the concept of the beginning of time [14]. In these manuscripts, Lemaître imagined a quantum sphere that gave birth to the cosmos through a series of radioactive decays. It was therefore fitting that his ideas proved to be a stepping stone for a group of nuclear physicists.

Another block in the theoretical foundation of cosmology was set in 1948 with Gamow’s and Alpher’s formulation of big bang nucleosynthesis [15]. In this work the authors showed how the lighter elements – hydrogen, helium, and lithium – could be formed in the first minutes of a hot dense universe. That same year Alpher and Herman estimated the temperature of a corresponding radiation afterglow at 5 K [16].¹ The following year, an English astronomer by the name of Fred Hoyle first coined the term “Big Bang” to describe a universe with a violent beginning.² He was himself a proponent of the steady-state theory which surmised that matter is created continuously as the universe expands. At that time, galactic redshift remained as the main observational evidence for the big bang scenario. In 1957 Hoyle together with Margaret Burbidge, Geoffrey Burbidge, and William Fowler, showed how elements heavier than lithium can be produced in stellar fusion [17]. This extensive paper established the theory of stellar nucleosynthesis and showed how even the heaviest elements can be produced during explosions of supernovae. At this point all the scientific ingredients of a hot big bang scenario had been prepared; the idea only lacked observational evidence.

In the 1960’s a group of physicists at Princeton University began to search for the presence of thermal radiation remaining from the primordial fireball. Their research was directed by Robert Dicke, who by then had made contributions to a wide range of physics, including radar development and atomic theory, but focused now on gravitation theory. A member of that group, Jim Peebles, independently derived the results of Alpher and Herman [19].

¹Surprisingly consistent with the correct value, $T_{\text{CMB}} = 2.73$ K.

²It has been suggested that Hoyle meant for the term to be derogatory.

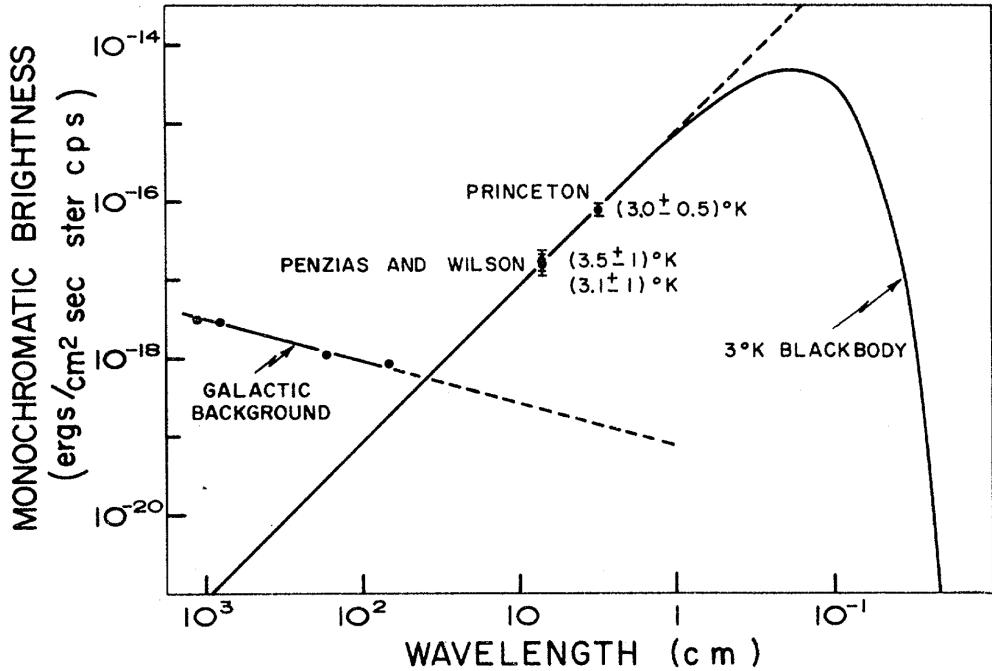


Figure 1.1: The results of the two New Jersey measurements, published in 1966 [18], showing intensity as a function of wavelength with a 3 K blackbody spectrum plotted for comparison. Both the Bell Labs (7.35 cm) and the Princeton measurement (3.2 cm) were performed safely within the Rayleigh-Jeans limit, almost two orders of magnitude below the peak value of a 3 K blackbody. Figure reproduced courtesy of P. J. E. Peebles.

Using microwave receiver technology that Dicke had developed twenty years earlier, called the Dicke radiometer, the Princeton group set about measuring this relic radiation on the roof of Guyot Hall in 1964. The experimental effort was led by Roll and Wilkinson [20]. As the Princeton group was commencing its measurements another New Jersey duo, Penzias and Wilson, had begun using a Dicke radiometer for radio astronomy. Battling an unknown noise contaminant, the pair had exhausted all avenues of reason, as they resorted to the sweeping of pigeon droppings inside their monstrous receiver horn. The Bell Labs researchers eventually made contact with the Princeton group which helped them understand their predicament. Penzias and Wilson had serendipitously discovered the cosmic microwave background, an incredibly uniform blackbody signal coming from all directions on the sky [21]. Subsequent work by physicists at Princeton helped define the results and their theoretical implications

[18, 22]. The discovery of the CMB brought the big bang universe to the forefront of modern physics and effectively obviated the steady-state hypothesis.

Figure 1.1 shows the results from the first Bell Labs and Princeton measurements overlaid on a 3 K blackbody spectrum. Despite common belief, the 1964 Bell Labs measurements did not represent the first evidence for a uniform cosmic afterglow. The study of CN molecular spectra, published as early as 1940, suggested “a maximum effective temperature of interstellar space” of about 1–3 K [23, 24] and an excess temperature of space was reported during the commissioning of the Bell Labs receiver [25, 26]. Regardless of who should be acknowledged for the initial discovery, the study of this susurrant signal continued, and we now know that the CMB is almost a perfect blackbody with temperature $T_{\text{CMB}} = 2.726$ K [27]. Its spectral radiance as a function of frequency, ν , follows the form

$$B(\nu) = \frac{2h\nu^3}{c^2} \frac{1}{\exp(h\nu/k_B T) - 1}, \quad (1.2)$$

where h and k_B are Planck and Boltzmann constants respectively, T is the blackbody temperature, and c is the speed of light in vacuum.

During these early days of cosmology, indirect evidence for the existence of a dark matter component had begun to emerge. Dark matter is hypothesized matter which interacts gravitationally with normal matter, but does not absorb or emit electromagnetic radiation. The first reference to dark matter was made by astronomer Fritz Zwicky, who observed that galaxies in the Coma Cluster have peculiar velocities which are inconsistent with velocities predicted by the virial theorem, given the estimated mass of visible matter [28, 29]. Observations of spiral galaxy rotation curves in the 1970’s, notably by Vera Rubin [30, 31], showed that angular velocities outside the galactic bulges were much larger than expected, given mass estimates.

The predicament of those times is neatly summarized by the first sentence of a 1974 paper by Ostriker et al. [32]:

“There are reasons, increasing in number and quality, to believe that the masses of ordinary galaxies may have been underestimated by a factor of 10 or more.”

More recent measurements suggest the ratio of dark matter to ordinary matter is about six to one. Unlike ordinary matter, dark matter is believed to reside in halos that are often concentric with galactic centers.³ The presence of dark matter is seen not only in the velocity profiles of galaxies but also by mapping the peculiar motion of clustering galaxies. By mapping the phase space of galaxies we learn about the growth of structure in the universe and the governing dynamics. Perhaps the most beguiling evidence for dark matter is found in the beautiful composite images of the bullet cluster [33].

The popular literature contains a wealth of publications with brilliant descriptions of the early days of cosmology [14, 34, 35, 36]. As astronomical observatories accumulated data, it became clear that the universe contained a large number of galaxies similar to our own. So far, nothing suggests our own galaxy is much different from the estimated hundreds of billions of galaxies in the observable universe. Similarly, our location in the Milky Way, our own galaxy, seems arbitrary. This gives some credence to the Copernican principle.

The Copernican principle states that we, as observers of the universe, do not have the benefit of a privileged vantage point. As modern sky surveys suggest that our local universe is isotropic and homogenous over the largest scales we conclude that so is the universe as a whole. Upon closer inspection we notice, however, that galaxies tend to clump together in halos with great lifeless voids in between. Figure 1.2 shows the distribution of approximately a hundred thousand galaxies across two patches on the sky as measured by the 2dF Galaxy Redshift Survey [37]. Similar distributions can only be seen in simulations that include a dark matter component that dominates baryonic energy densities at the ratio of six to one [38, 39].

Penzias’ and Wilson’s discovery of the cosmic microwave background, which later won them a Nobel prize, seeded the main bough of observational cosmology. Numerous experimental

³Dark matter halos are normally assumed to be spherical, not disk shaped like the word might suggest.

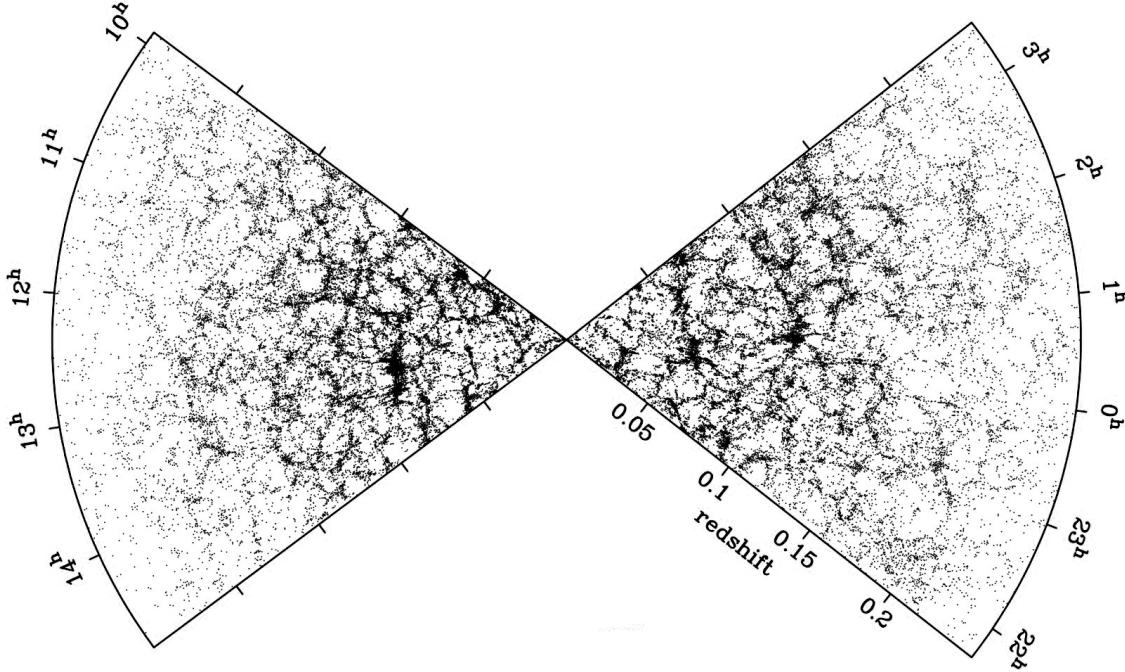


Figure 1.2: Results from the 2dF Galaxy Redshift Survey showing the distribution of galaxies projected onto the plane. Reproduced with permission from the 2dF Galaxy Redshift Survey Team [37, 40].

endeavors ensued, with efforts attempting to measure the uniformity [41] or spectral shape [42] of this fossil signal. Arguably the most famous of these seedling experiments is the Cosmic Background Explorer (*COBE*) satellite. The satellite was launched in 1989 with three instruments designed to measure different properties of the cosmic microwave background.

Using only 9 minutes of data spanning wavelengths of 1 cm to 0.5 mm, the FIRAS instrument measured a background radiation which was fit well by a 2.7 K blackbody spectrum [43]. During a 1992 meeting of the American Physical Society, measurements of CMB anisotropies were revealed, causing much stir in the scientific community. Publications of the main results followed in the *Astrophysical Journal* [44, 45]. After four years of observation, the coarse resolution DMR instrument had constructed a full sky image of the cosmic microwave background with fluctuation in the temperature of about ten parts per million.⁴ This was, and continues to be, the strongest argument for isotropy and the

⁴Only the CMB dipole anisotropy had been measured before the DMR results.

Copernican principle. At this point, observational cosmology had been established as an avenue for answering fundamental questions about the nature of the universe. The hot big bang scenario was no longer contested.

1.1 Modern Cosmology Theory – μ Review

Modern Cosmology Theory is hardly summarized in a few page document. The end of this section references more extensive review articles.

Using the aforementioned evidence and the theoretical foundation laid by Einstein's general relativity [46] we build the most general metric consistent with an isotropic expanding universe:

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

with $k \in \{-1, 0, 1\}$ for negative, zero, and positive curvature.⁵ This is called the Friedmann-Lemaître-Robertson-Walker metric (FLRW for short) after the people that first studied it in the early 20th century [47]. The time-dependent parameter a is referred to as the scale factor. It describes the expansion of the universe.⁶ The redshift, z , of an object that is receding from the observer due to the expansion of space can be related to the scale factor according to

$$1 + z = \frac{a_0}{a(t)}, \quad (1.4)$$

where a_0 is the scale factor at the current epoch, and $a(t)$ is the scale factor at an earlier time, when the photons were emitted. Using Einstein's field equations, which connect the geometrical properties of spacetime to its energy content, and the FLRW metric one can derive the two independent Friedmann equations

$$\left(\frac{\dot{a}}{a} \right)^2 = \frac{8\pi G}{3}\rho - \frac{k}{a^2}, \quad (1.5)$$

⁵Note that any other values for k can be absorbed by rescaling a and r .

⁶We set $c = 1$ in our definition of the FLRW metric, more specifically, $ds^2 = -c^2dt^2 \dots$

and

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p), \quad (1.6)$$

where ρ is energy density, p represents pressure, and G is Newton's gravitational constant. The equations describe interactions between a homogenous fluid and spacetime curvature in an FLRW universe and as such are applicable to the evolution of our universe on the largest scales. In this framework, photons, matter, and dark energy, established by observations as the key ingredients of our universe, develop differently as time progresses. Writing the equation of state for a perfect fluid as $p = w\rho$, we derive an equation for the scale factor in a flat universe:

$$a(t) = a_0 t^{\frac{2}{3(w+1)}}. \quad (1.7)$$

An ideal photon fluid is described by $w = 1/3$, whereas matter follows $w = 0$. The scale factor will therefore grow more rapidly in a matter dominated universe. From this one concludes that, being a mixture of photons, baryonic and dark matter, the universe will undergo phase transitions as different energy forms govern the development of the scale factor.

Current evidence, including observations of the cosmic microwave background, suggest that the energy budget of the universe is very close to the critical density required for a net zero curvature [48, 49], a flat universe. This implies $k = 0$, and we see from the Friedmann equations that the critical energy density required for a flat universe is

$$\rho_c = \frac{3H^2}{8\pi G}, \quad (1.8)$$

where $H \equiv \dot{a}/a$. Current best estimates of H yield a critical density of $\rho_c \approx 1 \times 10^{-29} \text{ g/cm}^3$, approximately one proton per cubic meter. Common interpretation of the Friedmann equations was that a universe with critical energy density would be perfectly balanced at the transition between a recollapsing universe and one that would expand forever. At the turn of the century, early observations of the cosmic microwave background combined with

theoretical priors seemed to favor a flat universe,⁷ yet matter inventories could only account for about a fifth of the critical density. Meanwhile, estimates for the age of the universe fell short of the age of some globular clusters [50]. Then, measurements of Type Ia supernovae came into view, suggesting that the universe was accelerating its expansion rate, as if some mysterious force was pulling it apart with an ever increasing rate [51, 52]. The scientific community must have been collectively scratching its head after this remarkable turn of events. Shortly thereafter, measurements of the cosmic microwave background indicated a spatially flat universe undergoing accelerated expansion [48], while measurements of baryon acoustic oscillations provided independent evidence for this cosmic picture [53, 54]. The phenomenon driving this acceleration has been given the name dark energy – a proper name, since it encompasses several theoretical possibilities. Dark energy represents a persistent cosmological component, which unlike baryonic and dark matter, is not watered down by the continuous expansion of space. The parameter Λ in Einstein’s field equations, originally used to balance forces such that a static universe could be realized, is now used to represent this vacuum energy density, and is colloquially referred to as the cosmological constant. The equation of state for dark energy is proposed as $p = -\rho$ which leads to $\dot{\rho} = 0$, or $\rho = \rho_0$. From this we surmise that

$$a(t) = e^{Ht}. \quad (1.9)$$

This implies an accelerated expansion of the universe, $\ddot{a} > 0$, a condition which might also have characterized another period in the history of the universe. First, it is worth emphasizing some of the problems facing the standard cosmological models.

Observations of the CMB reveal an incredibly uniform sky, yet, due to the finite speed of light two regions on opposite poles of the sky cannot have been causally connected at the time when the CMB photons were emitted. This is referred to as the horizon problem. This spectacular color coordination seems highly unlikely, unless these two regions were in causal contact for the time duration required to achieve thermal equilibrium.

⁷COBE DMR showed that the CMB is very uniform, and the inflationary paradigm predicts a flat universe.

Another issue, a type of fine-tuning problem, can be discerned from rearranging terms in the first Friedmann equation and defining $\Omega = \rho/\rho_c$:

$$\rho a^2(1 - \Omega^{-1}) = \frac{3k}{8\pi G}. \quad (1.10)$$

The right hand side of this equation is a constant, the left hand side must therefore also remain constant as the universe evolves. In the time between the infant universe and the present, the term ρa^2 must decrease by many orders of magnitude. As current measurements suggest $\Omega \approx 1.0$ the above equation would indicate careful tuning of the density parameter. This is the flatness problem. These deficiencies of the big bang model, along with others, motivated the development of the cosmic inflation paradigm.

The idea behind cosmic inflation was conceived by Alan Guth while studying the magnetic monopole problem, yet another complication involving conditions of the early universe. The creation of magnetic monopoles is expected at the high energy densities characteristic of an infant universe. Guth found that a scenario involving rapid expansion of the universe, sourced by a scalar field, could reduce the density of magnetic monopoles to present-day densities; magnetic monopoles have so far eluded detection. Soon thereafter Guth realized that inflation would also resolve the horizon and flatness problems [34, 55]. The theory of inflation quickly gained momentum, with solutions to outstanding problems proposed by Albrecht, Linde, Steinhardt, and others [56, 57, 58].

Scalar fields are widely used in physics. A prominent example is the hypothesized Higgs field whose existence was arguably confirmed with the discovery of the Higgs boson in 2012 [59, 60]. Scalar fields are invariant under translation, yet they can possess both kinetic and potential energy. The pressure and energy density associated with a scalar field are denoted as

$$\rho_\phi = \frac{1}{2}\dot{\phi}^2 + V(\phi), \quad (1.11)$$

$$p_\phi = \frac{1}{2}\dot{\phi}^2 - V(\phi). \quad (1.12)$$

In the context of cosmology, the field ϕ is normally referred to as the inflaton. From Equation 1.10 we see that Ω is forced towards unity as long as ρa^2 grows sufficiently. Defining inflation as a period of accelerated expansion, $\ddot{a} > 0$, in an era when the scalar energy density dominates all other contributions, we can use the Friedmann equations to rephrase the condition as

$$\dot{\phi}^2 < V(\phi). \quad (1.13)$$

Accelerated expansion requires that the scalar potential energy dominates the kinetic energy. Such conditions can arise if the potential is very flat; a scalar field would then roll slowly down the potential. In some sense, inflation represents a dynamical cosmological constant. During inflation, regions that were disconnected become causally connected while Ω tends to unity and exotic particles, such as magnetic monopoles, are redshifted to very low densities. Inflation suppresses some of the complications associated with the hot big bang model. It must conclude, however, with the conversion of the inflaton energy density into normal matter during a period called reheating.

Slow-roll inflationary models make simplifying assumptions about the dynamics of a single scalar field. In turn, these allow one to succinctly capture the requirements for inflation and parametrize the corresponding predictions for observables. The following discussion relies heavily on the wonderfully written textbook by Lyth and Liddle [61]. First we use the continuity equation for a perfect fluid:

$$\dot{\rho} + 3\frac{\dot{a}}{a}(\rho + p) = 0, \quad (1.14)$$

from which the second Friedmann equation is derived. This allows us to define the equations of motion for an inflaton field:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3} \left[V(\phi) + \frac{1}{2}\dot{\phi}^2\right], \quad (1.15)$$

$$\ddot{\phi} + 3\frac{\dot{a}}{a}\dot{\phi} = -V'(\phi). \quad (1.16)$$

Under the slow-roll assumption, it is assumed that terms can be dropped from the equations to arrive at a much simpler set of equations:

$$H^2 \simeq \frac{1}{3m_P^2} V(\phi), \quad (1.17)$$

$$3H\dot{\phi} \simeq -V'(\phi), \quad (1.18)$$

where we have used $H = \dot{a}/a$ and $m_P^2 \equiv (8\pi G)^{-1}$ to present more than one parametric convention. It is useful to define the slow-roll parameters,

$$\epsilon = \frac{m_P^2}{2} \frac{V'}{V}, \quad (1.19)$$

$$\eta = m_P^2 \frac{V''}{V}, \quad (1.20)$$

where V now means $V(\phi)$. Within the slow-roll approximation, it can be seen that $\epsilon \ll 1$ and $|\eta| \ll 1$ are necessary conditions for inflation.⁸ The slow-roll parameters encapsulate the simplest set of models that produce smoothing, flattening, and monopole dilution, the basic requirements of the inflationary paradigm. As it turns out, the slow-roll recipe produces a little too much smoothing. The final ingredient draws from quantum mechanics.

Even though the average density of our universe corresponds to a proton per cubic meter, this seemingly lifeless universe was able to facilitate the growth of galaxies, solar systems, and planets – structure as it is known to us. To explain this we invoke primordial density perturbations as quantum fluctuations of the inflaton, or more generally as perturbations to the FLRW metric. Unlike the field itself, the quantum fluctuations are not translation invariant. These perturbations, seeded during inflation, survive the transition from this speculative embryonic era to a largely uncontested hot and dense infant universe developing according to established laws of physics. The perturbations are therefore imprinted into the spectrum of the cosmic microwave background, and with proper care, can shed light on

⁸For example, take the time derivative of Equation 1.15, plug the result into Equation 1.16, and use the result to express V'/V .

inflation. In short, quantum fluctuations are stretched to cosmic length scales during rapid expansion in the inflationary era. However, in order to extract that information we must first understand how perturbed energy densities develop in the hot big bang scenario.

Gravitational attractions in the photon-baryon plasma tend to form halos with infalling matter while increasing photon pressure impedes this process and erases anisotropies. Different modes of compression and rarefaction develop at the speed of sound in the plasma [62]. This tug of war, referred to as acoustic oscillations, continues until the universe has expanded enough, and therefore decreased in temperature, for protons and electrons to form hydrogen. This is called recombination and happens at a redshift of $z = 1090$, approximately 380,000 years after the big bang [63]. The event defines a veil at the edge of our horizon, sometimes referred to as the last scattering surface. The universe now becomes comparatively transparent to light and the photons stream freely in every direction with a fraction bombarding our detectors today. These are the CMB photons, and they carry with them information about the fundamental oscillation modes at the time of recombination. By studying the properties of the CMB, we can examine the matter composition of the universe, investigate the nature of dark energy, and probe energy densities that are far beyond that of any terrestrial experiments.

The universe became electrically neutral at the time of recombination. However, we observe that the hydrogen in the interstellar medium is now largely ionized. This matter reionization occurred at around redshift of $z \approx 10$, and is thought to have been sourced by ultraviolet radiation from the first luminous objects.

CMB temperature anisotropies, $T(\hat{\mathbf{n}}) = \delta T(\hat{\mathbf{n}})/T_0$, are naturally decomposed using spherical harmonics according to

$$T(\hat{\mathbf{n}}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \phi), \quad (1.21)$$

where

$$a_{\ell m}^T = \int d\Omega T(\theta, \phi) Y_{\ell m}(\theta, \phi) \quad (1.22)$$

represent the expansion coefficients of a spherical harmonic decomposition. Assuming rotational invariance, the following relation must hold

$$\langle a_{\ell m}^T a_{\ell' m'}^{T*} \rangle = \delta_{\ell \ell'} \delta_{m m'} C_{\ell}^{TT} \quad (1.23)$$

where $\langle \cdot \rangle$ implies an average over the statistical ensemble defined by infinite sky realizations drawn from the same underlying theory. The angular power spectrum, C_{ℓ}^{TT} , then represents the variance in power in a given ℓ -mode. It follows that

$$\langle T(\hat{\mathbf{n}}) T(\hat{\mathbf{n}}') \rangle = \frac{1}{4\pi} \sum_{\ell=1}^{\infty} (2\ell + 1) C_{\ell}^{TT} P_{\ell}(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}'), \quad (1.24)$$

where $P_{\ell}(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}')$ is a Legendre polynomial of order ℓ and the following mathematical identity has been used:

$$P_{\ell}(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}') = \frac{4\pi}{2\ell + 1} \sum_{m=-\ell}^{\ell} Y_{\ell m}(\hat{\mathbf{n}}) Y_{\ell m}(\hat{\mathbf{n}}'). \quad (1.25)$$

The angular power spectrum has been measured by numerous experiments (see Figure 1.3). Recently, the *Planck* collaboration presented its first cosmological results [49, 64, 65, 66, 67]. The *Planck* survey covered the full sky and consequently the *Planck* derived TT power spectrum estimate spans a remarkably wide ℓ -range, corresponding to angular scales of 180 degrees down to approximately 3 arcmin. In Figure 1.3 the acoustic oscillations can be seen as a series of peaks and troughs starting at degree angular scales, coinciding with $\ell \approx 100$. Normally the power spectra are plotted as a function of ℓ , the multipole moment. The conversion to corresponding angular scales is found by the approximate expression $\theta \approx \pi/\ell$ [rad]. The CMB temperature anisotropies represent variations in the intensity of light from the last scattering surface. It turns out that this light can also be polarized.

The polarization of the CMB is caused by Thomson scattering from quadrupolar density anisotropies at the epoch of photon decoupling. Polarization is only generated by quadrupolar anisotropies as other temperature distributions, such as a dipole, do not share a common

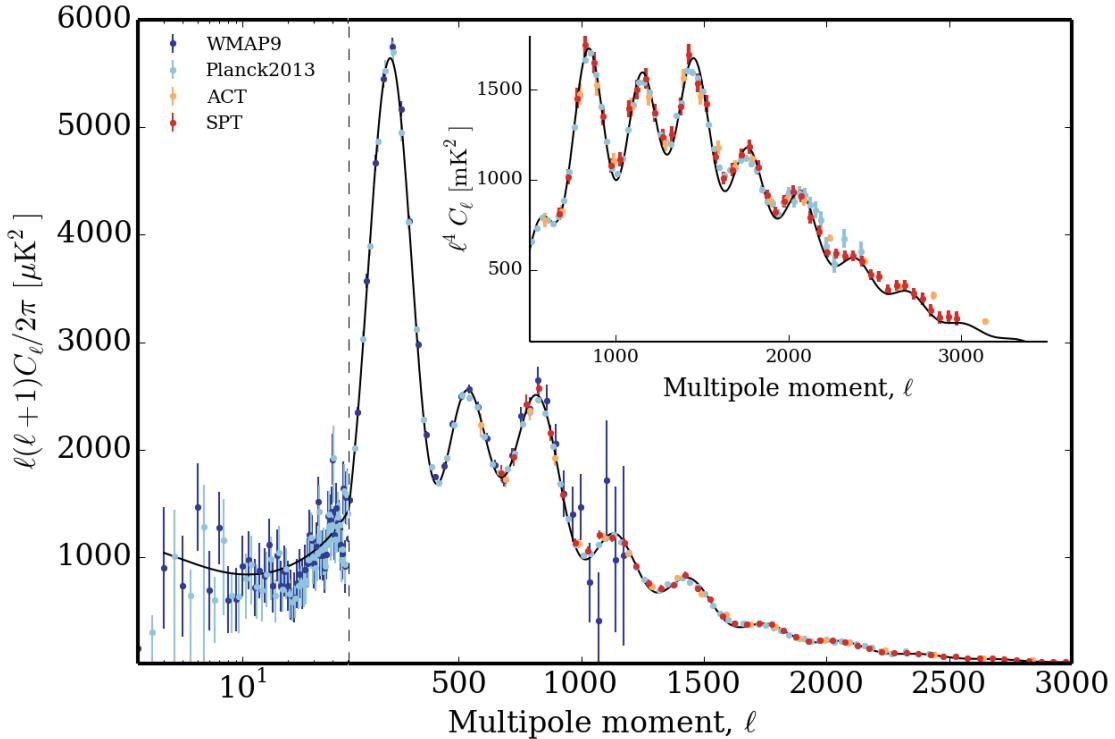


Figure 1.3: The temperature power spectrum from a few prominent experiments. The solid line shows the best fit power spectrum derived from *Planck* data in conjunction with polarization data from WMAP, high- ℓ experiments, and experiments measuring baryon acoustic oscillation [49]. *Inset*: The power spectrum plotted using a different scaling on the y-axis to highlight its oscillatory nature and the available measurements. The leftmost peak at $\ell \sim 800$ is the third acoustic peak. Data obtained from the Legacy Archive for Microwave Background Analysis (LAMBDA) [68].

plane that is perpendicular to the direction of propagation (see Figure 1.4). These anisotropies can be generated by perturbing the FLRW metric that describes our cosmology and these perturbations can in turn be decomposed into scalar and tensor components which have different properties. Scalar perturbations give rise to density waves that source the bulk of the temperature and polarization anisotropies. The density waves are rotationally symmetric about the direction of propagation, resulting in a polarization signal that is curl-free, hence the name *E*-mode polarization. *B*-modes represent the geometrical negative of these shapes, they appear to swirl around a fixed point on the sky. Figure 1.4 shows the geometrical shape of these patterns on the sky.

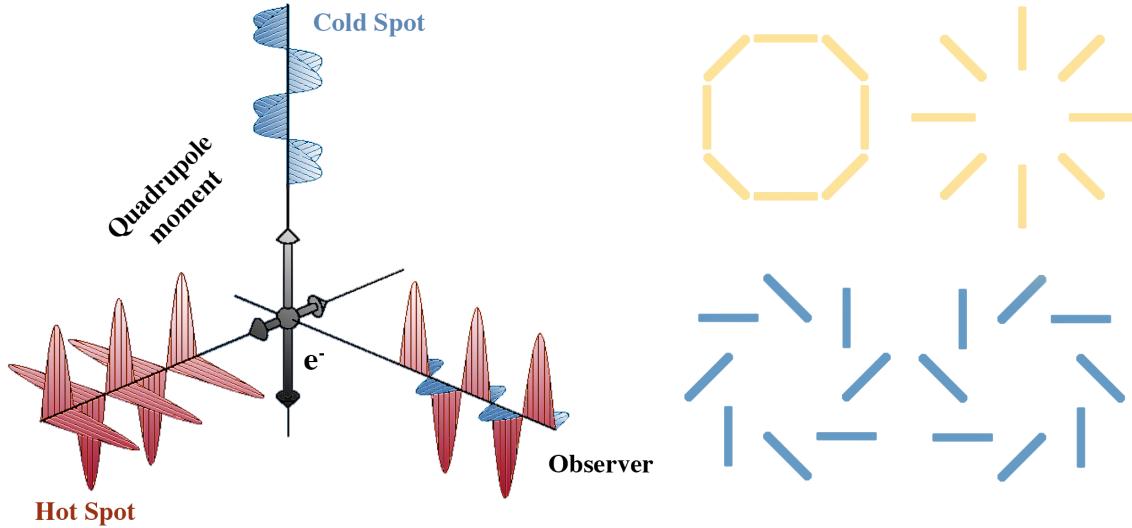


Figure 1.4: *Left:* Quadrupolar temperature anisotropies generate asymmetric oscillation patterns in electrons, which in turn radiate polarized light towards an observer perpendicular to the plane of the anisotropy. Figure presented with permission from [74] *Right:* The *E*- and *B*-mode polarization patterns as viewed on the sky. *E*-mode patterns (yellow) are parallel or perpendicular to a source while *B*-mode patterns (blue) have a specific handedness.

Polarization anisotropies generated by tensor perturbations, which are synonymous with gravitational waves, do not have this symmetry and *E*- and *B*-modes are sourced at roughly equal levels [69]. It is customary to decompose the spectrum of these perturbations in Fourier space using comoving wave-vectors. We use the parameters $\mathcal{P}_S(k)$ and $\mathcal{P}_T(k)$ to describe the power spectrum of scalar and tensor perturbations respectively. By assuming the perturbations only depend on the amplitude of the wave-vector, k , we are imposing isotropy. Characterization of the *E*- and *B*-mode power spectrum is ongoing [49, 70] with a number of CMB experiments currently being developed and fielded to search for the *B*-mode signal [71, 72, 73]. These experiments are designed to push the current *B*-mode sensitivity down by an order of magnitude.

The intensity and polarization of the cosmic microwave background at any point on the

sky are described by the Stokes parameters I , Q , U , and V . They are defined as

$$\begin{aligned} I &= \langle |\mathbf{E}_x|^2 \rangle + \langle |\mathbf{E}_y|^2 \rangle, \\ Q &= \langle |\mathbf{E}_x|^2 \rangle - \langle |\mathbf{E}_y|^2 \rangle, \\ U &= 2\text{Re} \langle \mathbf{E}_x \mathbf{E}_y^* \rangle, \\ V &= 2\text{Im} \langle \mathbf{E}_x \mathbf{E}_y^* \rangle, \end{aligned} \quad (1.26)$$

with E_x and E_y corresponding to the electric fields measured in the tangent plane, defined by a fixed orthogonal coordinate system with unit vectors $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$, and $\langle \cdot \rangle$ now corresponding to the time-averaged field. Whereas the temperature field is invariant under rotation around the unit direction vector, $\hat{\mathbf{n}}$, the Q , U , and V Stokes parameters are not. We can, however, form a linear combination of the Q and U Stokes parameters that transforms under rotation according to

$$Q(\hat{\mathbf{n}}) \pm iU(\hat{\mathbf{n}}) = e^{2i\Delta\psi}(\tilde{Q}(\hat{\mathbf{n}}) \pm i\tilde{U}(\hat{\mathbf{n}})). \quad (1.27)$$

This is a spin-2 function (see [75] for definition) and there exists a set of function bases known as the spin-s spherical harmonics that are appropriate for expanding these functions on the sphere. The expansion of $(Q \pm iU)$ using spin-2 spherical harmonics is similar to that of the scalar quantity $T(\hat{\mathbf{n}})$, with the expansion coefficients given by

$${}_{\pm 2}a_{\ell m} = \int d\Omega (Q \pm iU)_{\pm 2}Y_{\ell m}(\theta, \phi). \quad (1.28)$$

Here ${}_{\pm 2}Y_{\ell m}$ represent the spin-2 spherical harmonic functions. The E - and B -mode coefficients are then defined by the spin-2 coefficients according to

$$\begin{aligned} a_{\ell m}^E &= -({}_2a_{\ell m} + {}_{-2}a_{\ell m})/2, \\ a_{\ell m}^B &= i({}_2a_{\ell m} - {}_{-2}a_{\ell m})/2, \end{aligned} \quad (1.29)$$

with the corresponding power spectra of these scalar quantities as

$$\begin{aligned} C_\ell^{EE} &= \langle a_{\ell m}^E a_{\ell m}^{E*} \rangle = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m}^E a_{\ell m}^{E*}, \\ C_\ell^{BB} &= \langle a_{\ell m}^B a_{\ell m}^{B*} \rangle = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m}^B a_{\ell m}^{B*}. \end{aligned} \quad (1.30)$$

The CMB temperature and polarization power spectra encapsulate information about eras leading up to the epoch of photon decoupling. Up until this point, all data indicate that they are consistent with Gaussian and adiabatic perturbations, suggesting that the CMB power was imprinted while linear perturbations were still valid.

Using Newtonian physics one can show that the temperature anisotropies generated by perturbations to a homogeneous matter-dominated universe are scale invariant, with C_ℓ scaling like $1/\ell(\ell + 1)$ in the power domain [76]. This is the reason the power spectra are normally plotted multiplied with the $\ell(\ell + 1)$ prefactor. Scale invariance, however, does not imply a featureless spectrum and many features in the power spectrum represent birthmarks of our universe. On the largest scales, the Sachs-Wolfe effect, whereby gravitational effects modify the spectrum, dominates the anisotropies [77]. At degree angular scales the anisotropies are governed by oscillations in the densities of primordial plasma with the first peak of the power spectrum corresponding to perturbation modes that had entered the Hubble horizon and had time to develop overdensities. The location of the first peak therefore tells us about the geometrical properties of the universe. No causal physics occurring at or before recombination could imprint larger scales. These density perturbations were characterized by the propagation speed of sound in the primordial plasma, hence the term acoustic oscillations. Subsequent oscillations correspond to harmonics of these acoustic oscillations which are eventually suppressed on smaller scales due to Silk damping, photon diffusion during recombination [78].

The temperature power spectrum has been characterized from the largest scales down to $\ell = 10,000$ corresponding to about 1 arcmin angular scales. For $\ell > 3000$ the cosmic infrared background appears to dominate the power budget, with contributions from radio sources as

well as the thermal and kinematic Sunyaev-Zeldovich effect [79]. These are often referred to as secondary anisotropies. The secondary anisotropies offer a rare view of reionization and the growth of structure.

With results from the *Planck* satellite arguably extending the cosmic variance limit out to $\ell \gtrsim 1500$, the primordial temperature anisotropies have been somewhat rigorously quantified. Conversely, the polarized power spectra require much work. Figure 1.5 shows the *E*- and *B*-mode power spectra measurements and the best fit cosmological model (Λ CDM) based on *Planck* measurements. Most constituents of the inflationary paradigm predict a set of observable features in the CMB power spectrum, including a total energy close to the critical value ρ_c , a nearly scale invariant spectrum⁹, tensor perturbations, and therefore *B*-mode polarization. The first two of these have already been extensively verified. Most recently the *Planck* collaboration reported a deviation from scale invariance at the 5.4σ level [49].

Model dependent constraints on inflationary parameters are possible through measurements of the *TT* power spectrum alone [93, 94] as both tensor and scalar perturbations contribute to the *TT* power. However, different inflationary models predict varying amplitude of the spectral tilt and its derivatives. Conversely, inflationary models almost uniformly predict tensor perturbations at some level, which are manifest in a primordial *BB* power spectrum. Unlike others, the *B*-mode observables are generally not degenerate with a range of other cosmological parameters.

Unfortunately, a number of sources produce *B*-mode polarization and a positive detection does not mean that these were imprinted on the CMB at the last scattering surface. The most prominent sources include polarized foregrounds, both synchrotron and dust emission within our own Galaxy, as well as matter anisotropies within the last scattering surface that gravitationally lens some of the *E*-modes into *B*-modes. Galactic foregrounds tend to produce large-scale anisotropies ($\ell \ll 100$) and fall with ℓ whereas lensed *B*-modes do just the

⁹Models for slow-roll inflation predict small deviations from the Harrison-Zeldovich spectrum, characterized by $n_s = 1.0$.

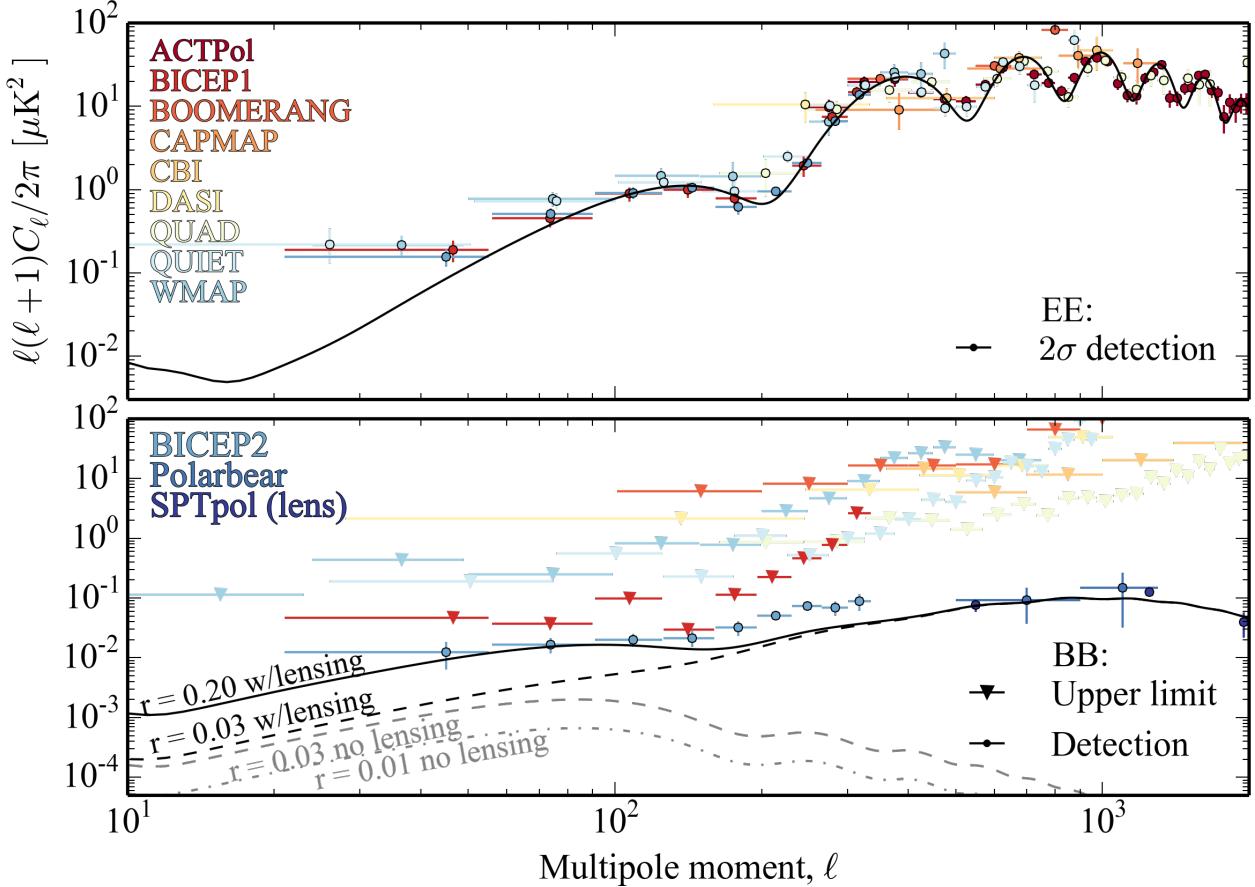


Figure 1.5: A combination of published *EE* and *BB* polarization power spectra from ACTPol [80], QUAD [81], *WMAP* 9 year release [82, 83], CAPMAP [84], CBI [85], DASI [86], BICEP1 [87], BICEP2 [88], BOOMERanG [89], QUIET [90], and SPTpol [91]. For the *EE* spectrum only 2σ detections are shown to increase clarity. The best fit Λ CDM model based on *Planck* parameters is shown in black [49]. It is possible that primordial *B*-mode power have been detected [88], and lensing of *E*-mode power into *B*-modes has been observed [92, 91]. The $r = \{0.01, 0.03, 0.20\}$ *B*-mode power spectra based on the best fit Λ CDM model are shown in black and grey.

opposite. Primordial *B*-modes are expected to peak at degree angular scales, corresponding to $\ell \sim 100$. This is where experiments looking for primordial *B*-modes focus their sensitivity.

Within the context of slow-roll inflation (see Equations 1.17 and 1.18) CMB observables can be conveniently linked to the slow-roll parameters. The RMS amplitude of scalar fluctuations has already been well constrained through measurements of the CMB; the temperature anisotropies are incredibly uniform, with only 100 part in million fluctuations. Based on this, we write $\mathcal{P}_S(k_0) \sim 100 \times 10^{-6}$, where \mathcal{P}_S is the scalar perturbation power spectrum and

k_0 corresponds to the scale of the measurement. Having established this measurement, the energy scale of inflation within the slow-roll scenario can be expressed as

$$V^{1/4} = \epsilon^{1/4} 9 \times 10^{16} \text{GeV}. \quad (1.31)$$

Similarly, the spectral index, $n_S - 1 = d \ln \mathcal{P}_S / d \ln k$, describing the scale dependence of the curvature perturbations follows

$$n_S = 2\eta - 6\epsilon. \quad (1.32)$$

A measurement of primordial B -modes constrains the amplitude of tensor perturbations, \mathcal{P}_T , and therefore the so-called tensor-to-scalar ratio, r , which in turn is related to one of the slow-roll parameters:

$$r \equiv \frac{\mathcal{P}_T(k_0)}{\mathcal{P}_S(k_0)} = 16\epsilon. \quad (1.33)$$

Finally, the spectral tilt of the tensor power spectrum can be succinctly written as $n_T = -2\epsilon$. Under the assumption that slow-roll inflation accurately captures the dynamics of the early universe, a measurement of the tensor-to-scalar ratio constrains the energy scale of inflation.

A great number of inflationary models have been proposed. Although many of these models fall under the umbrella of single field slow-roll inflation, there seem to be no limits to the complexity that inflationary models can assume. Within the framework of single field inflation, attempts have been made to gather all such models in the general description of effective field theory [95], yet some will argue that such generalizations represent addition of complexity that is not warranted. Although many claim that the observation of primordial B -mode polarization would cement inflation's role as a model of the early universe, it should be pointed out that the inflationary paradigm is by no means the only viable physical model that describes the early universe [96]. We also note that model complexity alone is not a

damning feature. Instead, models are falsified by comparing their predicted values for various observables with reality (see Section 1.2).

A plethora of textbooks and review articles describe the concepts briefly reviewed in this text. Some of the more common textbooks include [61, 97]. Review articles that describe cosmology theory in broad strokes include [98, 99, 100], while texts like [101] delve into the gory details of cosmic perturbation theory. There also exist review articles highlighting some of the differences between the various experiments involved in observational cosmology [102].

1.2 Observational Cosmology

As the first CMB measurements now approach their semicentennial anniversary, numerous experiments have contributed to our understanding of the relic radiation. The *COBE* satellite, launched in 1989, was the first of these to map out this radiation over the whole sky. More recent satellite experiments, *WMAP* and *Planck*, launched in 2001 and 2009 respectively, have made significant improvements in our understanding of the CMB and its anisotropies across the entire sky. Figure 1.6 shows a composite of the temperature anisotropies as measured by the three full-sky surveys. Note how the angular resolution improves with time.

Satellite experiments are expensive and rare, and other less costly experiments have equally furthered our understanding of the microwave sky. The BOOMERanG balloon-borne experiment, and others like it, measured the angular diameter distance to the surface of last scattering; coupled with constraints on the Hubble parameter, these measurements showed that the large-scale universe is flat [48, 103]. The DASI experiment was the first experiment to conclusively detect the polarization of the CMB [104]. Large aperture experiments, such as ACT and SPT, have identified the lensing field of large-scale structure through its effect on the CMB polarization [66, 105, 106].

At the beginning of 2014, a few experiments had published constraints on r from polarized maps of the cosmic microwave background (see Figure 1.5). The BICEP1 experiment held the upper limit from polarized data alone, $r < 0.65$ at 95% confidence, [107, 70] with

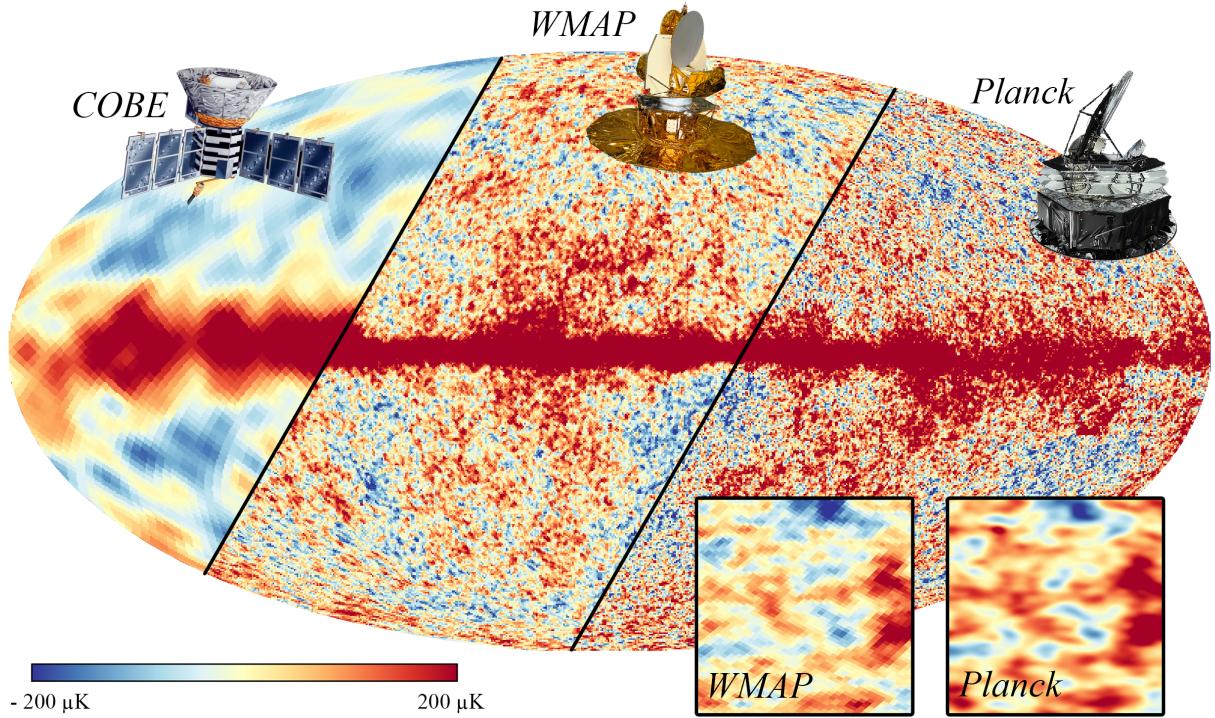


Figure 1.6: A composite Mollweide projection of the CMB anisotropies at approximately 100 GHz as observed by *COBE* DMR (1989), *WMAP* (2001), and *Planck* HFI (2009). The color scale covers $\pm 200 \mu\text{K}$ with the Galactic stripe saturating the scale. The insets in the lower right show *WMAP* and *Planck* maps over a 8 degree square patch on the southern sky. This marks the first instance of many in this document where the HEALPix package was used to analyze results [4].

competitive measurements from the QUIET experiment [71]. At that point, best limits on r were determined in a somewhat convolved manner, stemming from a combination of CMB temperature anisotropies and measurements of baryon acoustic oscillations, with the *WMAP* and *Planck* collaborations reporting $r \lesssim 0.13$ at 95% confidence [49, 82]. Those constraints are found by assuming no running of the scalar spectral index; relaxing that requirement gives $r < 0.26$.

The primordial B -mode signal is surpassed by lensing of E -modes into B -modes at small angular scales. The detection of this gravitationally lensed B -mode signal was first reported in 2013 by SPT [91] and independently measured by the POLARBEAR experiment [106].

In March of 2014 the POLARBEAR collaboration published a paper describing its

measurements of the B -mode power in an multipole range of $500 < \ell < 2100$, further constraining the impression of large-scale structure on B -mode polarization [92]. That same month, the BICEP2 collaboration announced a detection of B -mode power at degree angular scales, $30 < \ell < 150$, which appears to be consistent with $r \approx 0.20$ under the assumption of no foreground emission [88, 108].¹⁰ Many in the scientific community offered a tempered interpretation, pointing out that foreground models and their uncertainties suggest a comparable contribution from Galactic dust in this region. As the identification of a primordial cosmological B -mode signal would mark a watershed moment in the history of cosmology, the extraordinary announcement from the BICEP2 collaboration has invigorated other ground based experiments currently observing the microwave sky, including the Atacama B -mode Search (ABS) and the *Keck Array*. Concurrently, small scale experiments like SPTpol and ACTPol are characterizing the lensing B -mode signal amplitude while developing technologies and data analysis methods that allow them to push towards mass production of detectors and larger angular scales [109, 110]. The cosmological community is anxiously awaiting results on polarized foregrounds from the *Planck* HFI, planned for release at the end of 2014. These will undoubtedly help shed some light on the BICEP2 results. Finally, SPIDER is a balloon-borne experiment specifically designed to constrain the B -mode amplitude (see Chapter 2). The payload will launch from McMurdo station in December 2014.

Many experiments not mentioned here have also helped improve our understanding of the cosmic microwave background. Cosmology, however, is not derived from CMB data alone. Instead, observational cosmology is now informed by data spanning the entire electromagnetic spectrum.

Type Ia Supernovae

Much like Cepheid variables, Type Ia supernovae are astrophysical objects believed to have a fixed intrinsic brightness; both are referred to as standard candles. By measuring their

¹⁰At this amplitude, cosmic variance could shift r by ~ 0.05 .

apparent brightness we can estimate the physical separation. These catastrophic stellar events outshine entire galaxies for many days allowing scientists to map the recessional velocity of far-away galaxies. This method extends Hubble’s law out to redshift of $z \sim 1$ with current best estimates of the local value from the Hubble Space Telescope of $H_0 = 73.8 \pm 2.4 \text{ kms}^{-1}\text{Mpc}^{-1}$ [111]. Such datasets have revealed an inflation of Type Ia luminosity distance at high redshifts, suggesting accelerated expansion of the universe driven by dark energy [51, 52].

Baryon Acoustic Oscillations

Prior to recombination, interactions between radiation and matter left an imprint on the large-scale distribution of matter across the universe. The phenomenon is referred to as Baryon Acoustic Oscillations (BAOs) and its detection was first described in [53, 54]. Baryons and dark matter couple gravitationally, yet dark matter lacks photon interactions and, therefore, does not feel radiation pressure in overdense regions. Conversely, baryons are pushed outwards from initial overdensities by this photon pressure. As this happens, the universe is expanding and eventually the interaction subsides during recombination. Baryons are suddenly left behind as the CMB photons shed their subluminal counterparts. Most of the baryons fall back towards the overdense regions while a fraction remains in a shell as an echo from the former era.

By mapping the distribution of tens of thousands of galaxies over a range of angles and redshifts, scientists are able to measure this echo, the angular power spectrum of matter. The characteristic angular scale corresponds to an overdensity in galaxy correlations at 150 Mpc comoving scales. Current best limits are described in [112]. A great review article can be found in [113].

Big Bang Nucleosynthesis

In the canonical cosmological model, Big Bang Nucleosynthesis (BBN) is a process which took place in the first three minutes after the big bang. During this process, the chemical

abundances of all the lightest elements and their isotopes, such as ^2H , ^3He , ^4He , and ^7Li were generated [114]. Using spectroscopy, measurements of the chemical composition of metal-poor star forming regions are used to constrain the global composition of our universe. Although minor discrepancies exist, the overall BBN picture, as indicated by spectroscopy and the CMB, seems quite consistent [49].

Together with the CMB, Type IA supernovae, BAO, and BBN data form some of the main pillars of modern observational cosmology.

Concluding Remarks

Pioneering work published at the turn of the century proved that the density of the universe is near the critical density required for a flat universe [48]. Despite that, it would seem that only a small fraction of the universal energy density is provided by ordinary matter. A few years earlier, and somewhat contradictory to the statement above, evidence of an expanding universe emerged [51]. In 2013, the *Planck* satellite revealed maps of energy density variations as probed by CMB backlight. Finally, this year, the potential discovery of primordial *B*-mode polarization was announced [88]. The field of observational cosmology is clearly flourishing.

1.3 Thesis Work and Content

Throughout my graduate student career my primary responsibility has lain with the SPIDER experiment and the integration of its flight cryostat with other scientific hardware. A significant fraction of my time has been spent on designing, building, and characterizing sub-systems of this scientific instrument, many of which operate at cryogenic temperatures. During the commissioning phase of the flight cryostat, I worked on a thermal model that has been used to predict its thermal performance while also helping with the cryostat build process. One of my larger responsibilities included the design and build of two Fourier Transform Spectrometers that are used to characterize the spectral response of the SPIDER

telescopes. I also designed and built a capillary assembly that provides a steady supply of superfluid helium to a 1.8 K temperature stage in the SPIDER flight cryostat.

In addition, I worked on the in-flight characterization of the *Planck* satellite where my largest contribution was in the development of a simulation pipeline to characterize the fidelity with which we can reconstruct the spatial response of the instrument. I developed algorithms that merge multiple planet observations into a single beam map and probed the fidelity of that process with simulations run on the local Feynman computer cluster. Variants of these algorithms were incorporated into the official *Planck* analysis pipeline. Other work includes validation of the pointing solutions and photometric calibrations from measurement of planet flux densities.

The remainder of the thesis is structured as follows: Chapter 2 gives a general overview of SPIDER, a balloon-borne polarimeter designed to constrain the *B*-mode power of the CMB on degree angular scales over approximately 8% of the sky. The more detailed sections of that chapter describe work where I have made significant contributions. Other sections are presented in order to give a complete description of the experiment. Chapter 3 reports work on the calibration of the *Planck* High Frequency Instrument (HFI). Among other topics, that chapter describes a simulation pipeline which was constructed to constrain the error in the reconstruction of the *Planck* HFI beam response. Chapter 4 discusses calibrations and simulations of the optical response of the SPIDER telescopes building on concepts covered in Chapter 3. Finally, Chapter 5 offers brief conclusions and a discussion of the theoretical implications for a primordial *B*-mode discovery.

Chapter 2

SPIDER: A Balloon-Borne CMB Polarization Experiment

2.1 Instrument Overview

SPIDER is a Long Duration Balloon-borne (LDB) experiment designed to measure the polarization of the cosmic microwave background with unparalleled instantaneous sensitivity. Novel detector architecture allows for illumination of approximately 2500 detectors through an effective collecting area of roughly 0.5 m^2 . SPIDER will map the polarization of the CMB over a tenth of the sky during an estimated twenty-day Antarctic balloon flight. Table 2.1 highlights some of the main characteristics of the experiment. The payload was integrated and proven flight ready in the summer of 2013 in Palestine, Texas. The United States Federal Government furlough in October of 2013 inhibited ballooning missions from McMurdo Station that year, delayed the launch of the experiment by a year, and cast a pall over the entire project. Apart from political vicissitudes, the project members have had to deal with fires, strikes, dust storms, flooding, hurricanes, and countless power outages. The experiment, in both design and instrumentation phase has been described in various journal articles and conference proceedings [89, 115, 116].

Figure 2.1 shows the fully integrated payload hanging from a launch vehicle during compatibility testing. Six monochromatic telescopes observe the sky through the 3 mm thick plastic windows at the top of the cryostat.¹ Each telescope houses hundreds of polarization sensitive detectors that are cryogenically cooled to 300 mK. The total mass of the payload as seen in Figure 2.1, including cryogens, five telescopes, and all power and telemetry systems, was measured to be 2450 kg.²

Signal from these detectors is sampled sequentially using time-domain multiplexing with Multi-Channel Electronics (MCE) clocked at 50 MHz [117, 118]. Each telescope is read out by an independent MCE system which is then connected to an MCE Computer (MCC) through an optical fiber. Accounting for settling times, decimation rates, filtering, and so forth, the readout system will sample detectors at approximately 120 Hz. The MCCs write these data to a redundant file system consisting of both solid state and spinning disks housed in pressure vessels. Two computers, named *itsy* and *bitsy*, form the heart of the SPIDER flight computer system. A watchdog monitors the status of these and switches priority if a computer malfunctions. The flight computers facilitate communications with the MCCs, the attitude control system (ACS), and the Support Instrumentation Package (SIP) supplied by the Columbia Scientific Ballooning Facility (CSBF). Custom housekeeping and motor control electronics form an intermediary between the flight computers and the various electrical readout devices onboard the payload. All communications with the payload go through the ground station computer which routes signal to the flight computers through satellite or a line-of-sight connection [119, 120].

2.1.1 Foregrounds and Noise

SPIDER will observe a 4000 deg^2 window on the southern sky where Galactic foreground contamination is expected to be minimal. Early estimates of a twenty-day flight suggested that a SPIDER-like platform could measure $r = 0.04$ with 99% confidence in the absence

¹The windows are made of Ultra High Molecular Weight Polyethylene (UHMWPE).

²The mass of your average Ford E-Series with a couple of people and a tool chest.

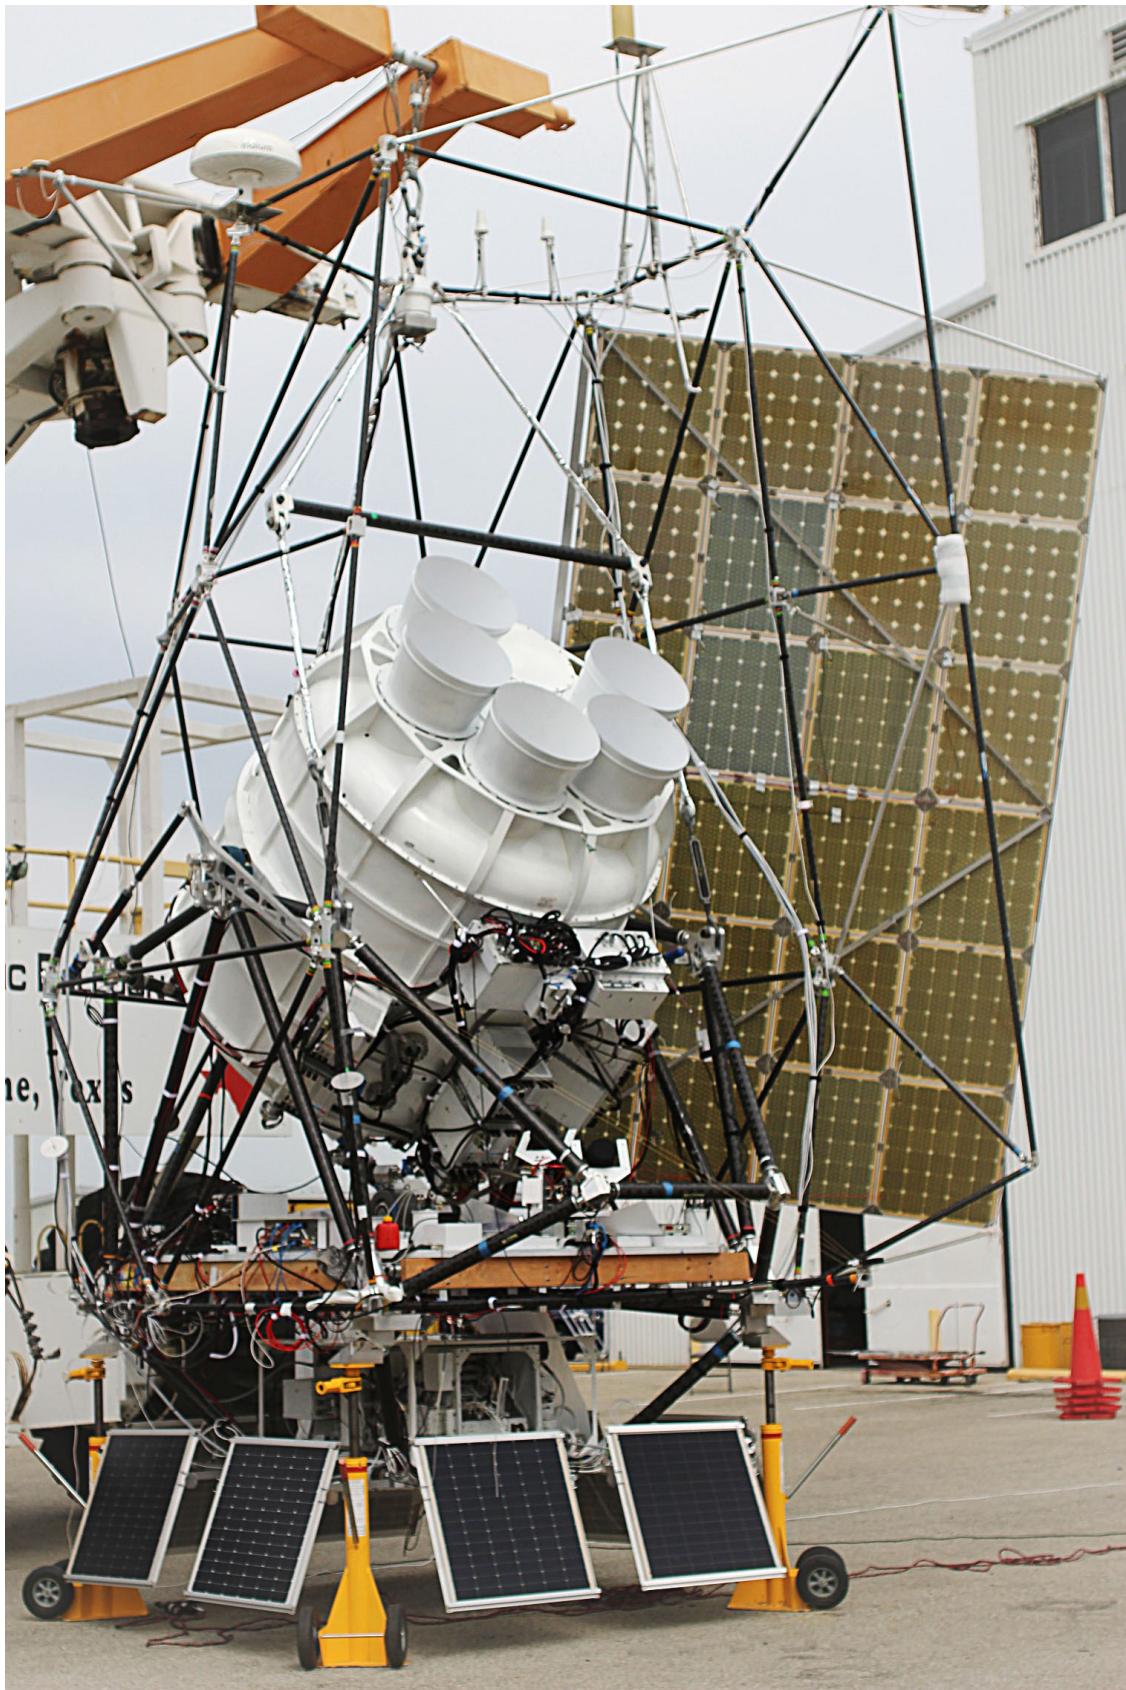


Figure 2.1: The SPIDER payload during compatibility testing in Palestine, Texas, Summer 2013. Courtesy of Steven Benton.

Table 2.1: Primary characteristics of the SPIDER experiment.

Sky coverage	8%
Scan rate (az)	6 deg/s at peak
Polarization modulation	Stepped HWP
Detector type	Antenna-coupled TES
Multipole range	$10 < \ell < 300$
Integration time*	17 days
Limits on r^\dagger	0.03

* Assuming 85% duty cycle.

† Assuming no foregrounds, at 99% confidence.

	Frequency [GHz]	
	94	150
Number of telescopes	3	3
Bandwidth [GHz]	22	36
Optical efficiency	30–45%	30–50%
Angular resolution* [arcmin]	42	30
Number of detectors†	690	1230
Internal loading ^{◇,b} [pW]	1.51	2.25
CMB loading ^b [pW]	0.35	0.30
Atmosphere loading ^b [pW]	0.09	0.19
NET per detector [$\mu\text{K}\sqrt{\text{s}}$]	120–150	110–150

*FWHM. † Assuming 80% yield. [◇] Including sleeve, window, and baffle.

^b Not accounting for optical efficiency.

of foregrounds [89]. Such statements were incomplete, as interstellar dust was expected to dominate an $r = 0.03$ primordial B -mode signal at the largest angular scales [121, 122, 123]. More recent work [124] showed that two flights of a SPIDER-like payload, utilizing three frequency bands, could constrain the amplitude of the B -mode signal to $r < 0.03$ even in the presence of significant foregrounds. Such statements relied on foreground models which, until very recently, have lacked data to characterize the amplitude of polarized emission from interstellar dust [125]. Emission from electrons accelerated by Galactic magnetic fields,

so-called synchrotron emission, is also expected to contribute significantly to the large scale power at 94 GHz [124] (see Figure 2.2).

SPIDER’s observation region on the southern hemisphere is expected to be relatively void of Galactic foregrounds. Assuming yield and noise characteristics described in Table 2.1, the instantaneous sensitivity at 150 GHz is expected to be $4 \mu\text{K}\sqrt{\text{s}}$, which is five times that of *Planck* HFI at 143 GHz. A twenty-day flight with an 80% duty cycle should lead to $0.2 \mu\text{K}/\text{deg}^2$ sky sensitivity, compared to approximately 2.5 and $15 \mu\text{K}/\text{deg}^2$ for the 143 GHz *Planck* HFI band and the 100 GHz *WMAP* band, respectively. SPIDER will be the most sensitive experiment to measure this region of the sky. Signal calibrations will be performed at degree angular scales using the already well characterized CMB temperature anisotropies. As the window for detecting primordial *B*-modes is centered on degree angular scales we bypass the need for extensive extrapolations of dipole calibrations performed by both *Planck* HFI and *WMAP*.

Choosing between characterization of foregrounds or the extragalactic *B*-mode signal, those working on CMB polarimeters would much rather constrain the latter. As polarized foregrounds are so poorly constrained, selecting the ideal observation region, defined by lack of foregrounds, represents something of a gamble. A tradeoff between focusing integration time over a small region on the sky or extending coverage to find the cleanest region is therefore established [126]. Unfortunately, the stochastic ensemble of dust particles in our Galaxy is overshadowed by another group of terrestrial clamor; the photon noise caused by everything between our detectors and the CMB.

CMB experiments strive to be photon noise limited, which means that the majority of a detector’s noise inventory is due to the unavoidable intensity and statistical properties of the CMB photons. This is easier said than done, as emission from the Earth’s atmosphere greatly exceeds the intensity of the CMB. The CMB is a perfect blackbody which, according to Wien’s law, peaks at around 160 GHz.³ As a result, most CMB experiments will distribute

³Note that CMB experiments are measuring dB/dT , the temperature derivative of Planck’s blackbody function. This derivative peaks at 217.5 GHz at a blackbody temperature of $T_{\text{CMB}} = 2.726 \text{ K}$.

their spectral sensitivity in the range of 50–300 GHz.⁴ Unlike the CMB, the intensity of Earth’s atmosphere is plagued by intricate time and frequency dependence, which is difficult to model [128]. This is the reason satellite missions are so desirable and the reason SPIDER will fly on a stratospheric balloon, even though it results in only 20 days of observation.

Unlike all ground-based experiments, SPIDER will operate almost in the CMB photon noise limit. At a float altitude of 37 km and an elevation angle of 40 deg, the loading from residual atmosphere is similar to that of a 1–2 K blackbody, compared to the 2.7 K temperature of the CMB. Unfortunately, in the case of SPIDER, the CMB photon noise equivalent power is surpassed by a factor of two by internal loading from the warm vacuum window which is required to prevent cryogenic loading from residual atmosphere. The SPIDER detectors also see loading from the baffles and cooled optics sleeve, see Table 2.1. As internal loading in stratospheric platforms dominates CMB power, satellite experiments will generally suffer less from noise.

Bolometric noise modeling and inventory is discussed in a wide array of publications [131, 132]. The most common noise terms involve statistical photon noise, Johnson noise in the thermistors, phonon noise from the detector’s thermal environment, and detector readout noise. See Section 2.8 for more extensive discussion of the SPIDER detector architecture. The noise-equivalent power, P_{photon} , in a polarization sensitive bolometer illuminated by a stream of photon with time averaged power, P , is [74, 133]

$$P_{\text{photon}} = 2h\nu P + \frac{2P^2}{\Delta\nu}, \quad (2.1)$$

where the first term on the right-hand-side corresponds to shot noise, fluctuations in arrival times of light quanta, and the latter term accounts for occupation bunching.^{5,6} In the

⁴Developments in cryogenic bolometer technology, which took place around the turn of the 20th century, allowed experiments to probe higher frequencies than those allowed by coherent HEMT receivers and extend well past 100 GHz while being photon noise limited [127].

⁵Note that P_{photon} has units of W^2/Hz .

⁶Many use the term NEP to express noise-equivalent power, in this discussion I have chosen to use the parameter P_{photon} as it is more compact.

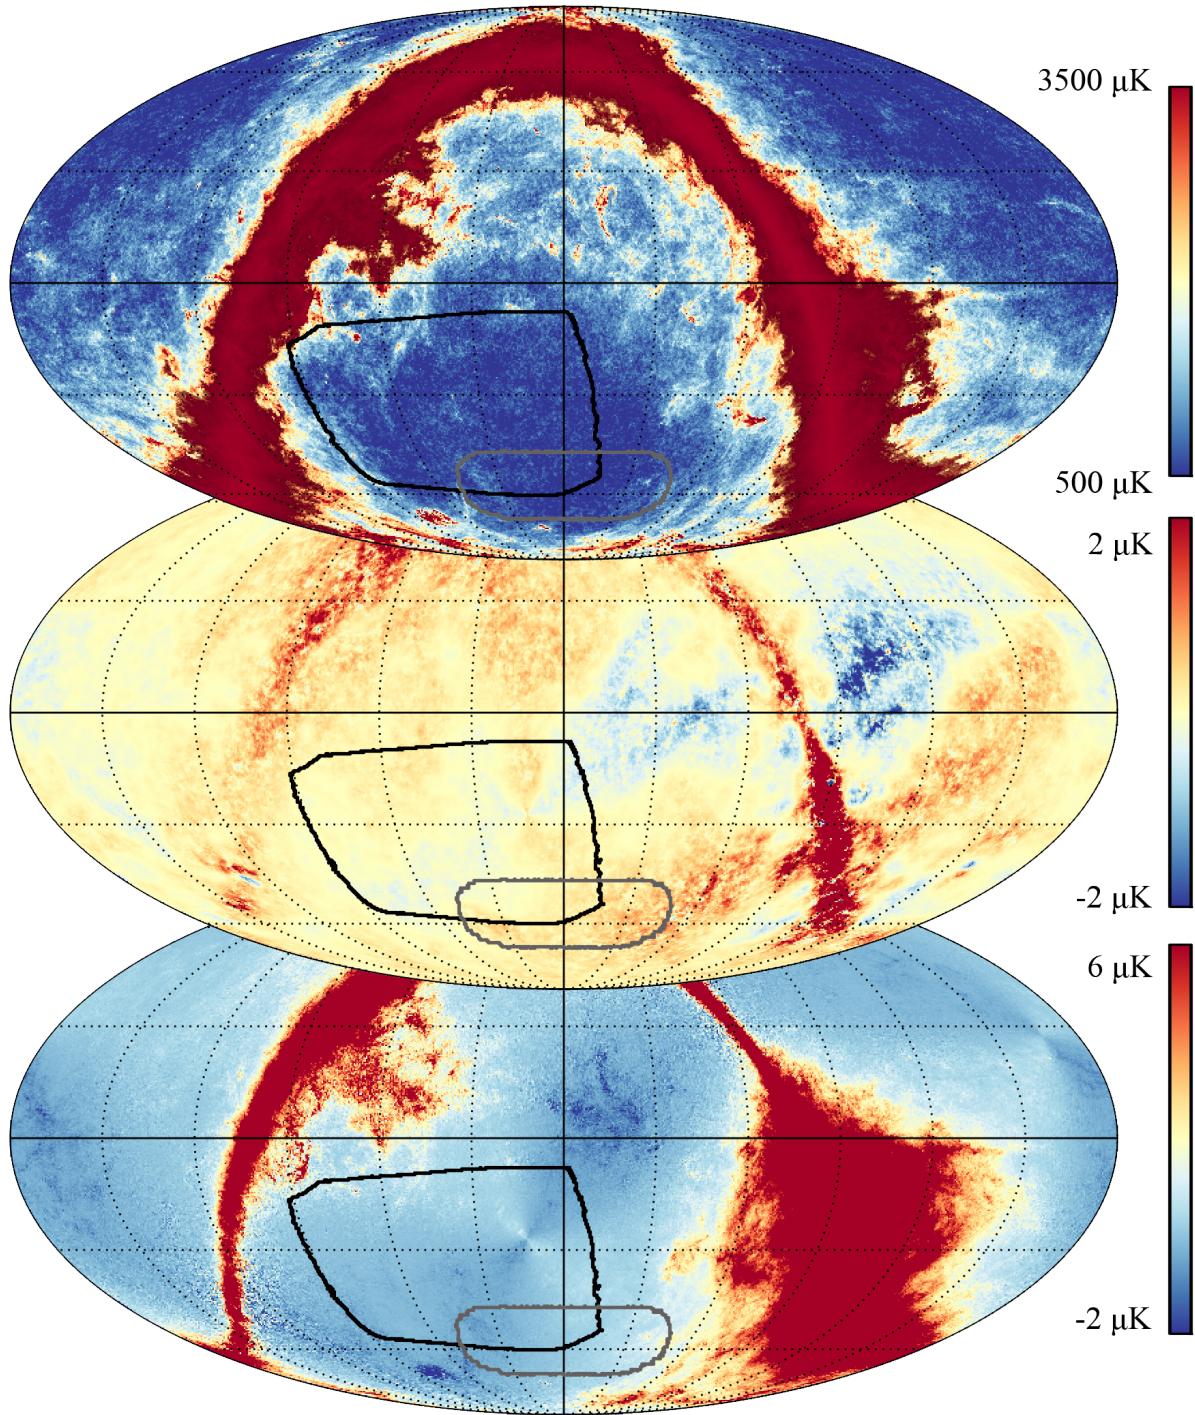


Figure 2.2: The outlines of the SPIDER and BICEP2 observation regions are shown in black and grey respectively. All maps are presented in equatorial coordinates using a Mollweide projection. *Top*: The *Planck* HFI 353 GHz temperature map. *Middle*: Predicted *Q* polarization of synchrotron emission at 90 GHz as predicted by the *Planck* Sky Model [129]. *Bottom*: Dust *Q* polarization at 150 GHz from the O'Dea model [130].

above expression, ν is the frequency of the incident photons, $\Delta\nu$ is the bandwidth under investigation, and h is Planck's constant. Thermal fluctuations intrinsic to the bolometer thermal environment are also unavoidable. For the following discussion we define $\Delta T \equiv (T - T_{\text{base}})$, where T_{base} is the temperature of the thermal reservoir, T is the temperature of the bolometer, and $\epsilon \equiv \Delta T/T_{\text{base}} \approx 0.1$ for a typical SPIDER detector. The phonon shot noise term, intrinsic to the thermal conductance, corresponds to

$$P_{\text{phonon}} = 4k_B T^2 GF, \quad (2.2)$$

where G represents the thermal conductivity of the link between the TES and its thermal reservoir, and F is a correction factor which accounts for non-equilibrium temperature differences across the path of conductance; $F \approx 0.5$ in the case of SPIDER. Similarly, resistor (Johnson) noise due to thermal fluctuations of charge carriers follows

$$P_{\text{Johnson}} = 4k_B T R. \quad (2.3)$$

Careful accounting shows that the photon term, Equation 2.1, dominates the contributions to the total noise budget of the SPIDER detectors.

The SPIDER detectors have been carefully tuned so that the unavoidable contribution of photon noise dominates the total noise budget of the SPIDER detectors.⁷ First, it is easy to see that $G = Q/\Delta T \approx P/\Delta T = P/\epsilon T_{\text{base}}$ if the photon loading is balanced by the thermal link to the heat reservoir; note that Q represents heat flux between the reservoir and the bolometer. Equating photon and phonon noise contributions, and assuming shot noise dominates occupation bunching, we arrive at the following expression for the temperatures of the bolometer and heat reservoir:

$$T_{\text{base}} \approx \frac{h\epsilon\nu}{2k_B F}. \quad (2.4)$$

⁷This is a simplification. The detector properties are also tuned so that we do not saturate the aluminum superconducting transition on the ground.

Evaluating this expression using a typical set of SPIDER device parameters, corresponding to $\nu = 100 \text{ GHz}$ and $\epsilon = 0.1$, we find that $T_{\text{base}} \approx 480 \text{ mK}$. In order for thermal fluctuation in the bolometers to be sub-dominant to the expected contribution from the photon shot noise, we have to cool them below 480 mK. This is the reason CMB detectors are cooled to sub-Kelvin temperatures. Using measured device parameters, including thermal conductivities and spectral bandwidth, we can model the relative contribution of noise terms for SPIDER’s detectors under flight-like conditions. The noise properties can also be probed directly, for example, by using a custom beam-filling cold load that can be bolted to the front of a test cryostat at Caltech. The temperature of the cold load, which approximates a blackbody, can be varied to simulate different thermal environments.

2.1.2 Recent Developments

Recent, potentially groundbreaking, measurements made by the BICEP2 collaboration have drawn a lot of attention [88, 108]. It remains to be seen whether other experiments can corroborate these findings. It is therefore interesting to compare BICEP2 and SPIDER. In many ways, these two experiments are quite comparable. They share the same science goal: to push down the limits on primordial B -mode amplitudes. The two employ a common detector architecture and a very similar telecentric optical design, originally drafted for the BICEP1 experiment. One is ground based, the other balloon borne, and since photon incidence dominates the noise budget for both experiments, the two have wildly different instantaneous sensitivities. As a result, twenty days of balloon borne observations yield map depths that are comparable to 2–3 years of ground based observations. Both experiments were conceived by a common scientific leader, Andrew E. Lange, who passed away in 2010.

Figure 2.2 shows a collection of sky maps with the SPIDER and BICEP2 observation regions outlined. The maps are based on *Planck* data, the *Planck* Sky Model [129], and the dust model of O’Dea et al. [130]. SPIDER will cover approximately seven times more sky than BICEP2. As polarized foreground models remain data starved, estimates of foreground

contamination in different regions suffer from significant uncertainties. Large sky coverage, two frequency bands, and almost disjoint observation regions should allow SPIDER to set limits on potential systematic inducing effects in the BICEP2 analysis. The dataset will certainly help advance Galactic foreground removal techniques [134, 135, 136]. Furthermore, assuming the BICEP2 measurements are sample variance limited, larger sky coverage will also help constrain the signal variance.

Figure 2.3 shows the power spectrum of Galactic foregrounds at 94 and 150 GHz as predicted by the *Planck* Sky Model [129] and O’Dea et al. [130]. The colored regions represent the variation obtained by calculating the power spectrum from various BICEP2 sized patches inside the SPIDER region. The models used for this analysis give results that are comparable to the analysis presented in [88]. The foreground signal is approximately one order of magnitude below a *B*-mode signal from $r = 0.20$. It has been pointed out that if foreground amplitude were off by a factor of 2–3, the foreground power would become comparable to the measured *B*-mode signal [125, 137].

Given SPIDER’s detector count, yields, scan strategy, and projected observation time, SPIDER will produce sky maps with 200 nK sensitivity per deg^2 compared to 83 nK quoted in BICEP2’s *B*-mode paper [88], which was obtained from 2–3 years of observation. However, neither of these quoted map depths account for information loss due to filtering of timelines. As ground based experiments normally have to perform aggressive filtering compared to balloon borne or satellite missions, we suspect that the effective map depth of these two experiments will be quite similar. It is worth mentioning the *Planck* HFI map depth, which is 1.0, 0.8, 1.3, 5.0 $\mu\text{K}/\text{deg}^2$ at 100, 143, 217, and 353 GHz respectively [138].

Recent developments in detector manufacturing technology and the push for greater sensitivities is giving rise to experiments that employ tens of thousands of detectors [110]. The SPIDER experiment is a step in that direction. The great success of the three orbital CMB missions can be largely attributed to technology which was prototyped on sub-orbital

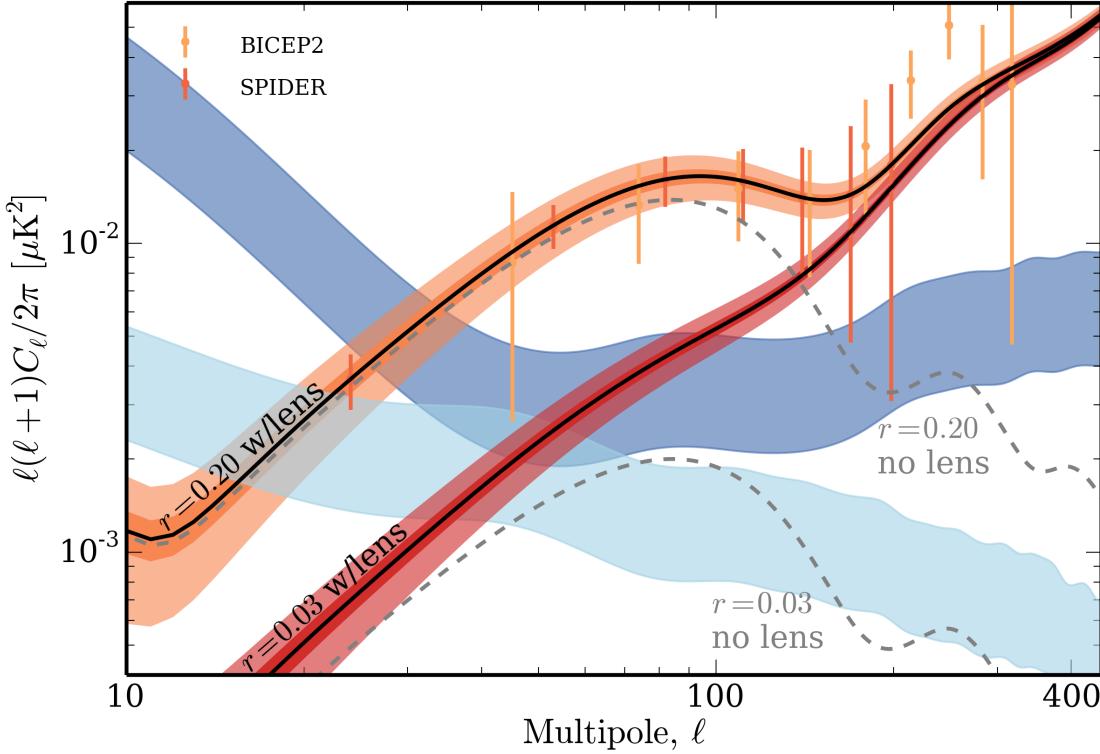


Figure 2.3: The amplitude of 90 and 150 GHz Galactic foreground emission in the SPIDER and BICEP2 observation regions, as predicted by the *Planck Sky Model* [129] for synchrotron emission at 90 GHz, shown with light blue hues, and O’Dea et al. [130] for dust emission at 150 GHz, shown with blue hues. Figure 2.2 shows some of the corresponding maps. The width of the colored regions represents variations in foreground amplitudes over an ensemble of BICEP2-sized maps drawn from the SPIDER region. Also shown are the BICEP2 measurements (yellow points) and the predicted sensitivity of a 20 day flight of the SPIDER experiment as derived from Fisher matrix analysis (red points). The colored regions that trace the $r = 0.03$ and $r = 0.20$ curves represent the cosmic variance for both 1% and 10% sky fraction. Fisher analysis data courtesy of Anne Gambrel.

missions. Current stage balloon borne experiments, such as EBEX, PIPER, and SPIDER will help prepare for the fourth generation inflation probe [139, 140].

2.2 Collaboration

The SPIDER experiment is a collaborative effort with contributions from many academic institutions. Primary funding comes from NASA through an *Astrophysics Research and*

Analysis (APRA) grant. Considerable financial support is also provided to our collaborators at University of Toronto through the Canadian Space Agency, with additional funding for this experimental effort from various other funding agencies, including the Gordon & Betty Moore and David & Lucille Packard Foundations. The SPIDER instrument was first proposed in 2005, but unsuccessfully. A second 2006 proposal for an APRA grant resulted in initial funding that began in 2007. A second APRA proposal for the SPIDER experiment was accepted in 2011. Quite a few experimentalists have made significant contributions during the design and build phase of this experiment, and considerable heritage derives from previous ballooning and ground based experiments, including BOOMERanG, BICEP1, BICEP2, and the *Keck Array*. Detector development, screening, and characterization is primarily done at the Jet Propulsion Laboratory (JPL) and the California Institute of Technology (Caltech). Half-wave plate and baffle development took place at Case Western Reserve University. Payload interfacing, including scanning control, pointing reconstruction, and power systems were designed and built by the University of Toronto, while detector readout systems are provided by the University of British Columbia. Cryogenic characterization and general payload integration has taken place at Princeton University. Scientists from many other institutions actively contribute to this experiment. Approximately 50 scientists have made significant contributions to this effort with about 15–20 active experimentalists at any given time.

2.3 Ballooning

Ballooning expertise for the SPIDER experiment is provided by the Columbia Scientific Ballooning Facility (CSBF), a NASA institution established in 1961. The facility conducts ballooning launches from various locations on the globe, including Esrange, Sweden; Alice Springs, Australia; and Fort Sumner, New Mexico. The first Antarctic launch was conducted in 1990 [141]. At the time of writing, 39 LDB flights lasting longer than 24 hours have flown from McMurdo Station, the southernmost anchorage in the Antarctic. The mean duration of

these flights is 20.1 days. Figure 2.4 shows a histogram of flight lengths as well as the best fit log-normal distribution function. Various reasons factor into decisions to terminate flight, with ease of payload recovery playing a large role.

Figure 2.5 shows the empirically determined probability density function of all scientific balloon trajectories launched from McMurdo Station in the last 24 years as well as the corresponding landing sites. That distribution of trajectories suggests a 79.2 deg mean latitude with 3.5 deg standard deviation. From Figure 2.5 it is clear that flights are preferentially terminated close to McMurdo Station.

Seasonal wind patterns in the polar vortex emerge during the Austral summer, allowing for circumpolar trajectories at 79 deg latitude with 9–12 day periods. The SPIDER payload will launch during this period, rising at an average velocity of 4 m/s, it will reach float altitude after approximately 3 hours of flight. During ascent the payload is expected to experience 0.5 g accelerations due to wind shear and to rotate in an uncontrolled fashion with angular velocities up to two revolutions per minute [142]. As the payload travels through the upper regions of the troposphere, the temperatures can go as low as 200 K [143].

Electrical components are tested in a chamber that simulates the thermal environment, which includes drops in both temperature and pressure, in order to ensure that they perform adequately during ascent. Thermal modeling complements these tests and allows us to determine proper placement of various flight computers and other electrical devices [143, 144, 145, 146]. No attempt is made to control pointing on ascent, as the forces experienced would put too much strain on pointing motors. This way, also, batteries will not drain during ascent. At a float altitude of 37 km, the pressure is approximately 5–10 torr with ambient temperatures hovering around 250 K. At this altitude, most of the Earth’s atmosphere is below the payload, which means that the sky is dark even during the day. Typically, the payload will drift in the polar vortex at a steady rate of about 10 m/s, albeit with large variations.

The SPIDER experiment will fly on 37 “Heavy”, a 20 μm thick, polyethylene balloon with

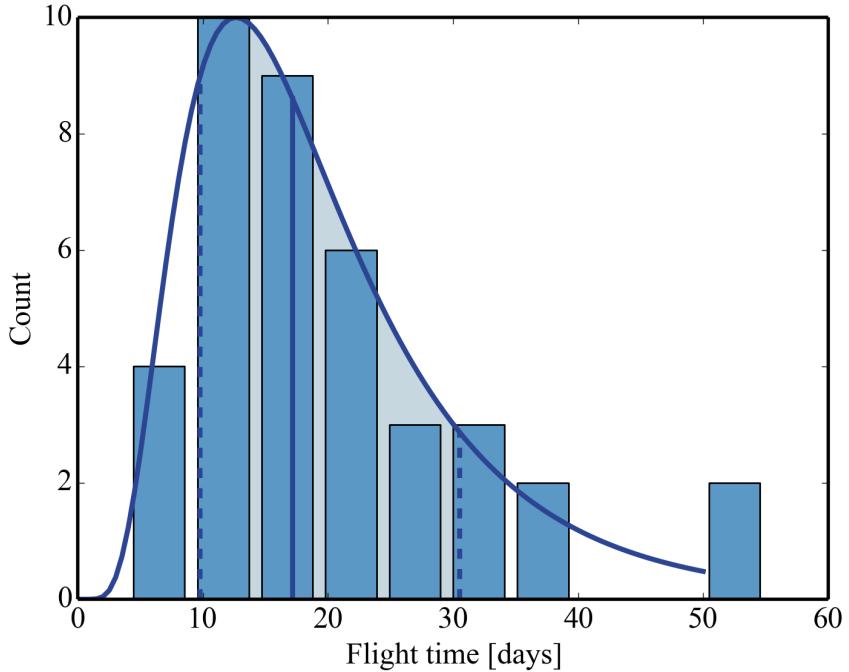


Figure 2.4: A histogram of flight times for all Antarctic LDB flights and a best fit log-normal distribution. The mean of all flight times is 20.1 days, while the best fit distribution suggests a slightly smaller mean and a considerable tail towards longer flight times. Since the turn of the century, there have been 26 flights, with a mean flight time of 23 days. The mean flight duration is biased by the need to recover payloads in a short amount of time. This means that flight terminations are predominantly near McMurdo Station, and for that reason, flights often terminate sooner than strictly speaking necessary.

8000 pound gross lift capacity. At mission termination, a detonator separates the payload from the balloon, resulting in free fall for approximately 5 seconds. A type of shock absorber, a so-called rip-cord, is deployed during this process. It is designed to limit the acceleration felt by the payload as the parachute is opening to no more than 5 g 's. The payload will then slowly decelerate as it glides through the troposphere hitting the ground at a velocity of 8 m/s approximately 10 minutes after flight termination. After landing is confirmed, another mechanism separates the payload from the parachute to prevent it from being dragged by surface winds.

Steady wind patterns, specialized infrastructure, lack of populace, and predictable geography make McMurdo Station and the Antarctic well suited for ballooning flights lasting longer than a week. However, there are no technological or meteorological constraints

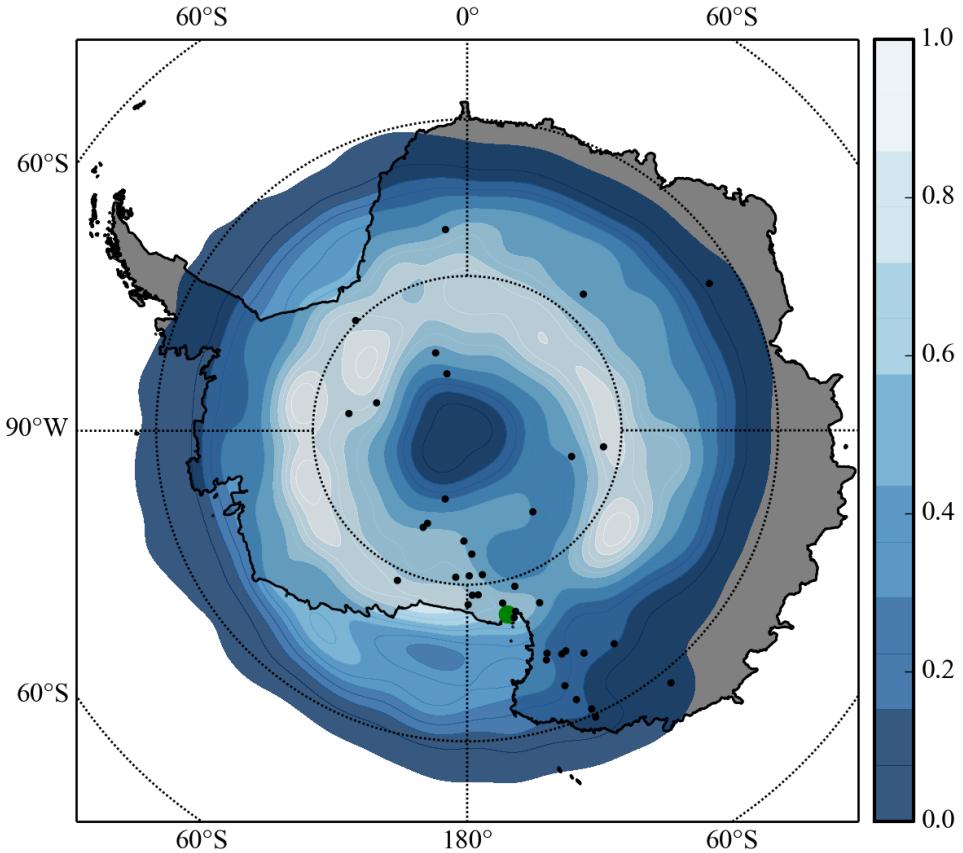


Figure 2.5: An empirical probability density function derived from all long duration ballooning flights in the last 24 years, or since the start of the Antarctic scientific ballooning [141]. Data within 24 hours of launch are omitted. Balloons are launched from McMurdo Station, marked by a green dot, and travel in a counter clockwise motion as seen on this map. The data suggest a mean latitude of 79.2 deg with 3.5 deg standard deviation. Black markers show the termination locations of all 39 flights used in this analysis. Data courtesy of CSBF [149].

preventing similar flights in the Northern Hemisphere. Successful circumpolar flights from Svalbard, at approximately the same latitude as McMurdo Station, have been made, with some in preparation [147, 148].

Long duration ballooning presents a cost-effective alternative to satellite missions, with reduced overhead cost and shorter development cycles. At float altitude of 37 km, atmospheric loading is minimal, resulting in almost space-like conditions. Ballooning, however, retains many of the mass, power, and telemetry constraints experienced in satellite missions, and, therefore, represents a great training platform for scientists.

2.4 Cryogenic Architecture

The following section is largely based on an SPIE proceedings article [150], with further results from an additional four years of cryogenic operations. We begin by reviewing the cryogenic architecture and follow that with a discussion of standard cryogenic operations and major results. We also describe a thermal model of the flight cryostat and present the phenomenology of a helium leak in the main tank.

The SPIDER science goals require cooling of six 1.3 m long telescope inserts with 30 cm apertures down to 1.8 K. The optical throughput and physical dimension of the science instruments determine the scale of the cryostat as well as the minimum parasitic load to the helium bath of several hundred mW. Early modeling suggested that the enthalpy of the helium vapor produced by the expected load would provide enough cooling power at roughly 20 and 110 K to eliminate the need for a separate liquid nitrogen bath. Additionally, it was found that the estimated load would allow for more than 25 day hold times assuming a tank volume of at least 1000 L. It was concluded that a helium-only system would reduce the total mass and simplify the design without compromising thermal performance.

During the design phase of the SPIDER experiment, it was also recognized that a 1000 L flight cryostat posed a manufacturing risk. To further inform the design of the cryostat, three test cryostats capable of housing a single telescope were built. These were built by a group of engineers at *Redstone Aerospace*, in addition to William Jones, Peter Mason, Tracy Morford, and Amy Trangsrud, all at Caltech at the time.⁸ Thermal characterization of these cryostats confirmed predictions from simple thermal models (see Section 2.4.3).⁹ Based on these results, design requirements were defined and agreed upon by Redstone Aerospace, the manufacturer of the SPIDER flight cryostat. See Table 2.2 for an account of main characteristics.

The cylindrically shaped flight cryostat has five main components that are illustrated in

⁸*Redstone Aerospace*, Longmont, CO.

⁹One of the three cryostats was found to suffer from a microscopic helium leak which was later fixed. The three test cryostats are now known as: “the SPIDER Test Cryostat”, “the Hienostat”, and “the BICEP2 Cryostat”.

Table 2.2: Flight cryostat main specifications.

Property	Value
Vacuum Vessel height	2.43 m
Vacuum Vessel diameter	2.11 m
Vacuum Vessel volume	5700 L
Main Tank net cryogenic volume	1284 L
Superfluid Tank net cryogenic volume	16 L
Mass of cryogenic assembly	850 kg
Hold time	≥ 25 days

Figure 2.6. Starting from the inside, the components are named: Superfluid Tank (SFT), Main Tank (MT), Vapor Cooled Shields 1 and 2 (VCS1, VCS2), and Vacuum Vessel (VV). The bulk of the cryostat is made of aluminum 1100 (VCS1, VCS2), chosen for its high thermal conductivity, and aluminum 5083 (MT, SFT), which maintains its strength after welding. The cryogenic assembly consists of a cylindrical 1284 L Liquid helium (LHe) main tank, connected through a capillary system to a 16 L superfluid tank.

VCS1 surrounds both tanks and serves as a radiation shield from warmer stages, while intercepting conduction and accommodating filters, which need to be maintained at low temperatures to reduce in-band parasitics. VCS2 provides additional radiation shielding from the VV, which is coupled to ambient temperatures. The dry weight of the cryogenic assembly is roughly 850 kg.

The superfluid tank will operate at approximately 1.8 K, defined by the atmospheric pressure at floating altitudes. Heat straps connect the superfluid tank to each of the telescope tubes, providing a base temperature for 10 STPL closed-cycle ^3He adsorption refrigerators that cool each focal plane.^{10,11} Of order 10 g of activated charcoal are installed inside a single refrigerator. When cooled to 4 K, the charcoal is capable of adsorbing all gaseous ^3He in the refrigerator. The adsorption refrigerator is equipped with a heat switch which can thermally

¹⁰STPL stands for Standard Temperature and Pressure Liter.

¹¹Refrigerators supplied by *Chase Research Cryogenics*, Sheffield, UK.

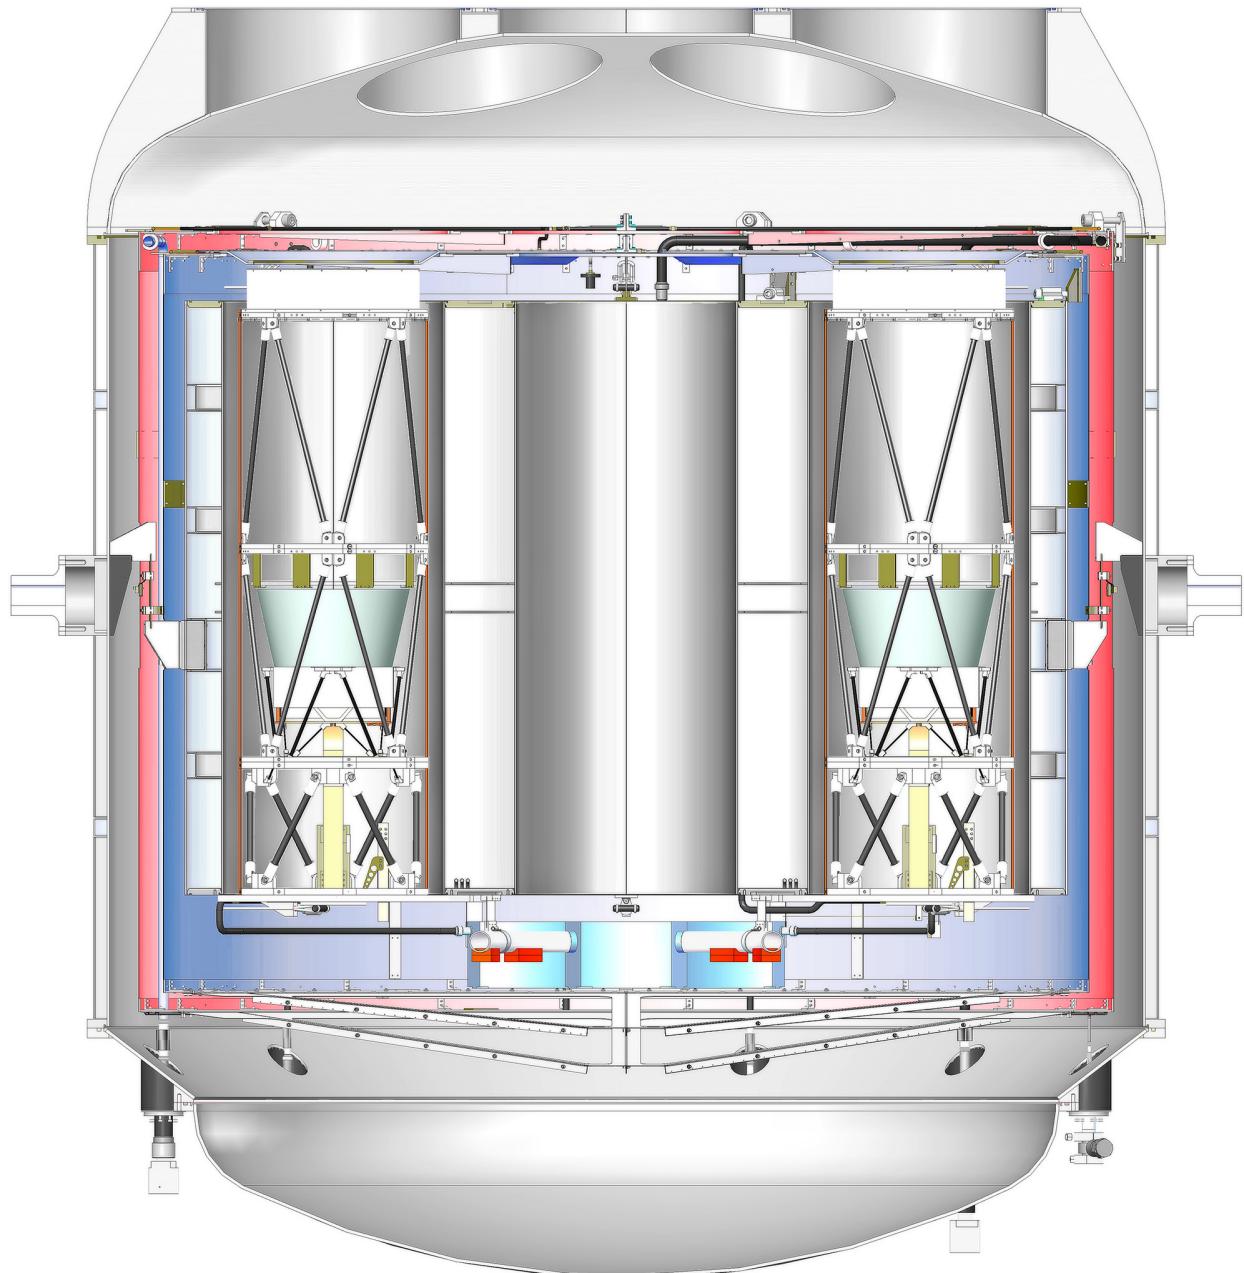


Figure 2.6: A cross section rendering of the flight cryostat with two inserts visible. The outermost layer, the vacuum vessel, surrounds the two vapor cooled shields, shown in red and blue. At the center, the main and superfluid tanks are colored in grey and light blue respectively. Although not visible, the vacuum vessel pumpout port is a KF50 fitting located at the bottom of the vacuum vessel. Another five plumbing lines exit at the bottom: the main and superfluid tank fill and vent lines, and the VCS vent line.

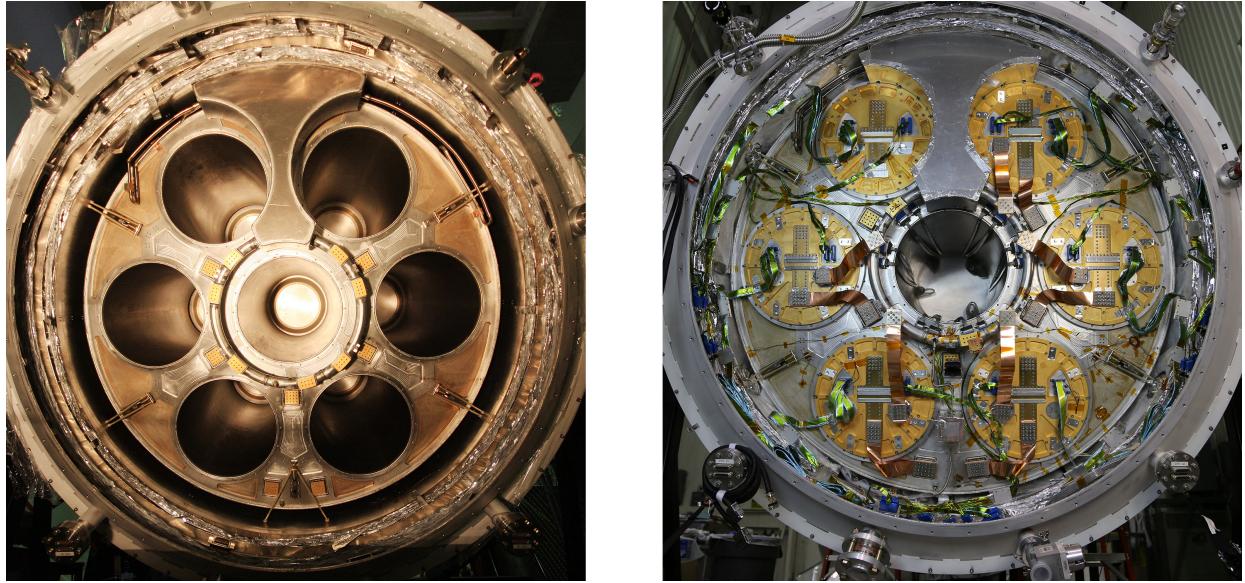


Figure 2.7: The view of the bottom of the main tank. *Left*: No inserts installed. *Right*: All six telescopes installed.

link the refrigerator to the 1.8 K base temperature of the superfluid tank. With the heat switch turned on, any ^3He gas in the refrigerator will cool, condense, and eventually fall into a small reservoir, located gravitationally below the heat switch contact. The heat switch is normally flipped some time after the charcoal has been heated up. This ensures that all 10 STPL are gaseous at the time when liquification begins. With heating of the charcoal pumps ceased, ^3He gas will begin to adsorb again, resulting in a reduction in the vapor pressure of the ^3He bath and, therefore, the temperature of the bath. A successful fridge cycle results in a 300 mK still temperature with hold times of 3–4 days.

2.4.1 Radiation Shields, Heat Exchangers, and MLI

Two intermediate aluminum 1100 vapor cooled radiation shields, VCS1 and VCS2, serve as thermal anchors for multi-layer insulation, filter blocks, and heat exchangers. VCS1 is supported by the MT, while VCS2 is supported from the VV. Six compact heat exchangers are symmetrically placed on the top sides of both VCS1 and VCS2. Cryogenic boil-off is forced to go through these flow-restrictive heat exchangers, cooling the respective stages, and

thus providing negative feedback [151, 152]. The heat exchangers are made of stainless steel blocks enclosing densely packed horizontal mesh (VCS1) or pellets (VCS2), both made of copper.

The outer sides of VCS1 and VCS2 are layered with aluminized polyester films (known as aluminized Mylar, MLI for Multi-Layer Insulation, or superinsulation), with 16 layers on the outside of VCS1 and 52 layers on the outside of VCS2. These $6.4\ \mu\text{m}$ thin layers of thermoplastic polymer provide a lightweight substrate for highly reflective 35 nm thick aluminum layers coated on both sides. The MLI is designed to provide sufficient radiation suppression to maintain large temperature gradients between stages. A 0.1 mm thick, spun-bound polyester sheet is placed between each layer of Mylar to reduce thermal conduction, which has the potential to dominate radiative effects if the layers have strong thermal contacts. The surface area of each stage ranges between 10–14 m^2 . The MLI packing density of 14 layers/cm reduces compression that would otherwise lead to undesirable conductance effects. This is roughly a factor of two lower than the quoted optimal packaging density for MLI [153, 154]. However, lower packaging densities enable more effective evacuation of the vacuum vessel.

2.4.2 Flexures and Plumbing

The main tank is supported by the vacuum vessel through six G-10/aluminum flexures symmetrically placed on the cylinder sides (see Figure 2.8). G-10, a lightweight, high tensile strength fiberglass, is an extremely poor conductor of heat, and, therefore, ideal for cryogenic flexures [155]. In order to dampen heat flow from the VV to the MT, the flexures are heat sunk at VCS1 and VCS2 using copper straps (see Figure 2.9). Similar but considerably smaller G-10 flexures, 33 mm long, 29 mm wide, and 0.8 mm thick, connect VCS1 to the MT and VCS2 to the VV at six points on the cylinder sides. Three axial flexures help align VCS1 with respect to the MT, and six flexures connect the SFT to the bottom of the MT.

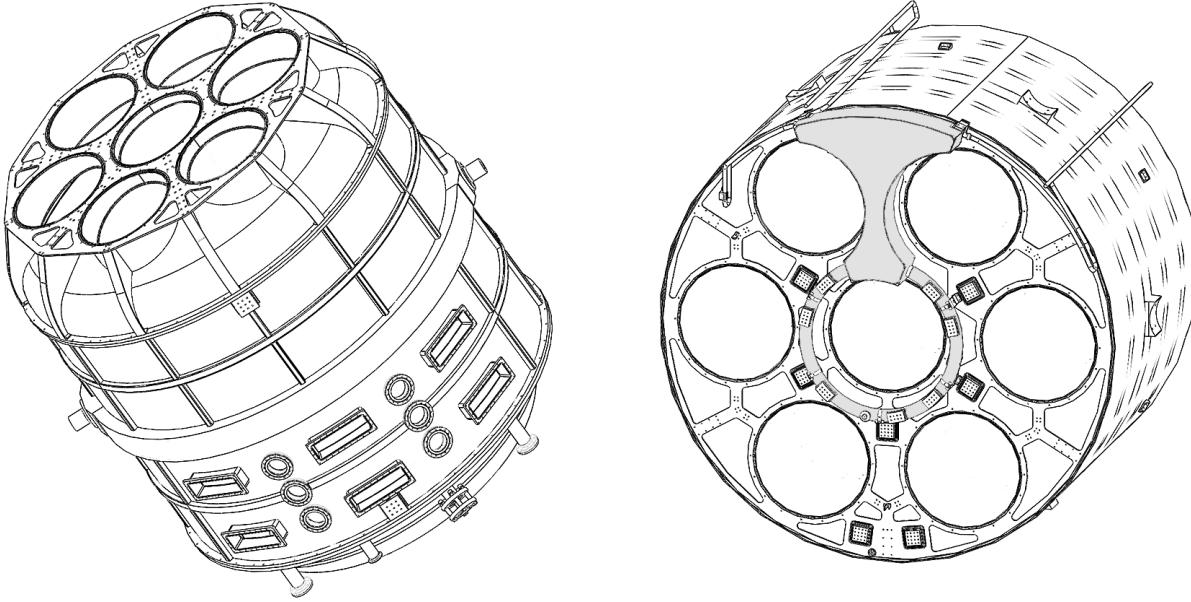


Figure 2.8: *Left*: CAD model of the vacuum vessel, the outermost stage of the flight cryostat. The dry weight of the cryogenic assembly, excluding telescope inserts, is 850 kg. All fill and vent lines exit the VV at the bottom. Ports on the front side of the cylinder provide hermetic connections to housekeeping electronics. Two trunnions on the center portion of the VV walls attach to the elevation drive on the gondola. The center insert will not be used. A cross section through the VV can be seen in Fig. 2.6. *Right*: CAD model of the main tank and superfluid tank assemblies as viewed from the bottom. Note the seven telescope inserts constituting extruded cuts through the MT. The SFT has a ring-like structure that connects to another larger cryogenic volume located under the MT. Explosion-bonded thermal contact areas, both on the MT and SFT, are connected to each insert to provide cooling power directly from the 4 K and 1.8 K baths. The cylindrical MT is 1.69 m in diameter, 1.14 m long, and weighs 220 kg. The insert diameter is 419 mm and the thickness of the MT walls varies between 4 and 6 mm. Observations will be performed with the cryostat tilted at 25–40 deg elevation such that the bulk of the SFT is above the ring like structure, which holds only about 0.5 L.

SPIDER will scan at a 25–40 deg elevation such that the liquid level of normal liquid does not lie in the plane defined by the bottom of the main tank. For this reason, and since the superfluid tank will never be entirely full, it is important that the helium in the SFT is superfluid to ensure that sufficient cooling power is supplied to each insert. This fact motivates the shape of the SFT, which can be seen in Figure 2.8, and the placement of thermal contact areas on the ring-like structure. Explosion-bonded aluminum-to-copper transition

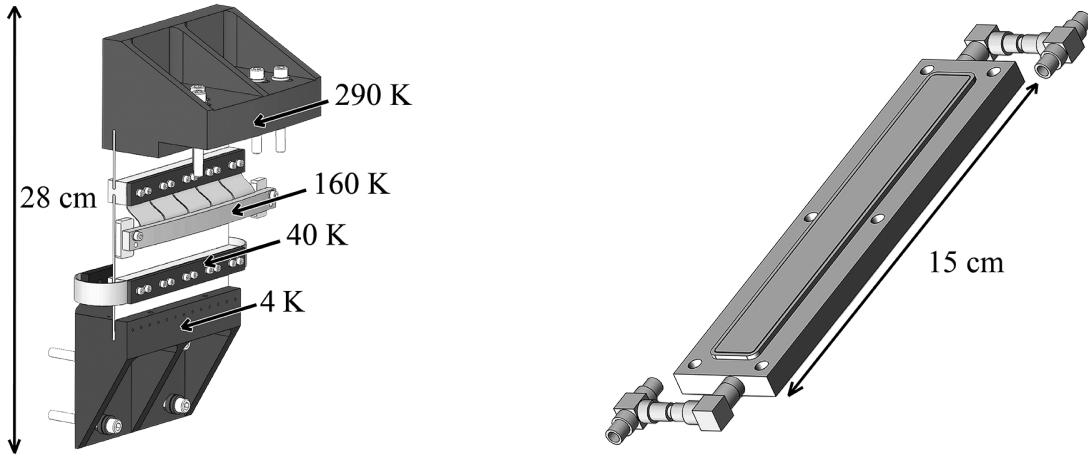


Figure 2.9: *Left:* One of the six main structural flexures which are symmetrically spaced around the main tank cylinder. These flexures support the main tank off of the vacuum vessel through a G-10 sheet which is thermally damped by copper heat straps connected to VCS1 and VCS2. The G-10 flexures are 1.6mm thick, 127mm wide, and 114mm long. *Right:* One of the six heat exchangers placed on the top of VCS2. The 15 cm long rectangular compartment is filled with copper pellets creating high flow impedance, which extracts enthalpy from the helium boil-off. The heat exchanger system forms a double annular structure on each of the vapor cooled shields. Helium vapor exiting the cryostat through the VCS vent must go through one of the six heat exchangers on both VCS1 and VCS2. The cylindrical symmetry of the heat exchanger system ensures that equal enthalpy is extracted by the heat exchangers.

plates provide reliable thermal contact areas from both the MT and the SFT.¹² Parts of the aluminum are milled out such that the copper is in direct contact with the cryogen. The remaining aluminum is then welded on to the corresponding tank. The transition plates are connected to custom-made copper heat straps, and supply the necessary cooling power to telescope inserts and sub-Kelvin cooling stages. The MT thermal contact areas can also be seen in Figure 2.8.

Plumbing lines are made of type 304 stainless steel, due to its low thermal conductivity and suitability for welding. There are five plumbing lines leading from the outside of the VV to either the MT or the SFT. These are: the MT fill and vent lines which have a 3/4 inch Outer Diameter (OD), the SFT fill and vent lines which have a 1/2 inch OD, and the VCS vent line which has a 1/4 inch OD. Aluminum-to-stainless steel transitions are made

¹²High Energy Metals, Inc., Sequim, WA.

Table 2.3: Cryogenic tasks required for cooldown – in chronological order.

Task	Time
Multiple purge cycles	36–48 hours
Pump with turbomolecular station	3–4 days
Fill main tank and equilibrate at LN ₂ temperatures*	4–5 days
Cool to LHe temperatures**	3 days
Total	11–14 days

* LN₂ = liquid nitrogen. ** LHe = liquid helium.

from explosion-bonded blocks that are welded in place. MT fill and vent lines will be capped off prior to launch, forcing helium boil-off to leave the MT through the heat exchangers. All plumbing lines, excluding the VCS vent line, are strategically heat sunk at VCS2 and not at VCS1. The length of the MT/SFT vent and fill lines is approximately 2.7 m, while the average travel of gas through the VCS vent line system is about 12 m. Vent lines are positioned on cryogenic tanks such that boil-off will be able to exit the cryostat when it is tilted at a 25–40 deg elevation angle and full of liquid.

2.4.3 Cryogenic Operations

The construction of the three SPIDER test cryostats informed the design of the flight cryostat. Thermal qualification of those systems showed that a helium-only cryostat with two vapor cooled shields could successfully cool an entire BICEP1-like telescope, while maintaining minimal in-band loading from warmer temperature stages, and a reasonable cryogen boil-off rate. From this it was concluded that the heat exchangers cooling the intermediate stages, and thereby providing negative feedback, were operating as expected. The testing also suggested that the multilayer insulation would provide sufficient insulation and that thermal conduction paths to sub-Kelvin stages were adequate to cycle adsorption refrigerators. The build phase

of the SPIDER flight cryostat began in the Fall of 2008 and was completed in the Summer of 2009. The cryostat¹³ was delivered to Princeton in January 2010.¹⁴

Cooling the SPIDER flight cryostat down to nominal temperatures requires approximately two weeks of pumping and liquid nitrogen pre-cooling. The nitrogen pre-cooling effectively removes 90% of the combined enthalpy of the telescopes and cryostat using 300–400 L of liquid nitrogen. The remaining cooldown to 4 K requires approximately 500 L of liquid helium before the cryostat has fully equilibrated at 4 K. Table 2.3 highlights the main aspects of the cooldown schedule and their approximate duration. The considerable time requirement is set by the net volume of the vacuum vessel – approximately 6 m³ – as well as the total mass that is cooled to the 4 K base temperature – about 350 kg from the main tank and an additional 48 kg per insert. Of the 14 test runs performed since the cryostat was built, excluding the liquid nitrogen runs performed at Redstone Aerospace, the cryostat has been cooled down to liquid helium temperatures ten times. The remaining four runs only used liquid nitrogen. The total liquid helium used for those ten runs amounts to little over 30,000 L. For two of those runs, the cryostat has been fully populated with telescope inserts. Table 2.4 describes some aspects of the cryogenic performance during all of the liquid helium runs conducted so far.

Pumping

The vacuum vessel requires approximately one week of pumping before the initial liquid nitrogen cooldown can begin. During the pumpdown procedure, we find that backfilling with dry nitrogen helps to reduce the asymptotic absolute pressure. Our experience suggests three to four gaseous nitrogen (GN2) backfilling procedures are useful, but any subsequent purge does not seem to help reduce the asymptotic pressure. Figure 2.10 shows the pressure profiles as a function of time for a few pumpdowns performed with a fully populated optical system,

¹³Within the collaboration, the flight cryostat is referred to as “Theo,” after Theodosia Burr. The act of trying to understand Theo’s cryogenic behavior is sometimes jokingly referred to as “Theology.”

¹⁴At about the same time the BICEP2 cryostat, one of the three SPIDER test cryostats, began what would become a three year cryogenic run with an average liquid helium consumption of about 22 L/day [108].

Table 2.4: Cryogenic score sheet for all liquid helium runs of the SPIDER flight cryostat. From left to right, T_{VCS1} and T_{VCS2} represent the average temperature of the two vapor cooled shields, f_{VCS} is the flow rate out of the main tank as measured by a flow meter on the output of the VCS vent line in units of Standard Liters Per Minute (SLPM), t_{hold} is the hold time of the cryostat assuming an initial charge of 1000 L, f_{He} represents an approximate estimate for the quiescent helium background as measured by a leak checker placed close to the pump out port while the cryostat is in equilibrium at 4 K, d_{pump} is the number of days that we were actively pumping on the vacuum vessel before starting initial liquid nitrogen fill, and N_{ins} is the number of telescope inserts in the cryostat.

Run	Date	T_{VCS1} [K]	T_{VCS2} [K]	f_{VCS} [SLPM]	t_{hold} [days]	f_{He} [mbar/l/s]	d_{pump} [days]	N_{ins}
3	03/2010	30 ± 3	160 ± 15	N/A	N/A	2×10^{-7}	6	0
4	05/2010	28 ± 3	168 ± 15	20 ± 2	24 ± 2	4×10^{-8}	8	0
6	02/2011	26 ± 3	142 ± 15	20 ± 3	24 ± 2	4×10^{-9}	9	0
7	08/2011	28 ± 4	157 ± 15	21 ± 3	23 ± 2	8×10^{-9}	7	1
9	05/2012	31 ± 3	151 ± 15	23 ± 3	21 ± 2	6×10^{-7}	7	1
10	07/2012	32 ± 3	153 ± 15	24 ± 4	20 ± 2	2×10^{-7}	7	1
11	10/2012	33 ± 3	151 ± 15	23 ± 3	20 ± 2	1×10^{-7}	6	2
12	02/2013	31 ± 3	148 ± 15	25 ± 4	19 ± 2	N/A	12	2
13	06/2013	44 ± 3	160 ± 15	54 ± 5	9 ± 1	4×10^{-7}	7	5
14	02/2014	42 ± 3	157 ± 15	38 ± 4	13 ± 1	1×10^{-6}	14	6

cryogenic Runs 13 and 15. The idea behind dry gas purges is that the gas helps dislodge water molecules that have desorbed to various surfaces inside the vacuum vessel while the chamber was exposed to atmosphere. Under that assumption, the effectiveness of a dry gas purge will depend on the time and ultimate pressure levels as well as the temperature and mass of the gas that is used.¹⁵ We find that purging to higher pressure, say 760 torr rather than 100 torr, is generally more effective. We use an *Agilent TriScroll 600* for this initial “rough-out” phase of the pumpdown.¹⁶ Figure 2.10 shows a dramatic change in pressure profiles after the first GN2 purge during Run 13.

After approximately 36–48 hours of purge cycles, an *Adixen turbomolecular station* with

¹⁵We have not attempted to heat the purge gas to speed up outgassing during dry gas purges. Similarly, no UV light source has been placed inside the vacuum vessel to speed up the desorption.

¹⁶*Agilent Technologies*, Santa Clara, CA.

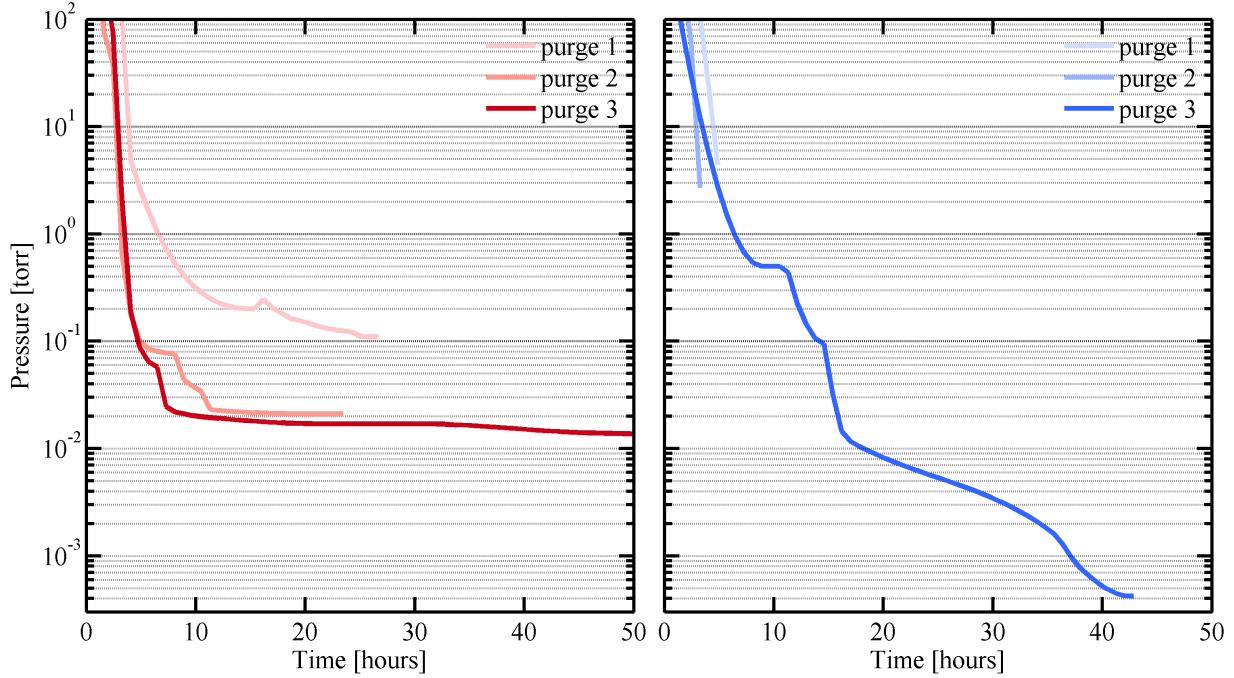


Figure 2.10: *Left:* Pumpdown profiles for the vacuum vessel prior to cooldown in Palestine, Texas, during Run 13. In retrospect, the pumpdown labelled “purge 1” lasted too long. Another dry nitrogen purge was warranted. *Right:* Pumpdown profiles for Run 15 (ongoing). The first two pumpdowns were quick (they barely make it onto the plot). The third pumpdown has the turbo pump placed significantly closer to the vacuum vessel. The change in pump rate is drastic.

an *ATP150 turbo unit* is installed on the vacuum vessel manifold.¹⁷ We pump on the system in this configuration for 5–6 days. During this process, we observe a marked difference in measured vacuum vessel pressure depending on the length and diameter of the pump out manifold. Ideally, the pump should be placed as close to the opening of the vacuum vessel as possible. This reduces the flow impedance to the pump. Figure 2.10 also shows how the pumping rate will increase when we change L/D^4 by a factor of 100–200, where L is the length of the pumping manifold and D is the effective diameter of the cross sectional area; during Run 15, we mounted the turbo directly to the vacuum vessel.

Figure 2.10 shows that the asymptotic pressure after a few days of pumping is of order 1 mtorr. At this point we expect a significant fraction of the residual gas in the vacuum

¹⁷ *Pfeiffer Vacuum Inc.*, Nashua, NH.

vessel to be water vapor. Water molecules at 1 mtorr have a mean free path of approximately 30 cm.¹⁸ This suggests that the system has reached the free molecular regime, defined as the pressure where gas particles collide more frequently with walls inside the vacuum vessel than with each other. As we continue pumping, we expect to reach an asymptotic pressure where we are limited by the outgassing rate of objects inside the vacuum vessel, and not pumping capacity.

2.4.4 Cryogenic Qualification

Once in equilibrium, VCS1 and VCS2, the intermediate radiation shields, stay at approximately 40 and 160 K respectively, while the main tank boil off corresponds to 80 L/day consumption and a flow rate of 38 SLPM.¹⁹ Both flow rates and equilibrium temperatures greatly exceed predictions by thermal models (see Section 2.4.5). The measured flow rate corresponds to 16 day hold time given the net volume of the main tank of 1284 L. As the payload ascends into the stratosphere, the shell of the vacuum vessel will drop in temperature to approximately 270 K [143]. Residual optical loading to the 4 K stage is also significantly reduced at float altitude, compared to sea level. Both effects are expected to cause significant reduction in loading at float. For example, predictions from some thermal models suggest a 200 mW reduction in loading to the main tank going from 300 to 270 K.

Custom Bolometer

Custom bolometers were constructed to characterize the radiative environment in the cryostat during Run 7 (see Figure 2.11). The bolometers consisted of a macroscopic, 158 cm^2 square aluminum 3003 plate stood off from the main tank using four 3.4 cm long G-10 rods with 3.2 mm diameter. The rods were epoxied into metal plates using *Stycast 2850FT*. The plate area corresponded to 1/280 and 1/15 times the area of the main tank and superfluid tank, respectively. Knowing these dimensions, we can estimate the loading to the tanks in the form

¹⁸The mean free path scales inversely with pressure.

¹⁹Five times the flow rate out of a healthy set of lungs.

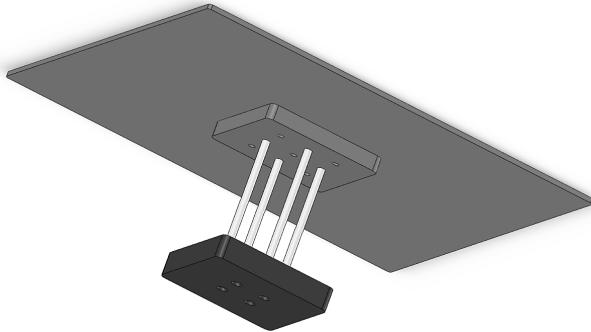


Figure 2.11: A CAD model of the custom bolometer. The 158 cm^2 square aluminum 3003 plate is stood off from the main tank using four 3.4 cm long G-10 rods.

of radiation or helium gas conduction. Two bolometers were manufactured this way. One was enclosed in an aluminum radiation shield with a macroscopic aperture facing the bottom of the main tank to minimize flow impedance for residual gas, including helium atoms.

We estimated the bolometer conductance using data from load curve measurements performed at varying times during Run 7. This analysis suggests that the loading to the shielded and control bolometers was $60 \mu\text{W}$ and $600 \mu\text{W}$, respectively. Naïvely scaling the control bolometer loading we estimate that the loading to the main tank and superfluid tank is 170 and 9 mW, respectively. This is not enough to explain the observed excessive loading to both cryogenic stages. We note that these results were found assuming that the radiative and gas environment is uniform within the cryostat.

In a subsequent cryogenic trial, Run 9, the bolometer was placed at the top of the cryostat. Data from that run suggested no significant change in thermal loading. It is worth noting, if these results were multiplied by a factor of four, we would arrive at a value similar to the discrepancy between predicted and realized loading. It is possible that the accommodation coefficient is different between the walls of the main tank and a relatively polished plate of aluminum used for the bolometer [156]. Because of the significant difference in loading between the shielded and non-shielded bolometer, we decided to place an additional radiation shield over the superfluid tank. The shield is an aluminum 1100 sheet metal formed in the shape of the SFT. Surprisingly, the loading to the SFT was not reduced by more than 20%.

SFT Load Curves

During Run 7, we also performed a load curve on the superfluid tank to estimate the equilibrium loading to the tank at 4 K. With the superfluid tank empty and evacuated, we heated the tank and watched the system reach a new equilibrium temperature.²⁰ Extrapolating the results of this load curve, the measurement suggested 40 mW cooling power from 4 K to the SFT equilibrium temperature of approximately 12 K. This result is consistent with the upper limits on SFT loading of 44 mW, which was established from the maximum superfluid hold time.

2.4.5 Thermal Model

Various thermal models were constructed both during the design phase and as the cryostat was being built. Such models allow us to make predictions about cryogenic performance and bracket the loading through different thermal links. For example, a comparison of early thermal performance with predictions of the thermal model suggest that the heat exchangers are operating with efficiencies close to unity. This is not uncommon [157, 158]. The SPIDER flight cryostat is a complicated system, and any thermal model describing it necessarily makes many simplifying assumptions. The primary indicators of cryogenic performance are hold time and the temperatures of intermediate stages. Intermediate stages should be as cold as possible to minimize loading on detectors through emission from filters [159]. The cryostat is modeled as a system with five isothermal temperature stages: the vacuum vessel shell, two intermediate vapor cooled shields, the main tank, and the superfluid tank. Heat transfer is modeled as conductive or radiative with accurate knowledge of dimensions and material properties.

We use both liquid nitrogen and helium during a cryogenic run. Table 2.5 lists some thermal properties for these two cryogens. Boiling temperatures represent the equilibrium temperature of the main tank. The liquid density at boiling point, as well as latent heat of

²⁰Note that this test was performed before the capillaries were installed (see Section 2.5).

Table 2.5: Thermal properties of ^4He and N_2 [160, 161].

	Boiling Point at 1 atm [K]	Density at B.P. [kg/m ³]	Gas Specific Heat [†] [J/gK]	Latent Heat of Vap. [J/g]
^4He	4.23	124.7	6.32–5.19	20.8
N_2	77.35	806.1	1.12–1.05	199.2

[†]In the range between boiling point and 300 K.

vaporization, are used to calculate hold times, while the specific heat is used to calculate the cooling power of the heat exchangers. The thermal model was originally written in MATLAB. The code was later verified and rewritten in *Python* by Zigmund Kermish. Some of the model results are discussed further in [159].

If the assumptions made in our model are valid, it should be able to predict equilibrium behavior at both 77 and 4 K. Model predictions were originally optimized using data obtained at liquid nitrogen temperatures only. Five free parameters were varied to minimize a Chi-squared penalty function: two describing the heat exchanger efficiency, two describing the effective conductivity of multi-layer insulation, and one describing the emissivity of bulk aluminum (see Appendix A). That model predicted equilibrium temperatures of VCS1 and VCS2 at approximately 40 and 150 K respectively, with at least 60 days of hold time! Measurements of VCS temperatures are quite consistent with these predictions, but the hold time is off by a factor of three; we observe approximately 20 days. It was clear that the thermal model needed revision to account for this discrepancy.

Despite the inability of the model to predict diminished cryogenic performance, it does give a reasonable order of magnitude estimate for the contribution of various forms of heat transfer between temperature stages in the cryostat. Dimensions of all components inside the cryostat are known to high precision. With VCS1 and VCS2 at 40 and 160 K respectively we estimate approximately 300 mW of loading to the main tank through G-10 flexures and about 60 mW of loading through the stainless steel tubing. Note, that an equilibrium flow

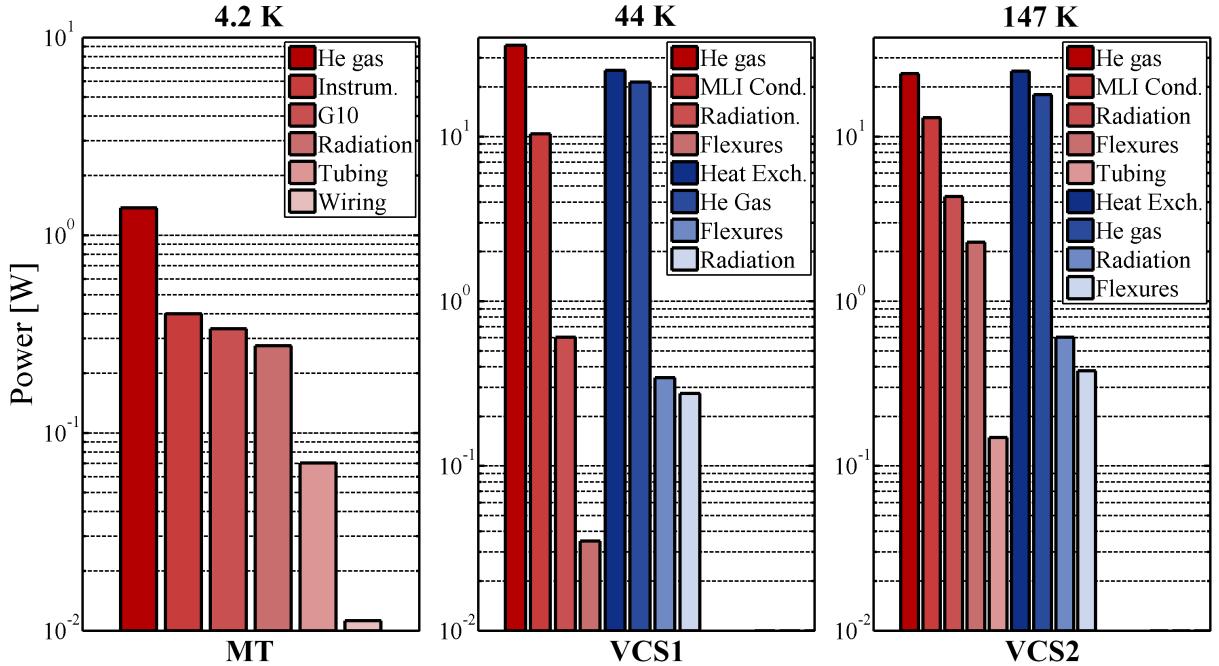


Figure 2.12: An example of predictions from the thermal model with a large component from helium gas conduction. Note that the total loading to the main tank, left, sums up to 2.5 W with VCS1 and VCS2 equilibrating at temperatures comparable to those of Run 14.

rate of 40 SLPM corresponds to a 2.5 W loading to the main tank. See further discussion in Appendix A.

An extension to the thermal model, which incorporates residual helium gas conduction in combination with light leaks, seems to reduce the discrepancy between thermal model predictions and realized values. Figure 2.12 shows predictions from a thermal model which includes a strong helium gas component. Note that the predicted loading to the main tank, approximately 2.5 W, is consistent with current behavior.

2.4.6 Helium Leak

The main tank of the SPIDER flight cryostat has at least one microscopic fissure that leaks helium. The leak was first detected in Longmont, Colorado, where the cryostat was built. We have had mixed success in locating and sealing this leak. The leak is observed at room

temperature, but only through pressurizing the tank with gaseous helium while pumping on the vacuum vessel. This prevents us from localizing the leak while at room temperature. The leak becomes more apparent as the cryostat is cooled to liquid nitrogen temperatures and early attempts at localizing the leak by covering the affected area with liquid nitrogen proved successful. We subsequently applied *Stycast 2850FT* to suspect areas and noticed a marked change in helium backgrounds (compare Runs 4 and 6).

However, with time, the leak rate increased and became harder to localize. The large number of aluminum welds as well as the complex shape and high density of threaded holes were all suspect.²¹ Additionally, bulk aluminum is known to have microscopic fissures extending through the entire thickness of a wall. Prior to Run 15, which is currently ongoing, Stycast was reapplied to suspect areas. This appears to have reduced the helium background levels down to values observed during Run 6. This could suggest that Stycast needs to be applied regularly.

Helium leak checks are performed at various stages in the cooldown of the cryostat as part of standard operating procedures. Although leak checkers report an absolute flow rate of helium atoms, a straightforward interpretation of these values between runs can be difficult.²² During Run 14, our leak checker would report an approximately 1×10^{-6} atm · cm³/s flow of helium atoms. Table 2.4 shows the approximate helium leak rates as reported by a leak detector placed on the vacuum vessel pump out port while pumping in parallel with a turbo station.

Helium atoms reduce the vacuum integrity of our system and can, with sufficient density, prevent the successful operation of the flight cryostat. To reduce the effect of residual helium, we install materials with large surface areas that will adsorb these helium atoms and prevent them from conducting heat in the cryostat. We have chosen to use activated charcoal which

²¹There are 17 m of aluminum welds on the main tank. A fraction of those were laser welded.

²²Leak checkers are great at detecting leaks, but not establishing repeatable leak rate measurements with high precision.

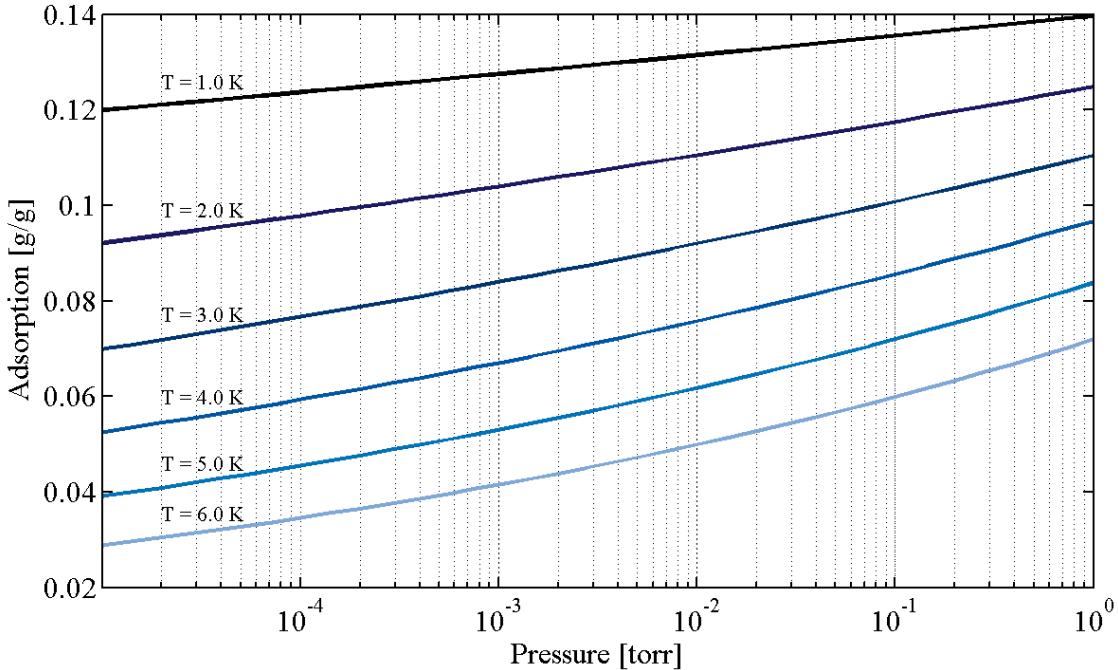


Figure 2.13: Adsorption isotherms of helium on activated charcoal as a function of pressure as calculated using Equation 3 in [162]. The ordinate shows how many grams of helium can be adsorbed on one gram of charcoal.

we epoxy to copper flexures bolted to 4 K surfaces. The procedure results in a 0.1 g/cm^2 surface density of activated charcoal.²³

Figure 2.13 shows an estimate for the adsorption isotherms of activated charcoal as a function of pressure. Activated charcoal is used for this purpose in a wide range of applications [162, 163, 164, 165, 166].

2.5 Capillaries

The SPIDER capillary system provides continuous flow of ${}^4\text{He}$ from the main tank to the superfluid tank. The system is critical for a successful flight because the hold time of the superfluid tank alone, when fully charged, is only 4 days, compared to the likely 20 day flight duration.

²³The charcoal is supplied by *Fisher Scientific*, Pittsburgh, PA. We use 1.40–3.35 mm pellet size in our measurements.

The general design of the SPIDER capillary system is based on a paper by DeLong et al. [167] wherein they describe a two stage ^4He cryogenic system for a dilution refrigerator. In that paper the authors establish an empirical relation between the room temperature flow impedance and the cooling power of the capillaries by studying a few different capillary assemblies. Their findings also suggest that for a fixed pumping speed the equilibrium temperature of a capillary filled pot rises with the throughput of the capillaries. The standard interpretation is that an increase in cooling power is balanced by heat input conducted through a superfluid film. In their measurements, the authors find that the critical power, i.e. the power needed to surpass the cooling capacity of the capillary system, was approximately $4.5\text{ mW}/10^{-4}\text{ mole/s}$. In other words, a flow rate of 10^{-4} mole/s between 4 and 1.8 K provides 4.5 mW of cooling power. This is roughly consistent with a 50% cooling efficiency, assuming a latent heat of evaporation for ^4He of $l = 93\text{ J/mole}$.

2.5.1 Brief Literature Review

In systems that are considered here, liquid helium flows as a result of a pressure differential. With the helium vapor pressure below the superfluid transition point, a non-negligible fraction of liquid will populate the quantum mechanical ground state intrinsic to the system once it exits the capillaries. A number of publications describe capillary systems. Wrubel et al. [168] describe a capillary assembly with four parallel polycarbonate capillaries and a needle valve for controlling the flow impedance. Das et al. [169] discuss a helium evaporator that is continuously fed by a fixed impedance line. Their measurements of room temperature flow impedance and corresponding cooling power seem to agree with the results described in the original capillary paper [167]. Fujiiyoshi et al. [170] describe a similar system with comparable results. It is evident that authors arrive at their preferred flow impedance through trial and error and reliance on empirical measurements, see for example [171, 172].

Some of the more theoretical aspects of helium superfluidity in capillaries are examined in a paper by Koh [173]. A simple model that describes the evaporation rate and film creep

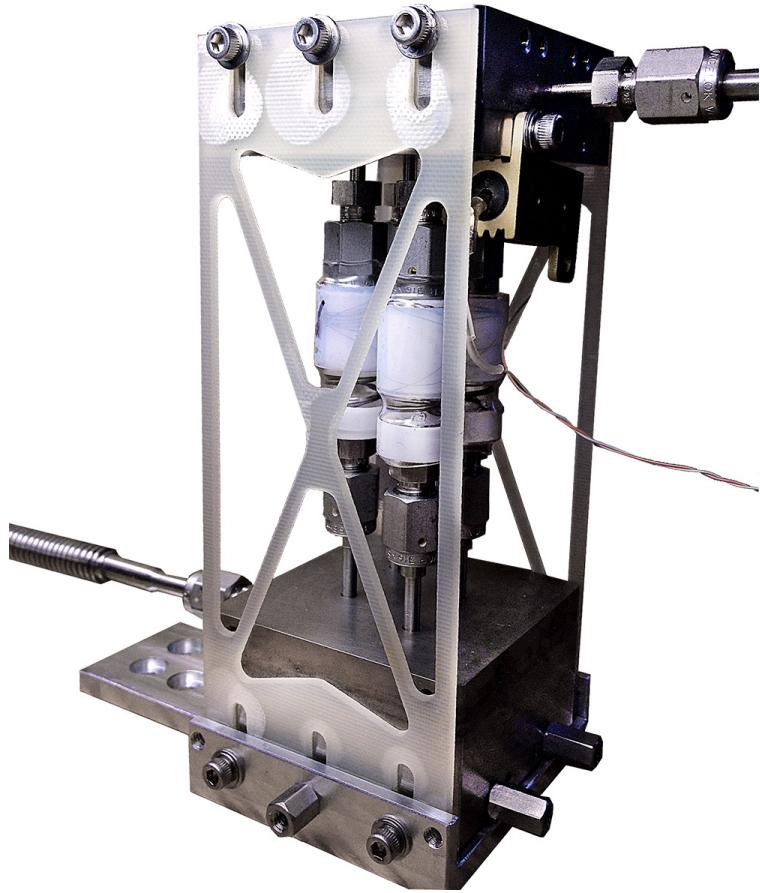


Figure 2.14: The capillary assembly installed prior to Run 10. Four capillaries connect the 4 K main tank box (bottom) to the 1.8 K superfluid box (top). The double volume structure is supported by two 1/32 inch thick G-10 flexures. The thermal load conducted through these flexures is negligible compared to the cooling power from the superfluid helium. Porous stainless steel Mott filters located below each capillary prevent ice and other dirt from entering and clogging this high impedance tubing. Superfluid helium exits the capillaries in the smaller of the two boxes (top) which is connected to the superfluid tank through 6 inch long bellows tubing with a 1/8 inch diameter.

in a sorption refrigerator is discussed in Lau et al. [174]. The model helps guide the design parameters of an orifice that prevents undesirable superfluid creep. A large body of literature has investigated the use of capillaries, which are sometimes called “fixed impedance lines,” in cryogenic applications.

2.5.2 Design

If successful, the capillaries will provide the superfluid tank with liquid helium as long as there is liquid in the main tank. The following list describes the main design criteria for the SPIDER capillary assembly:

- Continuous cooling power of approximately 60 mW to combat steady-state loading to the superfluid tank.
- Superfluid base temperature of at most 1.8 K for effective cycling of the six adsorption refrigerators that supply cooling power to the focal planes.
- No mechanical valves or moving parts.
- Robust operation for at least 50 days.

As described in Section 2.4.3, the steady-state loading on the superfluid volume has been measured under in-flight conditions. These measurements suggest a steady-state loading of 40 mW. Implementing a safety factor of 1.5, we conclude that the capillary assembly must provide at least 60 mW of continuous cooling power to the superfluid tank. The superfluid bath also has to reach a base temperature below 1.8 K. This facilitates effective operations of the ^3He closed-cycle adsorption refrigerators inside each telescope. Cycling the adsorption refrigerators creates an approximate 5 mW transient loading on the superfluid tank which lasts for about an hour. The superfluid tank will have to sustain transients from cycling six refrigerators once every 24–72 hours.

The SPIDER capillary system achieves this using three 35 cm long capillaries wrapped around Teflon spools and silver soldered into 1/8 inch *Swagelok* glands with gender preserving VCR fittings (see Figure 2.14).²⁴ The capillaries connect two approximately 50 mL volumes that are connected to the main tank and the superfluid tank through 1/8 inch bellows tubing, which was custom built to size.

²⁴ *Swagelok*, Solon, OH.

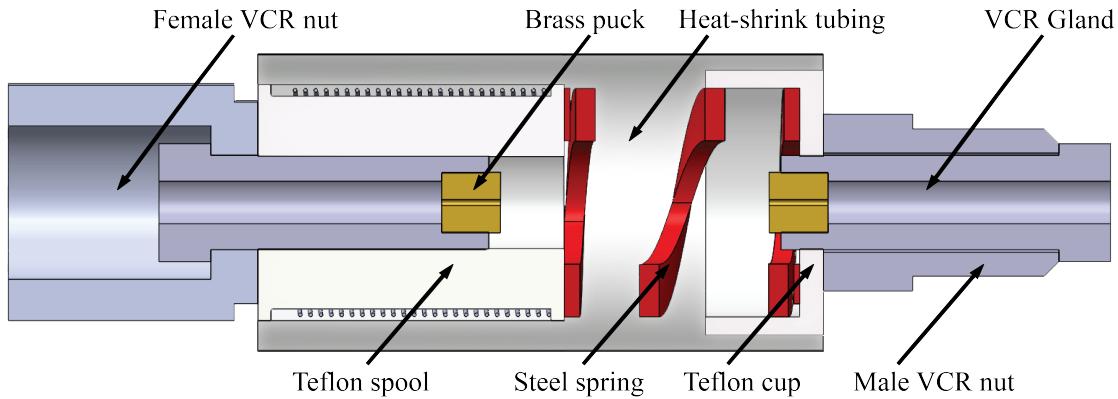


Figure 2.15: A cross section through one of the capsules. The ends of the capillaries are located inside the Swagelok VCR glands, approximately 1 cm from the brass pucks. This way the ends are shielded from mistreatment. After having silver soldered the capillary into the puck on the left, the capillary is threaded through a tiny hole at the edge of the Teflon spool. The spool then slides over the gland after which the capillaries are wrapped around the spool. With only a couple of inches remaining, the capillary is threaded through another hole at the edge of the spool, then through the Teflon cup, after which the other end is silver soldered into the brass puck on the right. The stainless steel spring is then carefully wrapped around the capillary and secured in the Teflon cup. Finally, heat-shrink tubing is positioned around both the Teflon components and spring to shield the capillaries.

The capillary material is extruded 304 stainless steel with a 0.0035 inch inner diameter and a 0.0025 inch wall thickness.²⁵ The dimensions were chosen to resemble one of the assemblies described in the original paper by DeLong et al. [167]. Having three capillaries reduces susceptibility to constrictions from ice slush in the main tank. Failure of one capillary, possibly due to an ice plug, should not affect the performance of the superfluid tank.

The stainless steel capillary lines are silver soldered into Swagelok VCR glands and then wrapped around Teflon spools which are sealed with transparent shrink wrap. Small, custom-made, brass pucks with 0.009 inch diameter holes fit snuggly into the Swagelok glands. A capillary is threaded through one of the pucks before soldering. Post wrapping, the capillary is routed back into the center of the Teflon spools and through stainless steel compression springs that provide structural support to the whole assembly while keeping heat conduction at a minimum. As the capillaries are internal to the springs, there is very little chance that

²⁵Capillaries provided by *Eagle Stainless*, Warminster, PA.

spring compression can pinch the capillaries. The springs allow us to gently modify the overall length of the assembly to fit perfectly between the two glands that are welded to each of the small boxes. The capillaries along with VCR fittings, Teflon components, and springs, are referred to as capsules. Figure 2.15 shows a cross section of an individual capsule. The design allows us to quickly replace capsules if necessary.

Stainless steel *Mott* filters are welded to the inside of the 4 K box so as to intercept large particles before they enter the capillaries.²⁶ The filters are spot welded to the inside of a lid before performing the seam weld that attaches this lid to the box. Unfortunately, the high packing density of Mott filters precludes seam welding. We measured the gap between the filters and the inside surface of the boxes using thickness gauges and found that it was less than 0.001 inch. If correctly installed, the particle capture efficiency of these filters is such that they collect 99.9% of particles whose diameter is larger than 20% of the capillary diameter. This should allow for effective operations of the capillaries without contributing significantly to the overall flow impedance of the system. The superfluid tank and capillary assembly are shown in Figure 2.16.

Measurements show that the design provides approximately 100 mW of cooling power to a superfluid volume while conducting only 2 mW between the two temperature stages. The net cooling power can be changed by simply altering the length of the capillaries in a way that does not require any other change to the design. This allows us to quickly arrive at an optimal cooling power by having a collection of interchangeable capsules.²⁷

2.5.3 Experimental Results

Superfluid is characterized by inviscid flow of zero entropy liquid with almost infinite thermal conductivity. It is observed when ^4He is cooled below 2.17 K, referred to as the λ -point. This phenomenon was first discovered in 1937 by Kapitsa, Allen, and Misener [175], and later

²⁶Filters purchased from *Mott Corporation*, Farmington, CT. We use media grade 20 filters.

²⁷A great number of backup capsules were constructed with wonderful care by Princeton undergraduates Will Taylor and Charles J. Titus, class of 2014.

garnered a phenomenological description for which Landau received the Nobel prize [176].²⁸ The elegant theory of superfluidity is contrasted by a medley of empirical evidence.

Capillary systems are commonly quantified by their cooling capacity, or critical power, Q_{crit} , which is related to the flow rate through the capillaries. The critical power can be estimated through the capillary impedance factor, an extrinsic quantity that can be measured at room temperature [167]. A pressure differential, ΔP , is set up between the two ends of a capillary system and the flow rate measured.²⁹ The impedance factor is then

$$Z = (1/\eta)\Delta P/\dot{V}, \quad (2.5)$$

where η is the dynamic viscosity of the fluid and \dot{V} is the volumetric flow rate. The warm flow impedance can then be compared to the cooling power of the capillaries measured at operational temperatures. For the SPIDER capillary assembly, we measure the flow impedance by pressurizing the main tank to 15.7 psia, using either nitrogen or helium. We then evacuate the superfluid tank and observe the pressure rise in the calibrated volume over hour timescales. The room temperature dynamic viscosity of helium, $\eta(\text{He}) = 20.0 \text{ kg/m/s}$, is larger than that of nitrogen, $\eta(\text{N}_2) = 18.6 \text{ kg/m/s}$. We have to account for this in our calculation of flow impedance. It is useful to note that dynamic viscosity is generally proportional to \sqrt{T} [177].

The critical power is measured by applying heat to an empty volume with a steady supply of superfluid helium. The base temperature of the superfluid bath increases with heat input until the superfluid flow is unable to compensate. A temperature runaway effect is then observed where additional heat lifts the temperature of the superfluid enclosure past the λ -point at a fast rate. The critical power can also be estimated from the equilibrium flow rate out of the superfluid tank. The measured critical power can be compared to the warm flow impedance to establish an empirical relation between the two parameters.

It seems difficult to estimate how superfluid helium flow rates depend on the geometry of

²⁸Kapitsa won a Nobel prize for his work in low-temperature physics. He shared that prize with Penzias and Wilson.

²⁹For example, by watching the pressure increase in a calibrated volume.

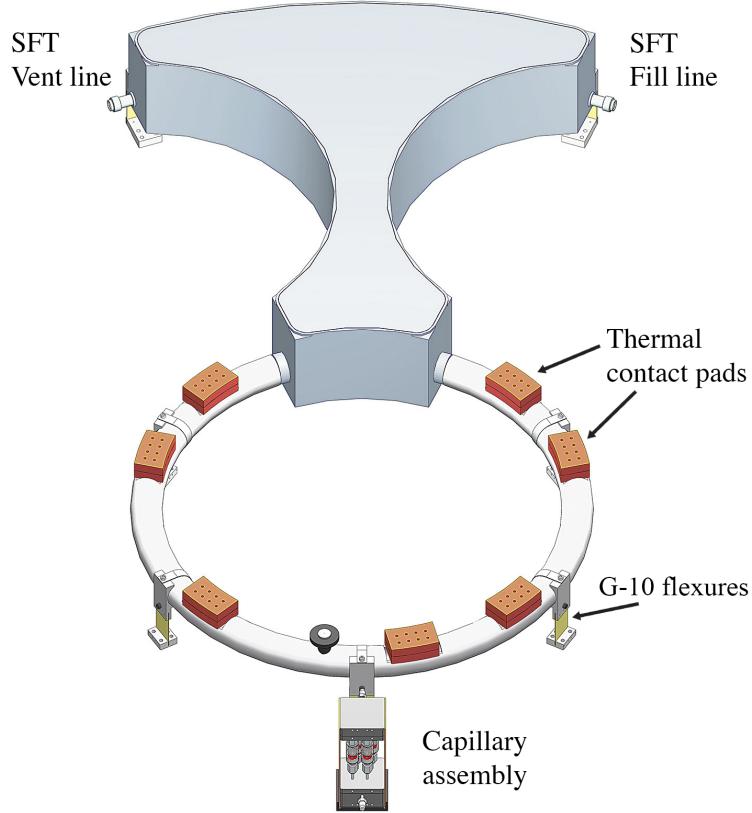


Figure 2.16: The superfluid tank and capillary assembly are mounted to the bottom of the main tank (not shown), which in this case would be rotated at a 45 deg angle with respect to the plane of the page. Thermal contact pads are welded into a ring like structure that sits below the main volume. The fill and vent lines exit at the top to the right and left respectively. The net volume of the superfluid tank is 16 L.

the capsules. Ideally, the Hagen-Poiseuille equation relates the pressure drop to the length and the diameter of the capillaries as follows

$$\Delta P = \frac{128\eta L \dot{V}}{\pi d^4}, \quad (2.6)$$

where ΔP is the pressure drop over the cylindrical tube with diameter d and length L such that $L \gg d$. Complications arise since the superfluid transition happens somewhere within the capillaries. In the two-fluid model, the above equation is still valid for the normal component whose flow should remain laminar. However, the superfluid component is likely to become turbulent. In this case, the appropriate expression for the heat flow is derived by

Schotte [178]. The two-fluid model in relation to capillaries is discussed further in Lages et al. [179].

Two distinct capillary assemblies have been run in the SPIDER flight cryostat. The current design was installed prior to Run 10. The size and complicated geometry of the superfluid tank, chosen to accommodate six telescopes, heat straps, and the capillaries themselves impedes thermal modeling. Both conduction and radiation contribute significantly to the total thermal budget of the superfluid tank. Heat is conducted through the stainless steel vent and fill lines and the G-10 flexures that suspend the superfluid tank from the main tank. Radiation from the inner vapor cooled stage and any light leaks from warmer stages also contribute to the loading at a significant level. Under equilibrium conditions, the enthalpy of the system is kept in balance by the evaporation of liquid helium. We can write the steady-state requirement as follows

$$\dot{m}H_{\text{vap}}(T_{\text{SFT}}) \approx \sum_i \int_{T_{\text{SFT}}}^{T=4K} C_i(T') dT' + \sigma_{\text{SB}} A_{\text{eff}} T_{\text{VCS1}}^4, \quad (2.7)$$

where \dot{m} is the mass flow rate out of the superfluid tank, H_{vap} is the temperature-dependent enthalpy required to vaporize a unit mass of ^4He , C_i represents the conduction through one of the stainless steel or G-10 thermal paths, and the last term on the right-hand-side represents the radiative loading to the SFT, with A_{eff} corresponding to the effective area of the SFT-VCS1 system. From this we see that the equilibrium temperature of the SFT depends on the flow rate. As the vapor pressure of the helium bath is set by the pumping capacity and the impedance of the capillaries, so will the equilibrium temperature of the liquid bath. If the impedance of the capillaries is sufficiently low, the flow from the main tank will negate evaporative liquid loss and an equilibrium state will be reached.

Table 2.6 shows the measured properties of the capillary assemblies for the six cryogenic runs performed with them installed in the SPIDER flight cryostat. During Run 9 we found that the cooling power of the capillaries was not sufficient to sustain a reasonable amount of superfluid. This was surprising, as we expected approximately 40 mW of cooling power

Table 2.6: Measured properties of various capillary assemblies. For Run 9 the cooling power of the capillaries was less than the total equilibrium heat input to the superfluid tank. The throughput of the capillaries was greatly increased between Runs 9 and 10 such that the pumping speed limited the equilibrium temperature of the superfluid tank. A flow restriction formed during Run 12, which limited the superfluid flow rate.

Property	DeLong	R9	R10	R11	R12	R13	R14
Impedance, $Z \times 10^9$ [cm $^{-3}$]	670	240	36	47	41	49	59
Critical power, Q_{crit} [mW]	15	18	154	N/A	N/A	N/A	N/A
Eq. flow, f_{SFT} [SLPM]	0.44	0.45	3.1	2.5	0.65	2.2	1.7
Zf_{SFT} [10 9 SLPM/cm 3]	295	109	112	118	27	107	100

based on the warm flow impedance measurement and the results of DeLong et al. [167]. Prior to Run 10, we chose to increase the throughput of the capillaries assuming that the critical power would greatly exceed the heat input to the superfluid system. At that point, we found that the pumping capacity of our scroll pump was limiting the minimum temperature of the system. Another, burlier, pump was acquired before Run 11 and one of the capsules was removed to reduce the throughput and the loading to the main tank. Run 11 performance was ideal.

Some flow restriction formed during Run 12, as we observed greatly reduced throughput given expectations from the warm flow impedance. Even at decreased flow rate, we were still able to efficiently cool both telescopes installed for that run down to 300 mK. The warm flow impedance measured prior to Runs 14 and 15 are significantly different from Runs 11–13. It is possible that one of the three capillaries has been damaged in a way which permanently restricts flow.

We note that warm flow impedance is generally a great predictor for superfluid flow. This is clear from the consistency of the product $Z \times f_{\text{SFT}}$, shown in the last row of Table 2.6.

For every cryogenic run, the cooling power has been measured by the method described above. We find that using the relation described in [167] our measurement of the warm flow impedance overestimates the critical power by a factor of three, but in a repeatable

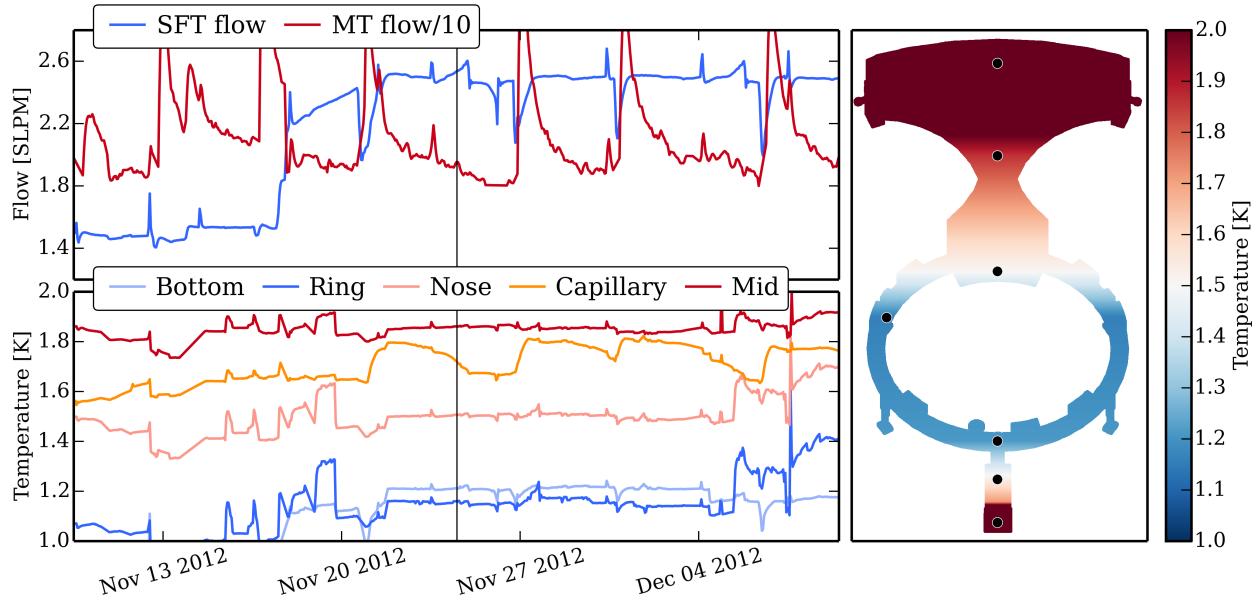


Figure 2.17: *Left above*: Flow rate out of the SFT and MT during a four week period during cryogenic testing in Princeton, Run 11. *Left below*: Temperature of various parts of the SFT as a function of time for the same time period. *Right*: A crude visualization of the temperature profile of the SFT at the time indicated by vertical line on the left-hand side plots. The black circles represent locations of thermometers.

manner. It is possible that the designs differ, for example, with regards to geometry, in some crucial way. Their measurement of the flow impedance is performed with 3.1 torr pressure differential whereas our measurement is performed with roughly 800 torr driving the flow. Direct measurements of critical power were only performed for Runs 9 and 10. We know, however, that the results scale linearly with equilibrium superfluid flow. Thermometers along the superfluid tank allow us to monitor the temperature profile of the superfluid tank, and to some extent, gauge the superfluid liquid level inside the tank (see Figure 2.17). The equilibrium flow rate is continuously monitored using flow meters calibrated for low flow rates which we install on the exhaust of our scroll pumps.³⁰

³⁰ Omega Engineering Inc., Stamford, CT.

2.5.4 Performance Characterization

The capillaries are at risk of becoming plugged, most likely with nitrogen ice, in the interim between liquid nitrogen pre-cooldown and the initial liquid helium fill. The process is inherently risky as residual nitrogen can solidify in the presence of liquid helium; it is fair to assume that full risk mitigation is impossible without moving parts. So far, we have not devised an operating procedure with a 100% success rate.

Given the long response time of the large cryostat, and the tight schedule afforded by McMurdo ballooning, flow constriction in the capillaries poses significant risk to the schedule. In the worst case scenario, plugged capillaries could require heating the cryostat back up to room temperature and breaking vacuum, causing a 2–3 week slip in schedule at minimum. So far, however, we have been successful in clearing plugged capillaries by suspending liquid helium cool downs and heating the capillaries (which are well isolated from the main tank) up to 300 K. This procedure requires approximately 2–3 days, and so far, it has a 100% success rate.

The initial cool-down stage involves filling the main tank with few hundred liters of liquid nitrogen while pressurizing the superfluid tank to 6 psig with gaseous helium and maintaining approximately 1 SLPM of flow out of a pressure regulator on the SFT vent. Once all liquid nitrogen has boiled out, the main tank is evacuated. The system is then backfilled with gaseous helium at room temperature. We maintain the capillary assembly at room temperature during the initial cool-down phase, approximately 270 K rather than 90 K. This should ward off flow restrictions. This has to be done carefully, however, as overheating might damage the capillaries.

Immediately following helium backfilling of the main tank we perform a capillary flow test to determine the status of the capillaries before initial liquid helium transfer. The flow test is performed with an approximately 90 K main tank pressurized to 15–16 psia with gaseous helium. The superfluid tank is then evacuated with a scroll pump attached to the SFT vent line manifold which is connected to an absolute pressure gauge. After evacuation, we valve

off the scroll pump and observe pressure rise in the 16–18 L volume of the superfluid tank and manifold. The pressure increase is due to flow through the capillaries, driven by the roughly 16 psi pressure differential. The rate of pressure increase as measured on the manifold depends in a non-trivial way on temperature of the main tank, capillaries, superfluid tank, and the room. As these sub-systems can have different temperatures at the time that the flow test is performed, we have to correct for this in estimating the performance of the capillaries.

Informed by the data at hand we can derive a dimensionless quality factor, Q , that scales linearly with observed superfluid flow,

$$Q \equiv \frac{\eta(T_{\text{cap}})}{\Delta P} \frac{V_{\text{SFT}} + V_{\text{man}}}{T_{\text{eff}}} \left(\frac{\Delta p}{\Delta t} \right), \quad (2.8)$$

$$T_{\text{eff}} \equiv \frac{\alpha T_{\text{SFT}} + T_{\text{room}}}{1 + \alpha}, \quad (2.9)$$

where V_{SFT} and V_{man} represent the volume of the superfluid tank and the outside manifold, respectively, ΔP is the pressure differential between the main tank and superfluid tank, η is the effective dynamic viscosity of gaseous helium traveling through capillaries at temperature T_{cap} , T_{SFT} is the average temperature of the superfluid tank, T_{room} is the room temperature, and $\Delta p/\Delta t$ is the rate of pressure increase in the superfluid tank during flow test. Finally, α is a parameter that we use to fit our model to the realized behavior. In our system, we find that setting $\alpha \approx 0.5$ maximizes the correlation between Q and realized flow rates.

This quantity is calculated immediately following a capillary flow test at liquid nitrogen temperatures and just prior to the first proposed liquid helium fill of a given run. It is used to estimate the health of the capillary assembly and determine whether the cryostat can be filled with liquid helium.³¹

³¹The initial helium fill is time consuming. The process normally starts in the early morning and is rarely over at a reasonable hour. The ability to quickly ascertain the health of the capillaries is essential.

2.5.5 Launch Configuration

Flow out of the main tank is guided through the vapor cooled shields. An ensemble of *Tavco* absolute pressure flow regulators are mounted on the end of the main tank and VCS vent lines.³² These regulators set the pressure in the main tank. A couple of Tavcos set to open at 13.5 psia are mounted on the end of the VCS vent line. We use two Tavcos to split the flow and therefore reduce the effective cooling power to each Tavco. Given their intrinsic flow impedance, we expect that the main tank will equilibrate at approximately 14.7 psia (1 atm) at float. Another Tavco set to crack at 17.5 psia is installed on the MT. It will only open in the event of MT overpressurization. This ensures that all flow out of the main tank is sent through the VCSs.

Fill lines of the superfluid and main tank are capped with custom-made burst disk assemblies that are set to burst at an absolute pressure of 25–32 psia. These were made by *Redstone Aerospace*, the same company that built the flight cryostat. The design involves an evacuated volume and two layers of burst disks. The inner disk cracks at 25–32 psi differential pressure and is backed by vacuum while the outer is set at a lower bursting pressure such that it will open as soon as the inner one does.

Atmospheric pressure changes rapidly during ascent. Such dramatic pressure changes can affect the thermal performance of the superfluid tank. Figure 2.18 shows the pressure profiles as a function of time for a 37H balloon, the balloon used for SPIDER, as well as the realized ascent profile for the EBEX and B2K experiments.³³ Based on the measured flow rate out of the tank during boil off procedures, we estimate that the superfluid tank has about 2.0 L of superfluid helium at equilibrium. An hour before launch, the scroll pump that is normally used to pump on the tank will be removed. A compact, battery powered, diaphragm pump, which is mounted on the cryostat, will then start pumping on the superfluid tank.³⁴ As this pump is not capable of sustaining 2 SLPM flow at 10 torr base pressure, the pressure in the

³²*Tavco, Inc.*, Chatsworth, CA, e-mail: tavcoinc@aol.com, fax: +1-818-998-8391.

³³EBEX and B2K data courtesy of Shaul Hanany and Bill Jones respectively.

³⁴*KNF Neuberger Inc.*, Trenton, NJ.

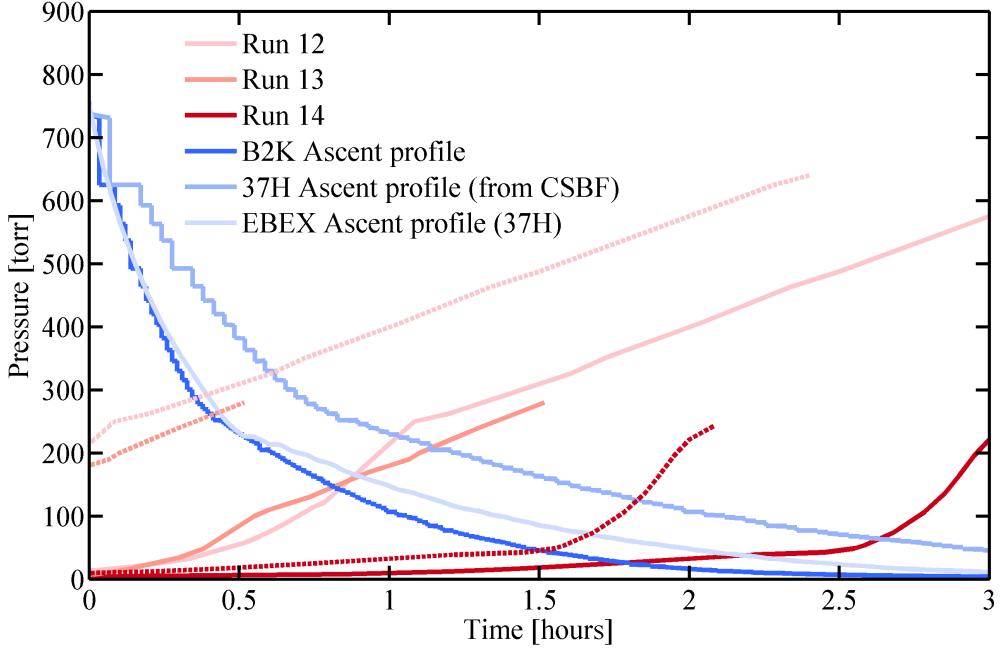


Figure 2.18: Atmospheric pressure as a function of time from launch. A typical flight reaches float altitude after approximately 3 hours. The plot also shows the pressure profiles of an emptied and valved off SFT as measured during cryogenic testing, see Run 12 and 13 curves. The Run 14 curve shows the pressure profile when, instead of emptying the SFT, we start pumping on the volume with the diaphragm pump. The dashed lines represent a one hour delay between the time when the test begin and the time the payload is launched. The three tests have been scaled to a common flow rate of 2.0 SLPM. In all cases, 10–20 torr pressure oscillations at approximately 2 Hz were observed once the pressure reached about 50–100 torr. This coincides with a change in the slope of the pressure profile.

SFT will begin to rise gradually, see Run 14 test in Figure 2.18.³⁵ During this period, the payload will also begin its ascent. Approximately 2.5 hours after the manifold reconfiguration, the helium vapor pressure will exceed the λ -point. Shortly thereafter, the pressure profile will change, see red curve in Figure 2.18. Of course, the atmospheric pressure is also dropping with altitude. When the atmospheric pressure becomes lower than the helium vapor pressure in the superfluid tank, a commandable valve will open and allow the atmosphere to pump

³⁵During Runs 12 and 13, instead of using the diaphragm pump, we emptied the SFT (using heaters) and waited for the pressure in the tank to drop to approximately 1 torr. We then valved off the SFT while measuring the pressure rise as a function of time. The pressure profiles in Figure 2.18 show the difference between emptying and valving off (Runs 12 and 13) and using a low power diaphragm pump (Run 14).

on the superfluid tank. The vapor pressure in the SFT will asymptote to the approximately 10 torr float pressure. We expect this will happen no later than three hours after launch.

In order to limit contamination from residual atmosphere, such as water vapor, we want to command the valve to open at as low pressure as possible. On the other hand, a violent pumpdown due to large pressure gradients and low flow impedance could have negative effect on the superfluid stage and therefore the sub-Kelvin stages in each telescope. We do not know the ideal balance between minimizing pressure oscillations and preventing constrictions due to residual atmosphere.

2.6 The Telescopes

The SPIDER telescopes are based on an optical design that is almost identical to that of the BICEP1 and BICEP2 telescopes [180, 181]. Each telescope is a monochromatic telecentric system (see Figure 2.19). The azimuthal symmetry of telecentric optics can in some cases help reduce systematics that affect experiments lacking such symmetries. The compact aperture allows for extreme baffling and aggressive sidelobe suppression. Finally, because of the small aperture, the entire optics assembly can be cooled to 4 K up to and including the cryogenic waveplate. Below the waveplate are two high density polyethylene (HDPE) lenses, the objective and eyepiece. These form a f/2.3 and f/2.6 refracting system at 94 and 150 GHz respectively. The two lenses are anti-reflection (AR) coated using Teflon, which has an index of refraction $n_{\text{Teflon}} = 1.57$. Light enters the telescope through windows with a 330 mm diameter.

The telescope frame is constructed from carbon fiber structural members epoxied into aluminum fixtures and mounted to aluminum rings at strategic intervals. The carbon fiber provides a conductive and lightweight structure with outstanding rigidity. Measurements show that the carbon fiber supports have a relatively high ratio of elastic modulus to thermal conductivity when compared to other polymeric and composite structures [155]. A flexible copper shim heat strap is routed internally, creating a thermal link between the focal plane

and the still of the adsorption refrigerator that is mounted to the baseplate. This heat strap ensures that the focal plane is cooled to 300 mK. Another heat strap connects the superfluid temperature stage to the magnetically shielding spittoon.

Both the objective and eyepiece curvatures follow a conic equation

$$z = \frac{cr^2}{1 + \sqrt{1 - (1 + k)c^2r^2}}, \quad (2.10)$$

with z describing the height profile of the lens and c and k representing the curvature and conic constant, respectively.³⁶ The parameters of the conic equation were optimized to minimize aberration in the focal plane using *Zemax*.³⁷ The Fraunhofer far field defined as D^2/λ corresponds to 36 and 54 m for 94 and 150 GHz respectively.

A 280 mm cold stop is mounted in front of the objective inside the telescope tube and cooled to 2 K. It is the coincidence of this optics stop with the objective which makes the system telecentric. Cooling the optical stop from 4 to 2.0 K was found to significantly reduce detector loading. This is understandable as approximately 25% of the antenna beam pattern terminates on the cold stop. The location of the stop corresponds to a relatively aggressive 6 dB edge taper. The stop is part of a blackened optics sleeve which spans the majority of the space between the two lenses. The sleeve is blackened using a mixture of *Stycast 2850FT* (cryogenic epoxy) and fine stainless steel powder. A set of eight blackened baffle rings is mounted on the inside of the optics sleeve, extending approximately 1 cm from the cylinder wall. These rings were implemented when measurements suggested 20% polarized reflectivity of the blackened surface. The baffle rings increase internal reflections and the effective absorptivity of the cooled optics sleeve.

A 4 (6) icm hot-pressed filter is mounted on top of the magnetically shielding spittoon of a 94 (150) GHz telescope.³⁸ At 120 and 180 GHz, these filters represent the lowest frequency

³⁶Note that c and k are dimensionless constants. The curvature of the aspheric lens increases with the absolute value of c while k defines the conic section. See [181] for further discussion.

³⁷*Zemax, LLC*, Redmond, WA

³⁸Use 30 GHz/icm to convert from icm units to GHz.

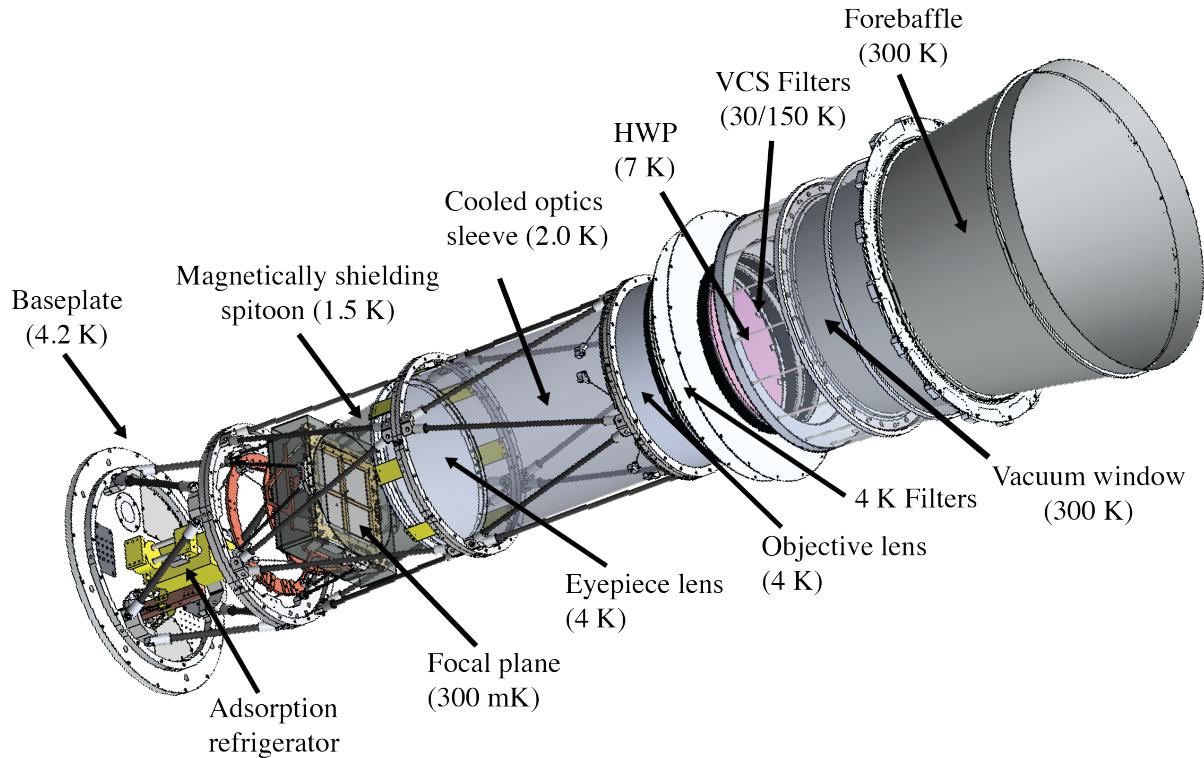


Figure 2.19: A CAD model of the entire telescope assembly, measuring 2.02 m end-to-end. The objective, 4 K filters, and half-wave plate are not clearly visible from this viewing angle. The VCS filters are omitted to reveal the surface of the HWP.

cutoff in the SPIDER filter stack. Another collection of filters and IR shaders are mounted above the objective and on both VCS1 and VCS2. These filters greatly reduce the radiative load that would otherwise saturate the bolometers and load the 4 K main tank. The SPIDER filter configuration is discussed in detail in [159].

Spurious signals can be generated in our timestreams from magnetic pickup in superconducting readout devices as the payload rotates in Earth’s magnetic dipole field. The telescopes use high magnetic permeability materials to reduce susceptibility to magnetic disturbances. Informed by finite element modeling, the focal plane, as well as the SQUID readout system are surrounded by a 1.8 K spitoon made of *Amumetal 4K* (A4K), a proprietary nickel and iron alloy.³⁹ The SQUID mux chips are further shielded by a combination of the niobium backshort, a niobium enclosure, and A4K sleeves. See Runyan et al. [182] for a comprehensive

³⁹Amuneal Manufacturing Corp., Philadelphia, PA.

description of the SPIDER magnetic shielding design. The telescope tubes are surrounded by a two-layer concentric A4K or *Cryoperm 10* magnetic shield assembly.⁴⁰ The high magnetic permeability materials essentially funnel magnetic field lines and reduce magnetic flux density at the focal plane.

The 12 kg concentric cylinder shield assemblies represent the first wave of resistance against unwanted magnetic fields. The theory of magnetic shielding in relation to cylinders is discussed extensively in [183]. The dimensions of the shield assembly are greatly constrained by the size of the telescope and the cryostat, such that only the thickness and number of shield layers could be varied as design parameters. As for all flight hardware, we also needed to limit the mass of the shields. We chose a 1 mm wall thickness two-layer assembly with 25 mm differences in outside diameters. These concentric cylinders are 110 cm long, corresponding to the length of the main tank, and the diameter of the inner shield is 38 cm. A two-layer design was found to give superior shielding to mass ratio over a three-layer system.

2.7 Half-Wave Plate

A simple demonstration with polarizing sunglasses should convince people that many sources of light are at least partially polarized. Unavoidably, objects inside our telescopes will emit or reflect partially polarized light. Similarly, systematics of various origin are likely to create spurious polarized signals in our detector timestreams. We use a rotating half-wave plate (HWP) to mitigate these effects. A half-wave plate is a carefully selected piece of birefringent material that can rotate polarization by a known amount. A slab of crystal is cut in the plane of its principal axis, such that two indices of the symmetric dielectric tensor, n_α and n_β , and the thickness of the plate, d , fulfill

$$2d(n_\alpha - n_\beta) = \lambda_0, \quad (2.11)$$

⁴⁰Cryoperm 10 was superseded by A4K during the build process.

where λ_0 is the band-center wavelength of the telescope. Rotating this plate by a known amount allows us to separate internal non-idealities observed as polarization from those which originate further skyward in the optical chain. The SPIDER HWP is extensively described in [184].

Birefringence is produced by asymmetric crystal structure which causes opposite polarizations to move at different velocities through the material. This effectively rotates the polarization of the light as shown in Figure 2.20. Birefringent sapphire, with an index of refraction of $n \approx 3$, was chosen for the SPIDER HWPs. This high refractive index would cause substantial reflections if it were not for Anti-Reflection (AR) coatings which are bonded to the sapphire. We use quartz and Cirlex (a polyamide) as AR coats for 94 and 150 GHz, respectively.

Half-wave plates are used by a large number of millimeter and sub-millimeter receivers, including ABS, EBEX, BLAST-Pol, and Polarbear [185, 109, 186, 187]. The addition of a HWP naturally complicates any optical system and considerable work has gone into modeling and characterizing HWP effects [188, 189, 190]. O'Dea et al. describes simulations that probe how HWP non-idealities could affect SPIDER [130].

The SPIDER half-wave plate is mounted on a gold-plated Invar ring with compression clips holding the sapphire in place (see Figure 2.20). The plate rotates on three bearings, one of which is spring-loaded so the delicate plate can be easily removed from the assembly. A cryogenic stepper motor spins a worm-gear that couples to the main gear of the plate. Optical encoders read out the angular orientation of the sapphire plate. The half-wave plate is mounted on top of the main tank, such that it equilibrates at approximately 7 K. An estimate described in [184] finds that the AR coating should equilibrate at about 2 K above the temperature of the plate. We therefore expect the AR coating and sapphire plate to stay below 10 K during normal operations. Each HWP assembly weighs 6.5 kg. The nominal scan strategy involves stepping the HWP by 22.5 deg every 24 hours. This effectively rotates Q into U in an arbitrary coordinate system.

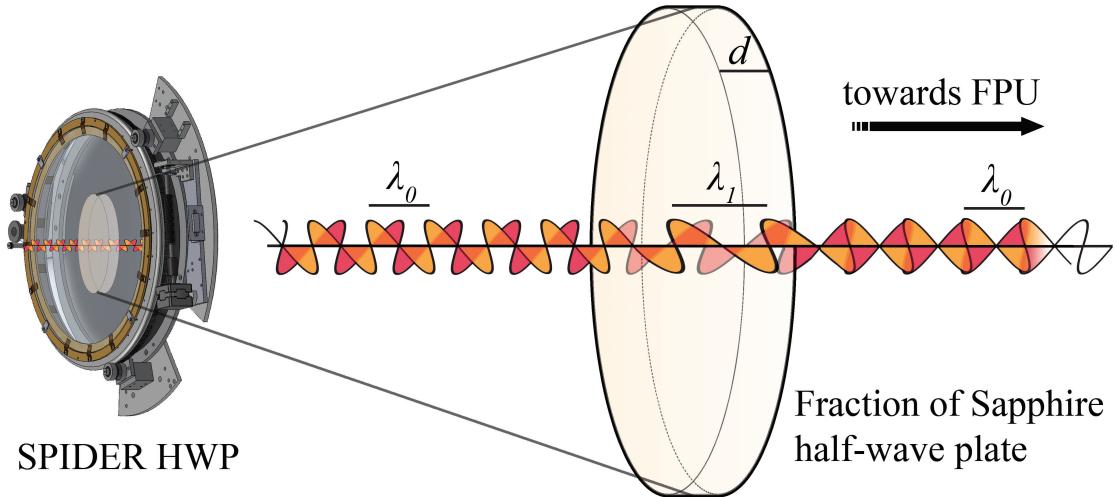


Figure 2.20: A birefringent crystal rotates the polarization of transmitted light. The exact effect depends on the angle of incidence, thickness of crystal, and frequency of light. The figure shows how the sapphire in the SPIDER HWP rotates the polarization vector of incoming radiation by 90 deg. Orthogonally polarized waves that were originally in phase are completely out of phase when exiting the plate.

To understand how a half-wave plate rotates polarization, and therefore reduces the effect of spurious polarization, we present the following derivation. For this discussion, the Stokes parameters Q and U are defined as (see Equation 1.26):

$$Q = \langle |\mathbf{E}_x|^2 \rangle - \langle |\mathbf{E}_y|^2 \rangle, \\ U = 2\text{Re}(E_x E_y^*), \quad (2.12)$$

with \mathbf{E}_x and \mathbf{E}_y corresponding to the electric fields measured in the tangent plane and $\langle \cdot \rangle$ representing the time-averaged field. Let us assume that a polarized signal, originating on the sky, is being observed by one of SPIDER's detectors as a pure Q polarization. In other words, $Q = Q_0$ and $U = 0$. Defining a right-handed coordinate system, $\hat{\mathbf{x}}\text{-}\hat{\mathbf{y}}$, on the focal plane, with its axis parallel to the symmetry axis of the phased antenna array, an arbitrary electric field at that point in time can be represented by

$$\mathbf{E}^f(\theta = 0) = E_x e^{i\omega t} \hat{\mathbf{x}} + E_y e^{i\omega t} \hat{\mathbf{y}}, \quad (2.13)$$

with θ indicating the orientation of the half-wave plate. The requirement of pure Q polarization is then satisfied if $E_x = \sqrt{Q_0}$ and $E_y = 0$. The half-wave plate is a uniaxial crystal with a fast axis and a slow axis, meaning that two disparately polarized electromagnetic waves propagating through the plate do so at different velocities. The thickness of the plate is chosen such that orthogonally polarized electromagnetic waves propagating through the plate have been phase shifted by $2 \times \pi/2$ after exiting (see Figure 2.20). Just after entering the plate, the two electric fields must have been out of phase by a half-wave, suggesting

$$\mathbf{E}^i(\theta = 0) = E_x e^{i\omega t + \pi/2} \hat{\mathbf{x}} + E_y e^{i\omega t - \pi/2} \hat{\mathbf{y}}, \quad (2.14)$$

where $\mathbf{E}^i(\theta = 0)$ represents the electric field of the light right as it enters the plate. If we assume that the fast and slow axis of the HWP are represented by another right-handed coordinate system with unit vectors $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$. Then, at this time, $\hat{\mathbf{a}}$ is aligned with $\hat{\mathbf{x}}$, and more generally we have $\cos \theta \equiv \hat{\mathbf{a}} \cdot \hat{\mathbf{x}}$. Assume that the half-wave plate is now rotated by some angle θ , such that the polarization axis of the plate are no longer aligned with the $\hat{\mathbf{x}}\text{-}\hat{\mathbf{y}}$ coordinate system. Decomposing the two electric fields along the polarization axis of the half-wave plate, $\hat{\mathbf{a}}\text{-}\hat{\mathbf{b}}$, it is clear that the electric field, as it has just entered the plate is

$$\begin{aligned} \mathbf{E}^i(\theta) = & \left(\cos \theta E_x e^{i\omega t + \pi/2} - \sin \theta E_y e^{i\omega t - \pi/2} \right) \hat{\mathbf{a}}, \\ & + \left(\sin \theta E_x e^{i\omega t + \pi/2} + \cos \theta E_y e^{i\omega t - \pi/2} \right) \hat{\mathbf{b}}, \end{aligned} \quad (2.15)$$

where we have simply applied the rotation matrix for counter-clockwise rotation in the $\hat{\mathbf{x}}\text{-}\hat{\mathbf{y}}$ Cartesian plane,

$$\mathcal{R}(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}, \quad (2.16)$$

to the electric field $\mathbf{E}^f(\theta = 0)$ as defined in Equation 2.13. As the electric field exits the plate, the components of the field, as decomposed along the $\hat{\mathbf{a}}$ and $\hat{\mathbf{b}}$ axis, have picked up a phase

shift, so that

$$\begin{aligned}\mathbf{E}^f(\theta) &= (\cos \theta E_x e^{i\omega t} - \sin \theta E_y e^{i\omega t - \pi}) \hat{\mathbf{a}}, \\ &+ (\sin \theta E_x e^{i\omega t + \pi} + \cos \theta E_y e^{i\omega t}) \hat{\mathbf{b}}.\end{aligned}\quad (2.17)$$

Note, that this is the same phase shift as we assigned in the step described by Equation 2.14. Finally, we can rotate back into the $\hat{\mathbf{x}}\text{-}\hat{\mathbf{y}}$ coordinate system by performing the inverse rotation using $\mathcal{R}(-\theta) = \mathcal{R}^{-1}(\theta)$ and remembering that $\cos(2\theta) = \cos^2 \theta - \sin^2 \theta$ and $\sin(2\theta) = 2 \cos \theta \sin \theta$. To save space we replace \sin and \cos with s and c respectively. This gives

$$\begin{aligned}\mathbf{E}^f(\theta) &= (c^2(\theta) E_x e^{i\omega t} - c(\theta)s(\theta) E_y e^{i\omega t - \pi} + s^2(\theta) E_x e^{i\omega t + \pi} + c(\theta)s(\theta) E_y e^{i\omega t}) \hat{\mathbf{x}} \\ &+ (-c(\theta)s(\theta) E_x e^{i\omega t} + s^2(\theta) E_y e^{i\omega t - \pi} + c(\theta)s(\theta) E_x e^{i\omega t + \pi} + c^2(\theta) E_y e^{i\omega t}) \hat{\mathbf{y}} \\ &= [(c^2(\theta) - s^2(\theta)) E_x + 2c(\theta)s(\theta) E_y] \hat{\mathbf{x}} - (2c(\theta)s(\theta) E_x + [c^2(\theta) - s^2(\theta)] E_y) \hat{\mathbf{y}} \\ &= [(c(2\theta) E_x + s(2\theta) E_y) \hat{\mathbf{x}} - (s(2\theta) E_x - c(2\theta) E_y) \hat{\mathbf{y}}] e^{i\omega t}.\end{aligned}\quad (2.18)$$

Comparison of the above result with Equation 2.13 shows that the end result is an electric field with its polarization rotated by an angle 2θ , corresponding to two times the rotation angle of the half-wave plate.⁴¹ Before the HWP was rotated, the signal was registering as pure Q polarization. After this rotation we find, using Equation 2.12,

$$Q = [\cos^2(2\theta) - \sin^2(2\theta)] E_x^2 = \cos(4\theta) Q_0, \quad (2.19)$$

$$U = -2 \cos(2\theta) \sin(2\theta) E_x^2 = -\sin(4\theta) Q_0. \quad (2.20)$$

The signal no longer looks like pure Q polarization. From this derivation it is clear that spinning the HWP by 22.5 deg will rotate a pure Q polarization into $-U$ polarization. By rotating the half-wave plate by a known amount, any polarized signal skywards of the plate

⁴¹Incidentally, the cumbersome notation might have caused people to look for alternatives, eventually resulting in the development of the Jones and Mueller matrix formalisms.

will be modulated, while spurious signals, internal to the half-wave plate, will not be affected. This procedure therefore reduces contamination from false polarized signals.

2.8 Detector Architecture

SPIDER uses arrays of Transition Edge Sensors (TES), a cryogenic bolometer technology [191], lithographed onto a silicon wafer. The detectors were designed and built by the Microdevices Laboratory at the Jet Propulsion Laboratory in Pasadena, in collaboration with the Caltech Observational Cosmology group. A network of dipole-antennas absorbs radiation coherently and deposits it thermally on a detector element instrumented with tiny partially superconducting thermometers. With 6×6 and 8×8 grids of dual-polarized detectors on each tile for 94 and 150 GHz, respectively, and four tiles per focal plane, each telescope employs either 288 or 512 detectors in total. The SPIDER detector architecture, focal plane, and TES theory are described in various publications [182, 192, 193]. In this section we present some of the basic concepts that are needed to understand the operational principles of a TES bolometer.

Transition edge sensors represent a class of cryogenic detectors that are cooled to superconducting temperatures and then voltage biased to a transition region where the resistance of the sensor is a strong function of temperature. The TES architecture is becoming ubiquitous in CMB experiments as almost all experiments currently in design phase are utilizing such detectors. Figure 2.21 shows the resistance of a typical TES element as a function of temperature. In order to achieve the temperature sensitivity offered by this technology, the detector has to be accurately placed in the superconducting transition region. This is accomplished by careful tuning of device parameters, voltage biasing, and inductive current sensing.

The bolometer signal originates in a dual-polarization grid of antennas which absorbs radiation and sums the signal coherently in a network of niobium microstrip lines. At 150 GHz, this grid is composed of 2×144 slot antennas which are placed a quarter wavelength above a

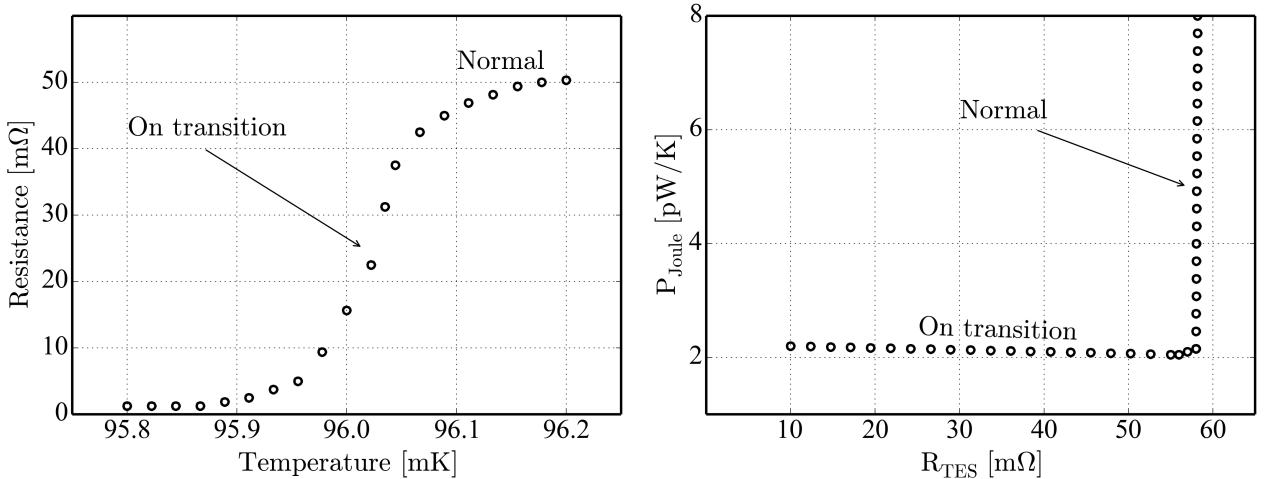


Figure 2.21: *Left:* Resistance as a function of temperature for a typical TES. The steep derivative near the transition temperature suggests that the device could function as a thermometer. *Right:* The Joule power dissipated in a typical TES as a function of its resistance.

superconducting niobium back short. This results in an angular response with a Gaussian fullwidth at half maximum of approximately 15 deg [193]. The signal is then passed through a third-order bandpass Chebyshev LC filter that defines both the upper and lower frequency cutoff, before it is terminated in a gold meander located on the detector island. The island is suspended with silicon nitride (Si_3N_4) legs and, therefore, thermally isolated and sensitive to temperature fluctuations. Heat from the microstrip lines raises the temperature of the voltage biased TES which we measure as a change in current. Figure 2.22 illustrates some of the main elements of the SPIDER detector architecture while Figure 2.23 shows photos of individual detectors as well as a full focal plane.

As Ohmic heating occurs in the gold meander located on the TES island the temperature of the island is raised. This heat is slowly conducted away through the weakly connected Si_3N_4 legs that support the island, see Figure 2.24. However, any rise in temperature of the voltage biased transition edge sensor changes its resistance and, therefore, the current flowing through the bias lines. With proper tuning of device parameters, the electrothermal coupling causes a swift reduction in TES Joule heating which dominates heat flow through

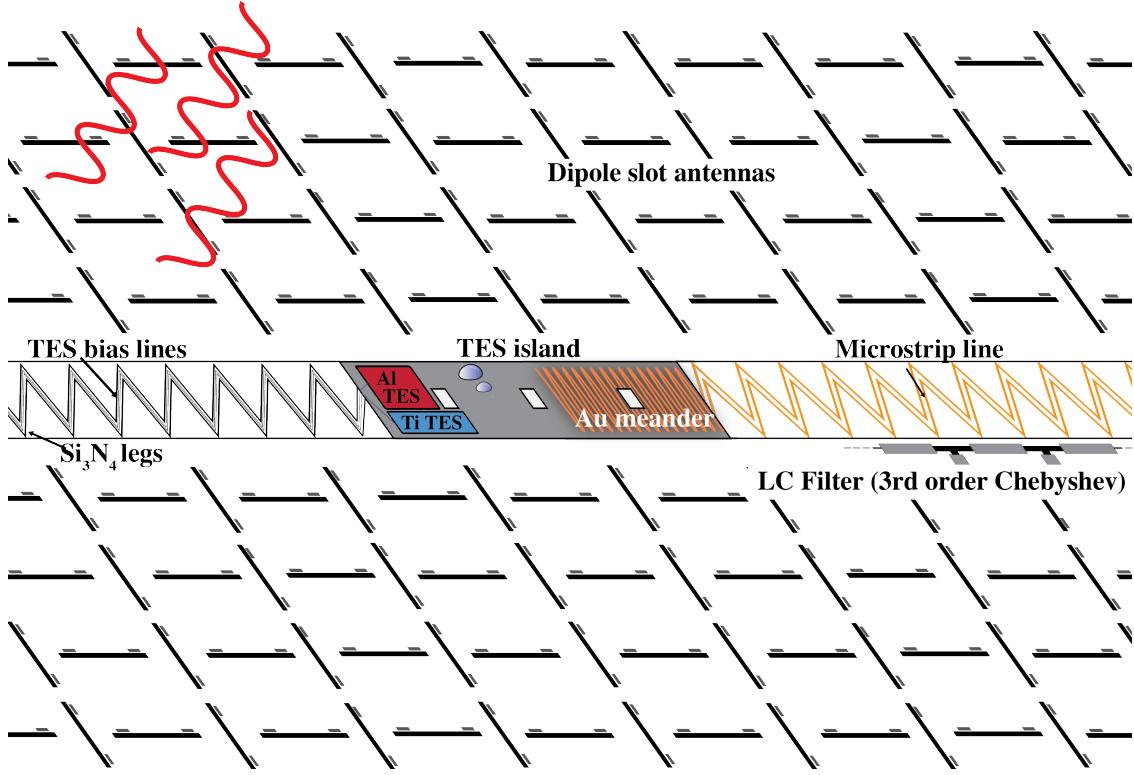


Figure 2.22: An illustration showing the main components of a SPIDER detector. Radiation excites electrons in dipole slot antennas that are phased coherently. Tiny microstrip lines, not shown here, connect the dipole antennas together and direct them towards a resistive gold meander that is located on the TES island. Ohmic heating on the island elevates the temperature of the TES detectors which registers as a change in current flowing through the transition edge sensors.

the legs. This behavior is described by a set of coupled differential equations relating the time development of current and temperature to physical properties of the detector assembly. Simple heat balancing suggests

$$P_{\text{legs}} = P_{\text{Joule}} + Q, \quad (2.21)$$

where Q is the optical power dissipated in the meander, P_{Joule} is the heating in the TES element, and P_{legs} is the heat conducted away through the Si_3N_4 legs. For the voltage biased TES element a change in P_{Joule} corresponds to a change in current according to

$$P_{\text{Joule}} = I_{\text{TES}}^2 R_{\text{TES}} = \frac{V_{\text{TES}}^2}{R_{\text{TES}}}. \quad (2.22)$$

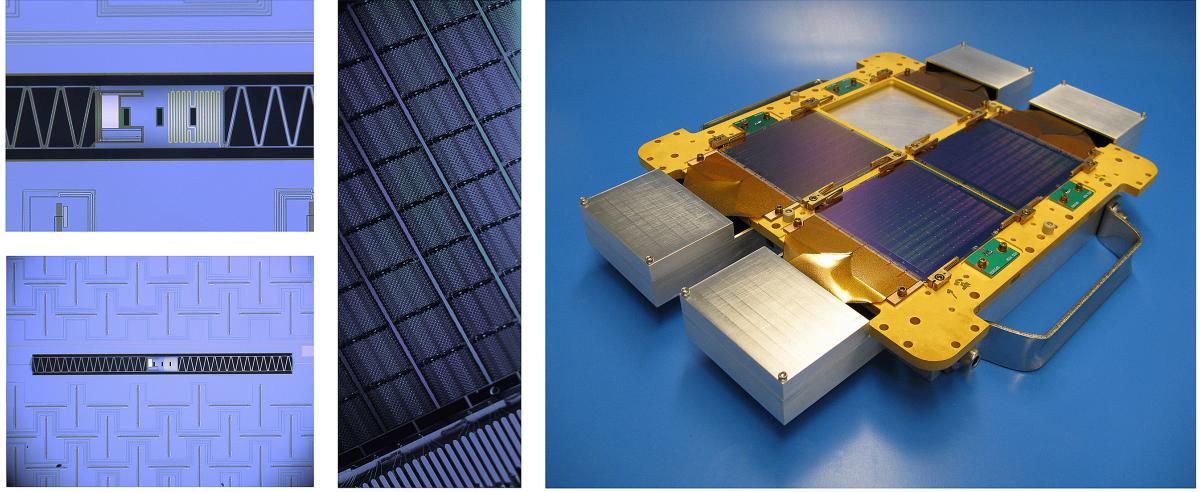


Figure 2.23: A collection of photos showing the SPIDER TES architecture. *Top left*: The TES island with its two TES detectors (left), gold meander (right), and the Si_3N_4 legs (far left). *Bottom left*: A fraction of the antenna array surrounding the detector island. *Middle*: A fraction of a detector tile, showing around 20 pixels. *Right*: A SPIDER focal plane populated with three detector tiles and the flexible superconducting cables that send bias line towards the SQUID assemblies located inside magnetically shielding enclosures.

At fixed bias voltage, any increase in meander heat dissipation must be compensated by an equal drop in the loading through the TES or by cooling through the legs [194]. Figure 2.21 shows how the Joule power dissipation remains essentially constant while the detector is on transition. This plateau is known as the saturation power. As long as the optical loading does not exceed this power, the detector remains on transition.

The natural time constant dictated by the heat capacity of the island, C , and its conductance, G , can be shown to follow $\tau = C/G$. With the TES electrothermal feedback it is found that the response time is enhanced by a dimensionless number L_I ,

$$L_I = \frac{\alpha_I P_{\text{Joule}}}{GT}, \quad (2.23)$$

where

$$\alpha_I = \left. \frac{T}{R} \frac{dR}{dT} \right|_R, \quad (2.24)$$

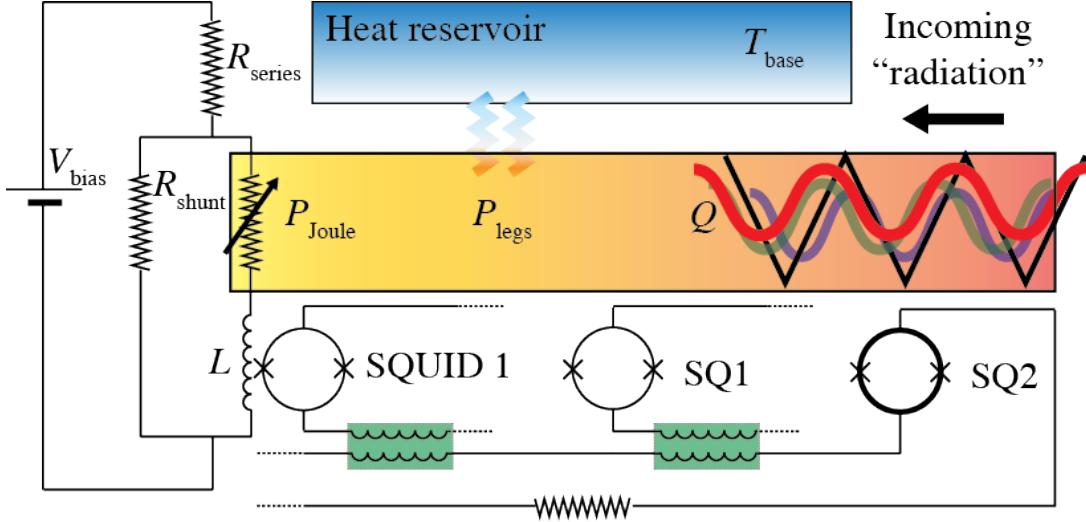


Figure 2.24: The approximate SPIDER TES thermoelectrical circuit. A shunt resistor maintains a steady voltage bias over the TES island since $R_{\text{TES}} < R_{\text{shunt}}$. A change in temperature of the TES corresponds to a change in resistance which is measured through inductive coupling using a superconducting quantum interference device (SQ1). A whole column of SQ1 detectors are read out inductively through a summing network which is connected in series with another SQUID (SQ2).

where R and T represent the resistance and temperature of the TES. The parameter L_I is called the loop gain and for a typical SPIDER detector we expect $L_I \approx 40$. Electrothermal feedback can reduce the time constant of detectors by at least two orders of magnitude, down to approximately 1 ms [195, 196].

A simplified electrical circuit of the detector is shown in Figure 2.24. A $3 \text{ m}\Omega$ shunt resistor is placed in parallel with the TES, which, when kept at transition, has at least five times the impedance of the shunt resistor. This means that most of the current of the electrical circuit flows through the shunt resistor, maintaining fixed voltage bias over the TES. From Figure 2.24 we infer the system of coupled differential equations

$$\begin{aligned}
 V_b &= I_{\text{TES}}(R_L + R_{\text{TES}}) + L \frac{dI}{dT}, \\
 C \frac{dT}{dt} &= Q + I_{\text{TES}}^2 R_{\text{TES}} - P_{\text{legs},(T,T_{\text{base}})}, \tag{2.25}
 \end{aligned}$$

where L is the inductance of the TES circuit and T_{base} is the temperature of the heat reservoir

on the other end of the legs. This system of differential equations is linearized and solved for small signal amplitudes in [194].

SPIDER uses a two element TES design, with saturation powers designed for ground based (Al) and in-flight (Ti) loading. The two TES elements are adjacent and connected in series on the TES island such that the optical responsivity is identical. Since we know the impedance of the titanium TES when driven normal, we are able to characterize the optical and spectral response of the detectors on the ground without sacrificing the improved sensitivity offered by reduced atmospheric loading. The architecture and characteristics of the SPIDER TES detectors is described in [192, 197].

2.9 Fourier Transform Spectroscopy

Fourier Transform Spectroscopy refers to frequency decomposition of electromagnetic radiation. An object that facilitates this is known as a spectrometer or a Fourier Transform Spectrometer (FTS). A Martin-Puplett interferometer is a variant of FTSs that directs polarized light through a wire grid beamsplitter. Its main advantage is wide spectral coverage afforded by the wire grid beamsplitter over more traditional dielectric splitters [198].

Ideally, the spectral bandpass of the SPIDER detectors resembles a top-hat function, unity in the passband and zero outside. Unfortunately, this can never be fully realized. The spectral response of the SPIDER detectors has to be characterized with considerable accuracy if the experiment is to reach its science goals. The effective bandwidth is required to analyze detector noise and to constrain spectral gain mismatch due to differential frequency response of a detector pair [199]. We estimate that approximately 1 GHz spectral resolution, between roughly 50 and 300 GHz, is required to adequately characterize the SPIDER bandpass. An FTS measurement can also identify non-idealities in the optical system, such as spectral fringing due to AR coat delamination. It is especially important to verify that the bandpass has minimal overlap with various atmospheric emission lines, such as those of oxygen (118 GHz)

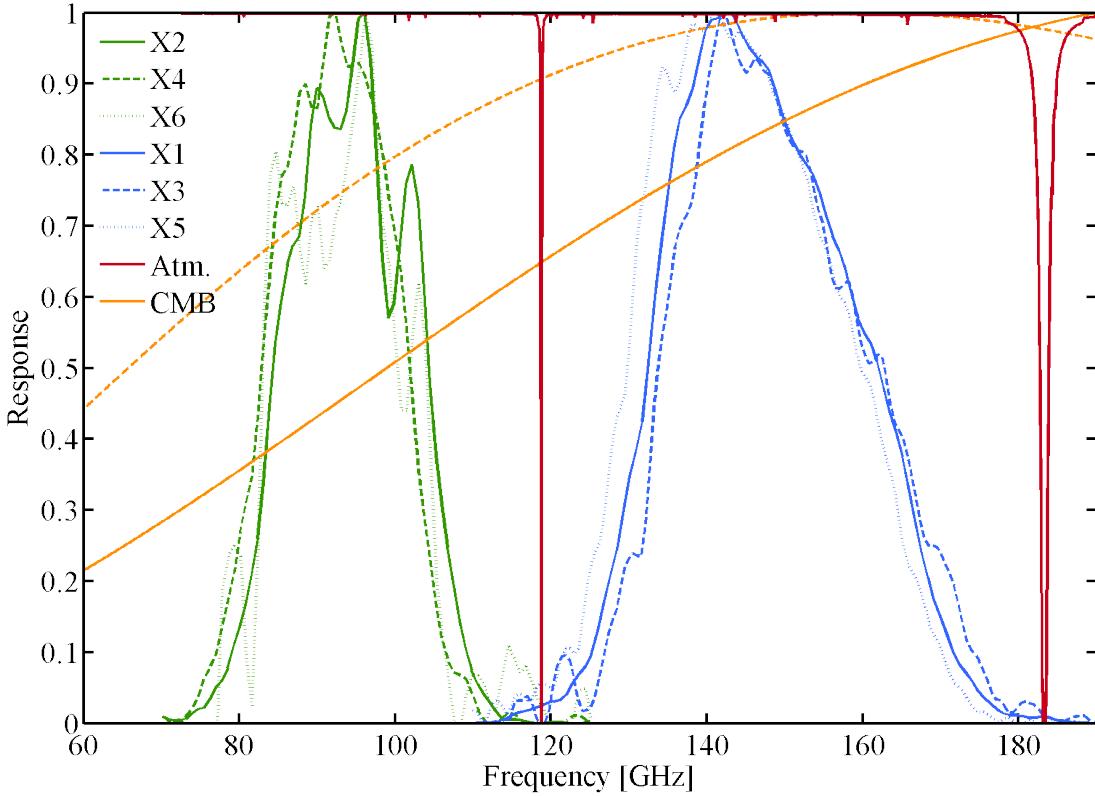


Figure 2.25: The average spectra of all six focal planes as measured in the Caltech test cryostat. Red curve shows an estimate for the atmospheric opacity at 30 km, corresponding to low float altitude. The lines at 118 and 183 GHz come from oxygen and precipitable water vapor respectively. Some percent level overlap between the these two absorption lines is present in both frequency bands. Two orange curves are shown as well, dashed line represents the first temperature derivative of the Planck blackbody function evaluated at $T_{\text{CMB}} = 2.726 \text{ K}$ while the dashed line is the Planck blackbody function evaluated at the same temperature.

and water (183 GHz). An FTS measurement determines $F(\nu)$, the relative detector response as a function of frequency.

Using $F(\nu)$, the effective band center of the SPIDER detectors is calculated according to

$$\nu_c = \int \nu F(\nu) d\nu \quad (2.26)$$

where $F(\nu)$ has been scaled so that its frequency integral is unity. The effective bandwidth is defined in various ways in the literature. Here we choose the definition from Kraus's famous

Table 2.7: All six telescope band centers, bandwidths, and color corrections for dust and synchrotron relative to the CMB.

Property	X2	X4	X6	X1	X3	X5
Band center [GHz]	94.2	92.2	95.7	147.2	148.9	144.7
Bandwidth [GHz]	26.1	25.3	20.6	40.2	40.3	40.8
$C_{\text{dust}} - 1 [\%]$	0.20	0.19	0.09	0.20	0.20	0.21
$C_{\text{sync}} - 1 [\%]$	-3.1	-2.9	0.4	-2.9	-2.8	-3.1

textbook [200]:

$$\Delta\nu = \frac{(\int F(\nu)d\nu)^2}{\int F^2(\nu)d\nu}. \quad (2.27)$$

Figure 2.25 shows the telescope-averaged spectra for all six telescopes. Note how the spectra of the 150 GHz telescopes are bounded by the oxygen and water emission lines. Similarly, Table 2.7 describes the telescope-averaged spectral characteristics of the SPIDER experiment as measured using our FTSSs. In particular, the color corrections C_{dust} and C_{sync} quantify how an absolute calibration on the CMB anisotropies, in units of K_{CMB} , will differ from an absolute calibration based on a region dominated with synchrotron radiation ($\beta = -3.0$) or dust ($\beta = 1.75$). Knowing the spectral bandpass and, therefore, the color corrections, could prove crucial for discriminating between CMB and foregrounds such as dust and synchrotron radiation. The color correction for the dust component following a power law with $\beta = 1.75$ is defined as

$$C_{\text{dust}} = \frac{\int d\nu F(\nu) (\nu/\nu_c)^\beta}{\int d\nu F(\nu) \frac{dB(T_{\text{CMB}}, \nu)}{dT} / \frac{dB(T_{\text{CMB}}, \nu_c)}{dT}}. \quad (2.28)$$

See Appendix D.3 for a more extensive discussion on color corrections. As SPIDER will likely be calibrated on temperature anisotropies, C_{dust} describes by how much we need to multiply our signal if we assume all of it comes from dust.

A schematic for a typical Martin-Puplett interferometer setup is shown in Figure 2.26. Light is emitted from a source at the focus of a 90 deg off-axis collimating mirror. The reflector directs the radiation towards a beam splitting wire grid which sends the light along

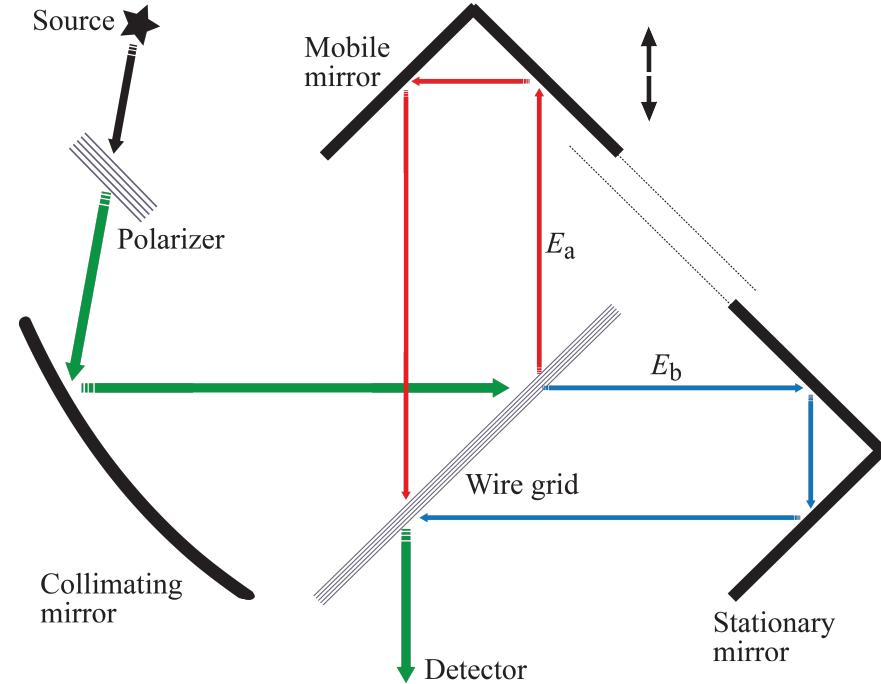


Figure 2.26: A diagram showing the path of light in a Martin-Puplett interferometer. Light is emitted from a source and subsequently passes through a polarizing wire grid. The source is placed at the focal point of a 90 deg off-axis collimating mirror which creates a collimated beam with a 6 inch diameter. The light passes through a wire grid beamsplitter, which split the light between the two axes of the system. Signal is modulated by the translation of one of the rooftop mirrors.

two different optical paths. After interacting with the polarizer for a second time the light is directed towards a detector. The relative phase difference of the two light bundles causes a modulation in radiation intensity as measured by the detector.

2.9.1 FTS Design

Two spectrometers for the SPIDER experiment were designed and built at Princeton University. The first spectrometer was finalized in the Spring of 2010 and sent directly to Caltech for calibration of telescopes mounted in the test cryostat.⁴² A second, lighter spectrometer was

⁴²This FTS has since been used at Caltech for SPIDER (95/145 GHz), Keck Array (95/145/220 GHz), BICEP3 (95 GHz), TIME (a 200–300 GHz spectrometer), MUSIC (a kinetic inductance detector operating in the sub-mm range), and MAKO (a pathfinder instrument for 350 micron imaging arrays). It was also used to characterize MUSIC in the field at the Caltech Submillimeter Observatory in Hawaii. A duplicate FTS, built by Zak Staniszewski, based on this original design has been used by BICEP2 and the Keck Array at

built in the spring of 2013. This spectrometer can be mounted on the flight cryostat and used to calibrate telescopes in their final configuration. Because of the complicated mounting scheme, the FTS needs to be lightweight.

The newer FTS design, including most optical components, is shown in Figure 2.27. The base of these spectrometers is a 0.5 inch thick light-weighted aluminum cast plate cut by a waterjet with threaded holes placed on a one inch square grid. We try to limit the height of the FTS by bringing the symmetry axis of the modulated beam as close to the plane of the breadboard as possible, with the strongest constraint set by the height of the linear stage. This makes the instrument more susceptible to reflections from the breadboard which we try to avoid by lining all elements with absorptive *Eccosorb AN-72*. As constructed, the symmetry axis is located 6.0 inches from the surface of the breadboard. Optical elements are aligned during the design process by placing fiducial crosshairs into the design. The FTS would then be aligned, assuming all parts are accurately machined. Further alignment is done by installing a source (such as an amplified broadband noise source) at the focal point of the collimating mirror and aligning optical components to maximize signal as registered in a broadband receiver at the output.

Appendix B contains some mathematical exploration of the SPIDER spectroscopic system which might prove useful to a person designing a similar instrument. The following pages describe various elements in the second FTS design.

Sources

Our hot source is a ceramic housing oven igniter located inside two thin stainless steel cylinders. Radiative power is controlled using a variable voltage regulator. We found that setting the regulator to 50–70 V is sufficient for $\mathcal{O}(10^3)$ Signal-to-Noise-Ratio (SNR) on the SPIDER bolometers. An infrared thermometer suggests that the temperature of the element

the South Pole and at Harvard. A modified version with a longer throw (higher resolution) has been used for the TIME spectrometer. Dr. Norman Jarosik, a Senior Research Physicist at the Princeton University Physics Department, supplied the wire grids used in most of those spectrometers.

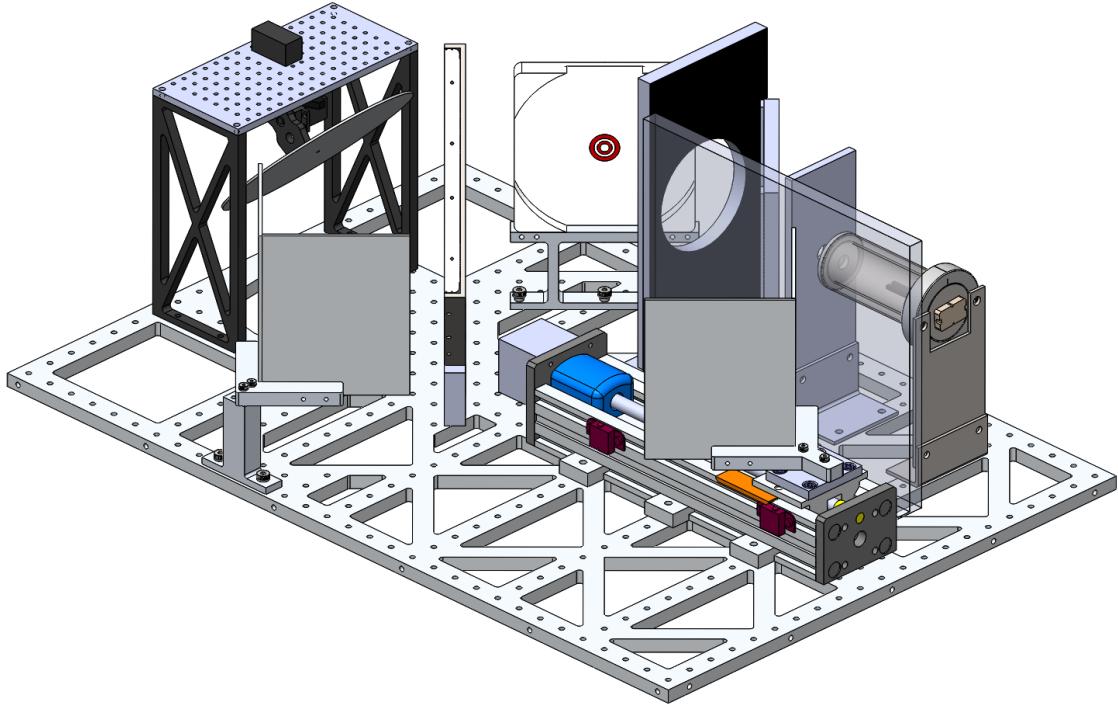


Figure 2.27: A CAD model of the entire FTS assembly without the cover. The breadboard measures 31×21 inches. The hot source shown in the right of the figure can be modulated by a chopper wheel. Baffling and *Eccosorb* intercepts stray radiation.

exceeds 800 K. The hot source can also be replaced by a Styrofoam box which we fill with liquid nitrogen. Styrofoam is transparent at millimeter wavelengths, so the box is filled with *Eccosorb HR-10* which is highly emissive.

Collimating Mirror

A 90 deg off-axis parabolic mirror collimates light from the source and directs it towards the wire grid beamsplitter. The surface of the mirror corresponds to a segment of a parabola which has been cut into a plate rotated by 45 deg with respect to the symmetry axis of the parabola, see Figure 2.28. A quick derivation shows that the mirror surface equation corresponds to

$$z = 4\sqrt{2}f + x - \sqrt{8f(4f + \sqrt{2}x) - y^2}, \quad (2.29)$$

where f is the focal length of the original parabola and x and y represent the translation from the parabola coordinate system to a system that is centered on the mirror plate. The effective focal length of the 90 deg off axis mirror is $f_{\text{eff}} = 2f$. The full derivation is given in Appendix B.3.

To reduce mass, the parabolic mirror is made from ABS (black) thermoplastic instead of aluminum.⁴³ The surface of the mirror is coated with aluminum after machining. We used a company named *VTI Vacuum Technologies* which specializes in RFI shielding solutions.⁴⁴ They refer to their process as vacuum coating; it generates a 6–8 μm thick coat with superior uniformity compared to conductive paint. Simple surface conductance measurements using a voltmeter give $R_{\text{surf}} \approx 0.3\Omega$. Figure 2.29 shows the surface of the aluminized mirrors before and after polishing as viewed at 30x magnification. Measurements of mm-wavelength reflectance suggest no difference between the aluminized material and a solid block of aluminum at the few percent level.

The parabolic shape is cut from a black ABS plastic block on a 3-axis CNC mill using a 1/4 inch ball mill which was stepped by 0.01 inches between cutting passes. The mill is controlled using G-code, see Appendix B.4. The base of the code derives from Thomas Essinger-Hileman [201] while William Dix, from the departmental machine shop at Princeton, oversaw the machining of the mirror.

Wire Grid

The wire grids are expertly wrapped by Dr. Norman Jarosik, a Senior Research Physicist at the Princeton Physics Department. We use gold-plated tungsten wire, with a $d = 5 \times 10^{-4}$ inch diameter, strung at a spacing of $\rho = 200$ wires/inch.⁴⁵ Since all electromagnetic waves should be attenuated before they penetrate the gold film, one can assume that the conductance of the wires is $\sigma = 45.2 \cdot 10^6 (\Omega\text{m})^{-1}$ to find that the cross polarization should be less than a percent

⁴³The thermoplastic has an approximate density of 1.05 g/cm³ compared to 2.7 g/cm³ for aluminum.

⁴⁴*VTI Vacuum Technologies Inc.*, Reedsburg, WI. It costs \$50 to coat the mirror.

⁴⁵Tungsten wire supplied by *Metal Cutting Corporation*, Cedar Grove, NJ.

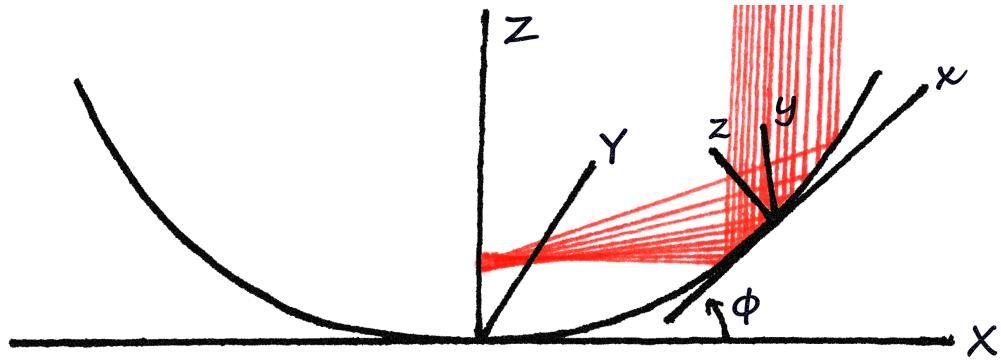


Figure 2.28: The shape of the collimating mirror as calculated from Equation 2.29. It should be noted that the mirror surface is only left-right symmetric, meaning that there is only one proper mounting orientation for the mirror inside the FTS.

over the SPIDER frequency range. Appendix B.1 describes estimates for cross-polarizing properties of wire grids as a function of diameter and spacing.

Beam Splitter

A wire grid rotated at 45 deg in the plane of the breadboard (azimuth) is used as a beamsplitter. The wires in the beamsplitter are oriented at a 45 deg angle with respect to the incoming beam in order to utilize the rotation of polarization angle generated by the rooftop mirrors. However, because of the azimuthal rotation of the wire grid, the viewing angle as seen from a beam that bounces off the collimating mirror is not 45 deg if the grid is tilted at a 45 deg elevation. Instead, a simple geometrical derivation will show that the wire grid needs to be elevated at $\theta_{el} = 35.25$ deg with respect to the surface of the breadboard. In our design, the beam splitter constrains the maximum width of a circular collimated beam that can propagate through the FTS. This circular beam has a diameter $D = 4.6$ inch. Its area therefore corresponds to roughly 50% of the effective area of the focal plane. Because the beam exiting the FTS is not completely uniform, and because its power cannot be easily distributed between the four detector tiles, a typical FTS measurement is performed by illuminating the focal plane in four different locations, one for each tile.

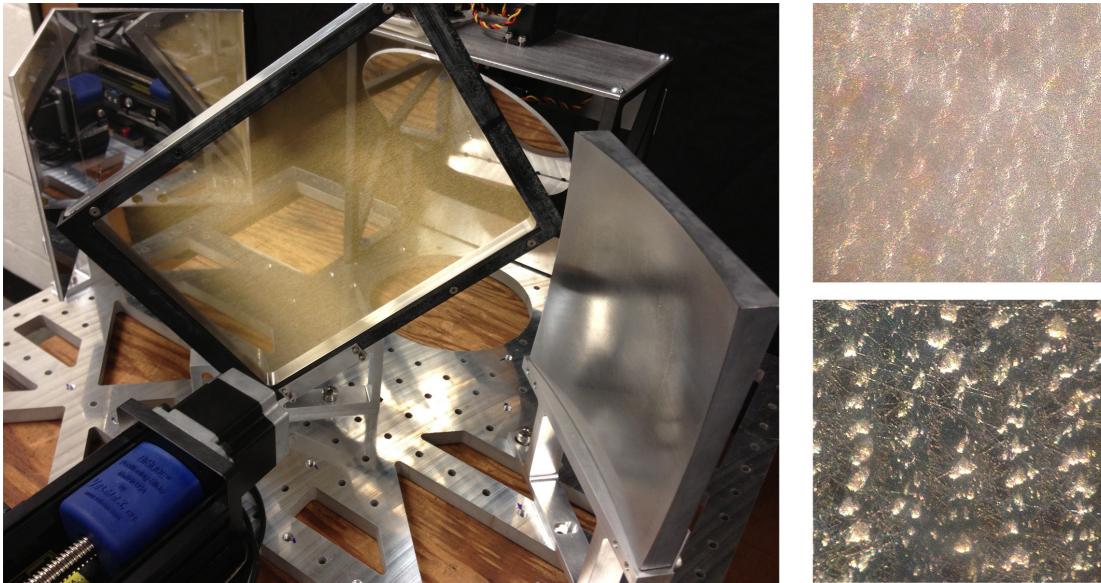


Figure 2.29: *Left*: The wire grid beamsplitter and the collimating mirror mounted on the FTS. *Right*: The surface of the collimating mirror as viewed at 30x magnification before (top) and after (bottom) hand polishing with abrasive grease. The width of the squares is approximately 2 mm. Note the black plastic surface which is revealed by polishing.

Rooftop Mirrors

The rooftop mirrors are 1/8 inch thick mirror finish aluminum plates screwed into precision machined Y-shaped mounting brackets. The brackets are attached to posts that place the center of the mirror at the proper elevation. The Y-shaped mounting blocks are machined with a high tolerance on the 90 deg angle formed by the two arms of the bracket. This ensures that the two sides of the rooftop mirror are perpendicular after they have been bolted to the mounting bracket. One of the mirrors is mounted to a linear stage. The rooftop mirrors are shown in Figure 2.27.

Carbon Fiber Enclosure

A lightweight carbon fiber enclosure allows us to make changes to the FTS configuration while it is mounted on the cryostat. We use 1.3 mm thick carbon fiber panels supplied by *Protech Composites Inc.*⁴⁶ The panels are glued together using carbon fiber angle brackets

⁴⁶*Protech Composites Inc.*, Vancouver, WA.

from *Dragonplate*⁴⁷ and *Scotchweld 2216* epoxy. The unibody cover is lined on the inside with *Eccosorb AN-72*.⁴⁸ The result is a surprisingly rigid assembly, weighing 3 kg, which can be mounted over the FTS by a single person.⁴⁹ This is especially useful when the FTS is mounted on the SPIDER flight cryostat while the cryostat is tilted at 45 deg elevation. Figure 2.30 shows the FTS with the enclosure. The carbon fiber represents a conductive surface for the *Eccosorb* mounting. This improves the shielding performance without compromising weight.

Linear Stage

A commercial solution is found for the linear stage with a *Velmax BiSlide* with 8.3 inches of travel driven by a NEMA 23 stepper motor.⁵⁰ An *Elgo EMIX23* magnetic encoder with 1 μm resolution and 6 μm repeat accuracy is mounted to the side of the stage. This particular FTS is 1.1 sided, meaning that the mirror can move only slightly past the white-light fringe in one direction with the majority of the throw used for a single side of the interferogram. In this configuration the FTS will have a 0.85 GHz spectral resolution. Interferograms are acquired with the linear stage moving at constant velocity, normally at a speed of around 2 mm/s. So far, all analysis of interferograms assumes the linear stage is moving at a constant velocity during the data acquisition. The encoder output, sampled at approximately 100 Hz, suggests that the commanded and realized velocities, when commanded to move at 2 mm/s, agree to within 0.01 mm/s with a 0.01 mm/s RMS error.

Output Mirror

Radiation that interacts with the wire grid beamsplitter for the second time is directed towards a mirror on a simple two axis rotation stage. This mirror directs light through an aperture in the breadboard and towards the telescopes. We use a pair of analog servomotors

⁴⁷ *Dragonplate*, Elbridge, NY.

⁴⁸ *Emerson and Cuming Microwave Products*, Randolph, MA.

⁴⁹ While standing on one leg perilously hunched over the flight cryostat.

⁵⁰ *Velmax Inc.*, Bloomfield, NY.



Figure 2.30: The covered FTS sitting on a lab stool. The mass of the entire assembly is 23 kg, making it possible to transport by a single person.

with 0.9 Nm stall torque and 0.3 deg per step.⁵¹ The two motors are controlled using an *Arduino Uno* prototyping board and a simple joystick interface with an LCD panel. As the maximum diameter of the beam that can propagate through the FTS is not large enough to illuminate all of the focal plane at the same time, the rotatable mirror allows us to move the beam over the focal plane. We find that we can illuminate the entire focal plane with $\mathcal{O}(10^3)$ signal to noise ratio by moving the mirror to four distinct positions.

2.10 Outer Frame

The flight cryostat is mounted on a lightweight gondola structure made of reinforced carbon fiber. The gondola and most of the outer frame components were designed and built by members of Barth Netterfield’s group at the University of Toronto [143, ?, 146, 145]. The carbon fiber inserts are epoxied into aluminum joints that are bolted together on custom-made aluminum joints. The gondola is designed to sustain large angular acceleration during the initial balloon jerk as well as a maximum 10 *g* parachute shock which could be realized during

⁵¹The servomotors are from *HITEC RCD*, Poway, California, named *HS-645MG Ultra Torque*.



Figure 2.31: A rendering of the gondola assembly without the flight cryostat. The reaction wheel is mounted over the enclosure for the CSBF support instrumentation package. The trunnions of the flight cryostat are mounted to the pillow block shown at the top of the figure.

flight termination. The gondola design, including material selection, was informed by a large ensemble of finite element studies and epoxy joints were validated with pull tests under flight like conditions. The gondola design utilizes the high specific strength and stiffness of carbon fiber resulting in a 143 kg structure that supports a fully loaded flight system weighing approximately 1250 kg when full of liquid and loaded with six telescopes [143, 202].

The gondola is connected to the balloon using a three-leg bridle hitch configuration with the suspension cables linking below a pivot. Figure 2.31 shows a rendering of the gondola, including the reaction wheel and suspension cables. A set of linear actuators, not shown in Figure 2.31, drive elevation arms affixed to the trunnions of the flight cryostat. This allows the telescopes to step in elevation from 15 to 45 deg [120].

The SPIDER observation region is defined as a quadrangle on the fixed sky, with edges defined as the great circles passing through those four points. Figure 2.2 shows the approximate SPIDER observation region in Equatorial coordinates. This region is also truncated using limits on azimuthal separation between the Sun and the boresight, 70 deg on the port side

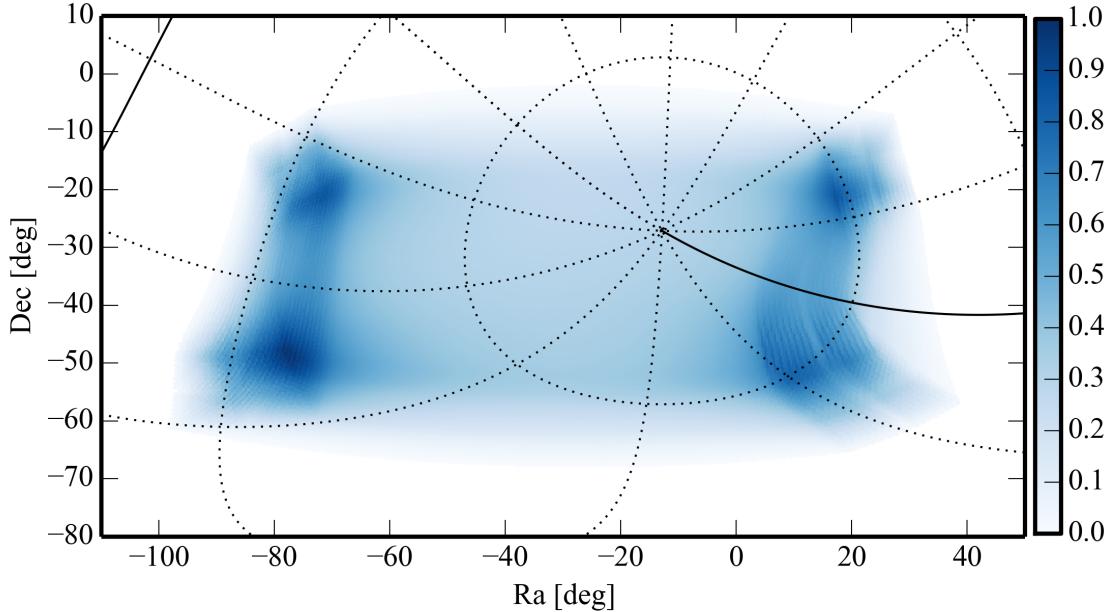


Figure 2.32: The integrated observation time as a function of location for an entire 150 GHz telescope over 24 hour duration, courtesy of Sasha Rahlin and Jamil Shariff [145, 159]. Approximately 75% of the integration time is spent in 50% of the map. Dashed lines show constant latitude and longitude in Equatorial coordinates.

and 90 deg on starboard. SPIDER will observe this region by scanning in azimuth with a 90 deg throw sinusoidal velocity profile and a peak angular acceleration of 0.8 deg/s². This corresponds to a maximum angular velocity of 6 deg/s in the center of the scan region and an approximately 47-second period.⁵² The solar panels are oriented such that, given the scan strategy, the time averaged solar flux is maximized. Azimuthal scanning is facilitated by the combination of a pivot located right under the flight train and a reaction wheel positioned under the cryostat. The pivot dumps angular momentum into the flight train, which causes the balloon to rotate. The pointing control loop is carefully optimized to minimize pendulations. This work is described further in [120, 145].

The SPIDER science scan involves a sinusoidal velocity profile to rotate in azimuth with 0.1 deg elevation steps approximately every third turnaround. Following a monotonic and periodic function, these elevation steps span the entire range of the elevation drive, 20–45 deg,

⁵²This corresponds to 21 mHz in the frequency domain.

in approximately 12 hours. Using this scan strategy, a telescope will cover almost the entire SPIDER observation region in 24 hours. This observation strategy is repeated every sidereal day. For that reason, cross-linking comes primarily from non-negligible sky rotation at these latitudes.⁵³ During observations, the half-wave plates will be stepped by 22.5 deg once every 12 hours, although the exact HWP stepping strategy has yet to be fixed. Figure 2.32 shows the integrated and normalized observation time for an entire 150 GHz telescope over 24 hours of observation [203].

Two independent power systems, each connected to a 4-by-3 solar panel array, supply all power for the duration of the experiment. The panels are wired into switched-mode DC power supplies charging a set of 40 Ah sealed lead-acid batteries. The panels supply 1440 W on average with 2100 W at normal solar incidence. We estimate a 20–25% average power margin with some initial battery drainage during ascent, when azimuthal pointing is unconstrained. The solar panels are mounted on a lightweight aluminum frame and connected to the gondola using constant torque, friction hinges. The panels are deployed manually and locked into position prior to launch. The panel mount was designed to sustain ground winds of up to 15 m/s.

SPIDER will employ a large number of in-flight pointing solutions [146, 204]. A combination of attitude reconstruction devices, including differential GPS, magnetometers, and Sun-sensors will provide in-flight pointing accuracy of approximately 1 deg. Pointing for science analysis is generated by a combination of data from gyroscopes and star cameras. Two star cameras will take 100 ms exposures every two seconds to determine the attitude of the outer frame at those instances. Gyroscope data is then used to interpolate attitude between those intervals. The pointing reconstruction has significant heritage from flight of the BLAST payload [187, 205].

⁵³As telescopes scan the sky, their angular response, or beam, will cross a given point on the sky at a certain angle, determined by the scan strategy. Integrated over the total observation period of that experiment, the beam will most likely cross this point on the sky a number of times. Cross-linking then refers to the distribution of crossing angles. A patch that has good cross-linking will be probed by a wide range of beam crossing angles. Ideal cross-linking symmetrizes the effective beam.

The relatively large size of the SPIDER beams implies modest requirements for pointing reconstruction and realized performance greatly surpasses science goal requirements [124].

Line-of-sight communications are possible for the first 18–36 hours of the flight. This provides a 1 Mbps link which should allow us to monitor housekeeping thermometry and timestreams from approximately 300 detectors. A dedicated link is used to calibrate and focus star cameras during this period. As the payload drifts past the Transantarctic Mountains, all communications are limited to low bandwidth satellite links. A 6 kbps link with the Pentagon run TDRSS satellite constellation should allow for 6–30 detector timestreams with sparsely sampled housekeeping data. The most likely scenario is a 2 kbps link through the commercially run Iridium system. This is only enough to provide sparsely sampled pointing and housekeeping data in addition to a single detector timestream. Great deal of automations is required because of the limited connectivity with the payload. This includes code that cycles fridges, controls pointing, tunes SQUIDs, biases detectors, and flags and remedies various detector performance issues. The development of the SPIDER flight code is a demanding project (see further discussion in [159]).

2.11 Current Status

The SPIDER flight cryostat and all six telescopes are packed and ready to ship from the Princeton High Bay. Payload integration and detector characterization is complete. The experiment will ship to the Antarctic on August 26, 2014. The first wave of scientists will arrive at McMurdo Station at the end of October.

Chapter 3

Calibration of the *Planck* High Frequency Instrument

In this chapter, we describe the absolute calibration and spatial response reconstruction of the *Planck* High Frequency Instrument (HFI). We begin by discussing the dipole calibration, which is used to obtain an absolute detector response. This procedure converts raw detector units into a meaningful flux quantity. The spatial response, also known as the beam response or the point spread function, relates the absolute calibration to different angular scales. We describe an algorithm which employs two-dimensional Gauss-Hermite functions to reconstruct the spatial response from planet observations. We also present a simulation pipeline which estimates error and bias in the beam reconstruction. Using this simulation pipeline, we show that the Gauss-Hermite reconstruction generates a beam response that is internally consistent. This extends the validity of the dipole calibration to all angular scales used in cosmological analysis of the HFI data. Finally, we compare the absolute calibration of the *Planck* HFI detectors to models of planet fluxes. Our analysis suggests that the planet flux models are consistent with the primary calibrations in all six HFI bands.

Table 3.1: Summary of the *Planck* HFI and its band averaged properties [138, 207].

General properties						
Property	Center Frequency [GHz]					
	100	143	217	353	545	857
Bandwidth [GHz]	33	46	65	101	171	246
Angular resolution* [arcmin]	9.65	7.25	4.99	4.82	4.68	4.33
Solid angle [arcmin ²]	104.2	58.4	26.9	25.1	25.4	23.0
Number of detectors	8	12	12	12	4	4
Number of Spiderwebs	0	4	4	4	4	4
Number of PSB's	8	8	8	8	0	0
NET [$\mu\text{K}\sqrt{\text{s}}$]	71	58	88	353	0.087 [†]	0.085 [†]

* FWHM. [†]MJy/sr/ $\sqrt{\text{Hz}}$.

3.1 Instrument Overview

The *Planck* satellite was launched from French Guiana in May 2009.¹ The primary science goal of the satellite is to measure the temperature and polarization anisotropies of the CMB over the full sky [206]. The satellite has two instruments with different scientific leadership. They are the High and Low Frequency Instruments (HFI/LFI). I have been a member of the *Planck* HFI core team since the summer of 2010. The satellite was commissioned by the European Space Agency (ESA) although significant financial and scientific contributions have been made by NASA and other institutions in the US.

Both instruments onboard the satellite are coupled to the sky through an off-axis Gregorian telescope. Two mirrors, a 1.5 m primary and a 1.0 m secondary, each cooled to 45 K, focus light into the 52 feedhorn coupled bolometric receivers that populate the HFI, see Figure 3.1.

¹Even though the satellite has completed its observations, data are still being analyzed at the time of writing. I will therefore use the *present tense* for most statements regarding the project.

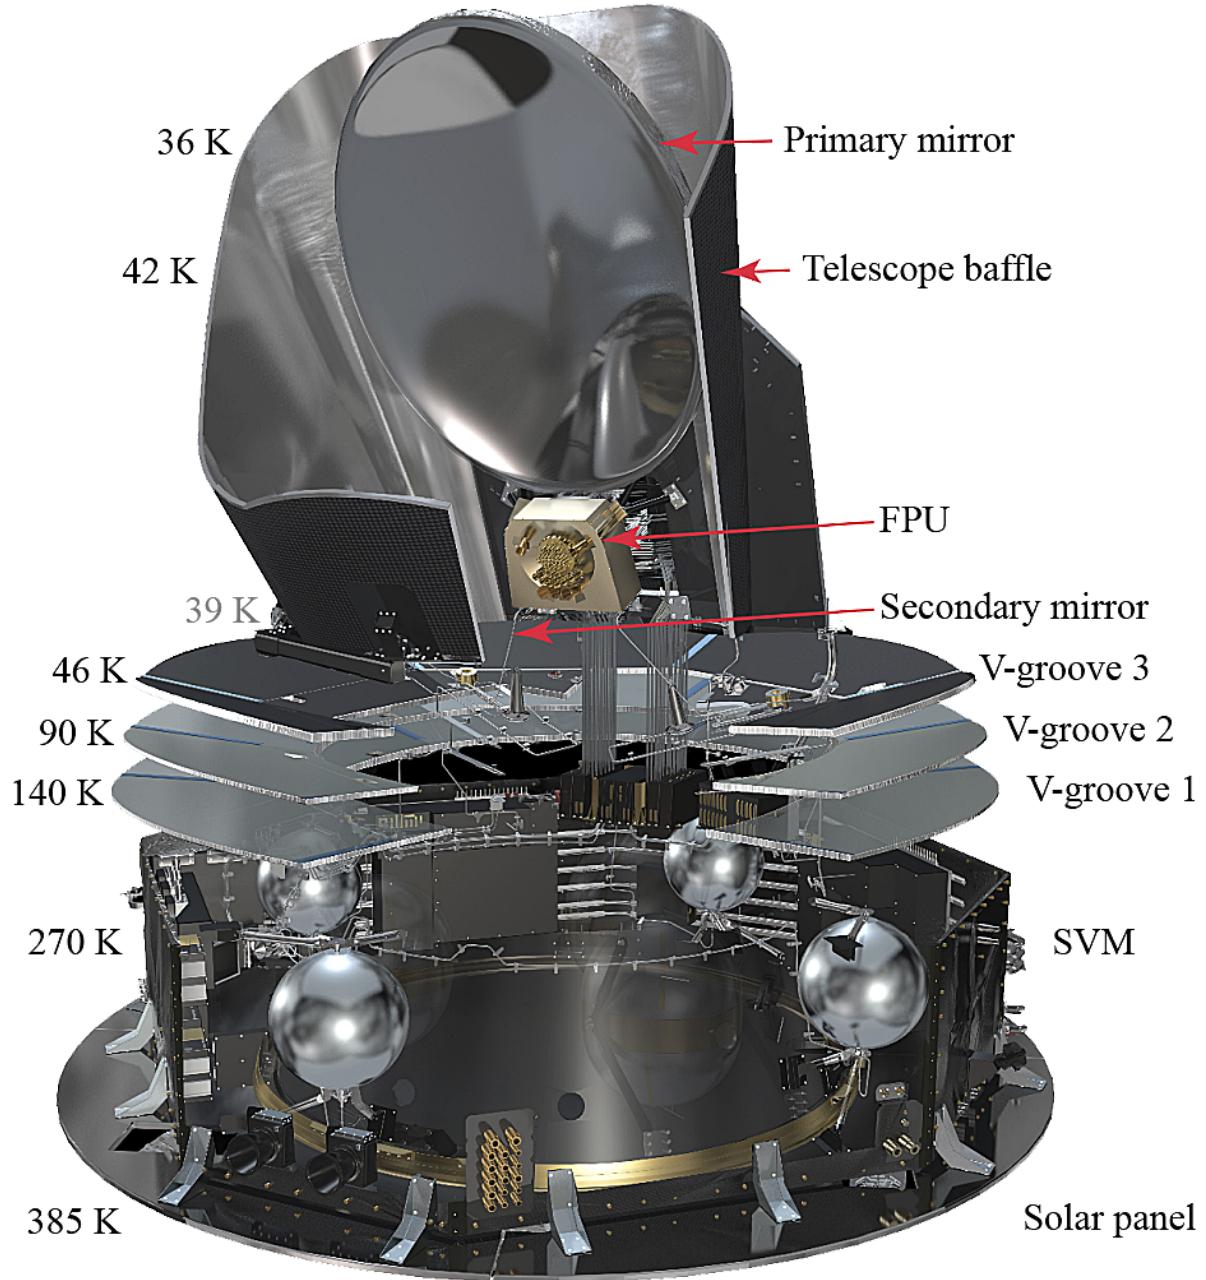


Figure 3.1: A cutaway drawing of the *Planck* satellite showing the cryogenic instruments inside the Service Module (SVM), as well as the 1.5 m primary reflector and the focal plane [208]. Two black star cameras, mounted on the side of the SVM, and a set of thrusters that facilitate orbital maneuvers are visible. Four spherical titanium and Kevlar composite structures, designed for 290 atm maximum pressure, hold the ${}^4\text{He}$ and ${}^3\text{He}$ needed for the open-cycle dilution refrigerators [209]. The satellite spins around its vertical symmetry axis while ensuring that the solar panels, mounted at the bottom, are always directly facing the Sun. This provides a stable thermal environment and maximizes power generation. Figure courtesy of ESA [210].

The spectral coverage of these receivers is spread over six frequencies, with band centers ranging from 100–857 GHz. Each bolometer has a unique identifier. For example, “100-1a” refers to the A-polarization channel of a 100 GHz detector [138]. Parts of the receiver elements are cooled to 100 mK through a series of passive and active coolers [209]. These refrigerators made *Planck* the coldest known object in space while it was operational.

The HFI frequency coverage is optimized for study of CMB temperature and polarization anisotropies and sensitivity to the Sunyaev-Zel’dovich (SZ) effect [211]. The higher frequency channels enable characterization of Galactic foregrounds and studies of the cosmic infrared background. More than half of the HFI detectors are Polarization Sensitive Bolometers (PSB) [212], the remainder are Spiderweb Bolometers (SWB) [213]. The PSB’s are distributed over the four lowest frequency bands, centered on 100, 143, 217, and 353 GHz [214]. These four bands are often referred to as the CMB channels, even though the 353 GHz band sees a lot of Galactic foregrounds and is not directly used to estimate the CMB power spectrum.

The two highest frequency bands, each with four SWBs, have band centers at 545 and 857 GHz. As 545 GHz corresponds to a 0.55 mm wavelength, these two bands are often called the sub-millimeter channels [215, 216]. The sub-mm channels are multi-moded in order to increase the throughput, and therefore, the signal-to-noise ratio of those channels at the cost of limited resolution [217]. Neither scan strategy nor sample rate justify finer resolution at those frequencies. Unfortunately, the multi-moded nature of the sub-mm channels seriously complicates any analytical description of their spatial response.

The HFI focal plane is populated with back-to-back feed horns that couple light from the secondary mirror onto the cryogenic bolometers . The feedhorns are designed to produce a Gaussian response function with a -25 – 30 dB edge taper [218, 219]. Table 3.1 highlights some of the main features of the *Planck* HFI.

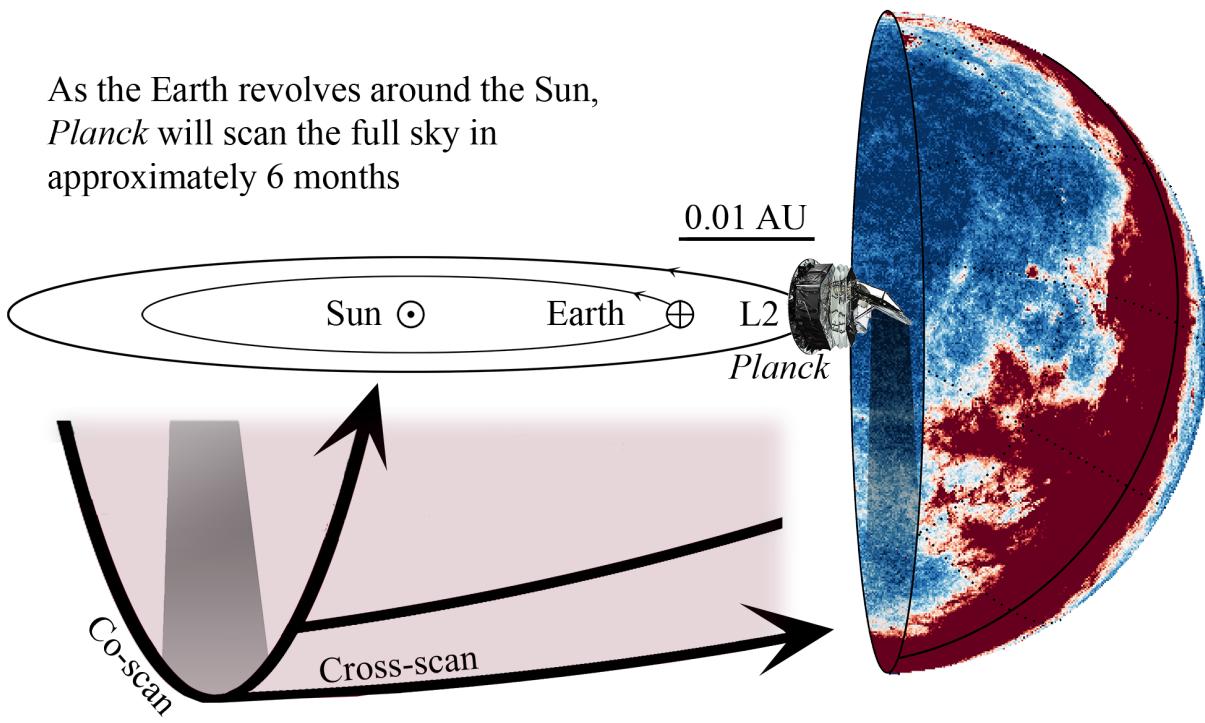


Figure 3.2: The location of *Planck* satellite, relative to the solar system and a depiction of *Planck*'s scan strategy – obviously not to scale.

3.2 Scan Strategy and Pointing

Planck is located at L₂, a Lagrangian point outside Earth's orbit with an identical sidereal period (see Figure 3.2). The colocation of Earth and the Sun on the sky as viewed from L₂ make it an optimal venue for satellites conducting full-sky surveys. *Planck* rotates around its symmetry axis at 1 rpm while stepping azimuthally, or in the ecliptic plane, by 2 arcmin every hour. This ensures that the satellite's solar panels are pointed directly at the Sun at all times, therefore, maintaining a stable thermal environment and minimizing stray radiation [220]. In the time between these azimuthal steps, the satellite will trace out approximately 60 circles on the sky. The set of such circles is referred to as a “ring”. Additionally, the spin axis precesses with a 7.5 deg amplitude over the duration of a survey to cover the poles [221]. Residual drifts and nutations are minimal and accounted for in pointing reconstruction [222]. With this scan strategy, *Planck* observes the whole sky in approximately 6 months.

After a preliminary characterization phase, full-sky observations began in August 2009 [223]. Science observations continued relatively unperturbed until cryogens were exhausted in January 2012, after 30 months of observations. This operational period corresponded exactly to the projected lifetime of the refrigerators [206]. Over the duration of the experiment, the HFI mapped the whole sky approximately 5 times. The final products are all-sky maps at a HEALPix resolution characterized by $N_{\text{side}} = 2048$, corresponding to 1.8 arcmin pixel width [4]. Given the satellite’s sky coverage and observation time, this corresponds to an average integration time of 0.5 seconds per pixel.

Planck’s scan strategy does not provide optimal coupling between large angular scales and temporal scales that are short compared to the $1/f$ noise knee of its detectors. This somewhat complicates calibration of the instrument. Furthermore, the scan-strategy provides minimal cross-linking, whereby the satellite scans over a fixed location on the sky with a varied ensemble of crossing angles. Finally, complications arise due to the relatively slow time response of the HFI detectors, compared to the 6 deg/s scan rate.

The HFI bolometers have time constants in the range of 2–10 ms, corresponding to 1–4 arcmin at the nominal scan rate [213]. Time-domain deconvolution of the time-response function makes the resultant signal response frequency independent and allows us to probe angular scales up to and including the beam size in pixelized maps. Because of filtering and errors in time-response deconvolution, the correlation of samples on the sky will depend strongly on their angle relative to the primary scan direction. Influences producing temporal correlations, such as cosmic rays, cause similar effects. The result is a dataset that has a propensity for ring-aligned residuals. Detailed discussion is offered in the most-recent *Planck* DPC paper [138] and the HFI time-response and beams paper from 2013 [207].

Data are matched with locations on the sky through post-processing of pointing information. The pointing data are primarily supplied by the two time-delay integration star cameras that are mounted on the side of the service module (see Figure 3.1). The cameras are sampled at 8 Hz. Higher resolution pointing timelines are subsequently generated using

quaternion interpolation. The satellite also houses gyroscopes which provide ancillary pointing information. The relative orientation of the detectors with regards to the star cameras is derived by fitting the detector point spread function to observation of the bright point sources such as planets.

It is useful to define a Cartesian coordinate system for a small patch of the sky, with axis aligned parallel and perpendicular the primary scan direction of the satellite. These axes are called the co- and cross-scan directions, and they are shown in Figure 3.2. Using this new terminology, we can make the following statement regarding *Planck*'s scan strategy: Samples separated in the cross-scan direction by more than 2 arcmin will be separated temporally by at least one hour.

3.3 Absolute Photometric Calibration

The *Planck* HFI is calibrated using the CMB dipole signal generated from the motion of the satellite relative to the rest frame of the CMB. This also causes aberration at all multipoles, a minute signal, whose detection was first reported in [67]. The absolute calibration provides a mapping from raw detector units to an effective CMB temperature, K_{CMB} . The satellite motion relative to the rest frame of the CMB can be separated into the orbital motion of the satellite as it rotates around the Sun (orbital-dipole), and the peculiar velocity of the solar system barycenter relative to the CMB rest frame (solar-dipole).

So far, *Planck* HFI has been calibrated using the solar-dipole assuming its properties, as measured by *WMAP*, are valid [224, 225]. For the 2014 release, *Planck* HFI will be calibrated using an orbital dipole template [226]. The calibrations of individual detectors are obtained by varying g_r to minimize the following ring-by-ring residual, \mathcal{R}_r ,

$$\mathcal{R}_r = \sum_{i \in r} |\mathcal{F}_i^{\text{PSM}} + \mathcal{D}_i^{\text{orbit}} - (g_r \times s_i + o_r + n_i)|, \quad (3.1)$$

where the subscript i refers to individual detector samples, g_r represents the conversion from

raw detector units to K_{CMB} , s_i is time ordered data, o_r is the ring-by-ring offset, which helps capture the long time constant noise properties, n_i is the time-dependent noise model, and $\mathcal{F}_i^{\text{PSM}}$ is a foreground model. Accurate information about the satellite's peculiar velocity relative to the CMB rest frame allows us to estimate $\mathcal{D}_i^{\text{orbit}}$ and constraint an orbital dipole model. This equation is linearized and solved by iteration.

Temporal variations in gain, i.e. variation in G_d with time, have been observed [138] and are due to non-linearity in the analog to digital readout electronics. This effect appears to cause approximately 3% variation in G_d over the duration of the experiment. This effect is rectified in the 2014 data release [226].

The absolute calibration of the dipole signal relates power on the largest angular scales – over 180 degree angles – to the digital readout units of the flight computers. In order to estimate the power spectrum of the CMB, this calibration has to be extended from the largest angular scales towards the fundamental limits set by the angular resolution of the instrument. The function that defines the angular sensitivity of the detectors in multipole space is called the beam window function. It represents the azimuthally averaged spherical harmonic transformation of the beam response on the unit sphere.

To justify this statement, we start by writing an approximate expression for a signal registering in our detectors as the convolution of the instrument beam response, $\Psi(\hat{\mathbf{n}})$, and the underlying temperature anisotropies²:

$$\begin{aligned} s_i &= \int d\Omega \Psi(\hat{\mathbf{n}} - \hat{\mathbf{n}}_i) T(\hat{\mathbf{n}}), \\ &= \int d\Omega \Psi(\hat{\mathbf{n}} - \hat{\mathbf{n}}_i) \sum_{\ell,m} a_{\ell,m} Y_{\ell,m}(\hat{\mathbf{n}}), \end{aligned}$$

where $T(\hat{\mathbf{n}})$ represents the temperature anisotropies as defined in Equation 1.21. We can then use Equation 1.24 to calculate the corresponding beam convolved power spectrum. This

²Neglecting time-response and frequency-dependent effects.

involves performing a double integral over the unit sphere

$$\begin{aligned}
\langle T(\hat{\mathbf{n}})T(\hat{\mathbf{n}}') \rangle_{\text{beam}} &= \iint d\Omega d\Omega' \Psi(\hat{\mathbf{n}}) \Psi(\hat{\mathbf{n}}') \times \langle T(\hat{\mathbf{n}})T(\hat{\mathbf{n}}') \rangle_{\text{true}}, \\
&= \iint d\Omega d\Omega' \Psi(\hat{\mathbf{n}}) \Psi(\hat{\mathbf{n}}') \times \frac{1}{4\pi} \sum_{\ell=1}^{\infty} (2\ell+1) C_{\ell}^{TT} P_{\ell}(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}'), \\
&= \frac{1}{4\pi} \sum_{\ell=1}^{\infty} (2\ell+1) C_{\ell}^{TT} \iint d\Omega d\Omega' \Psi(\hat{\mathbf{n}}) \Psi(\hat{\mathbf{n}}') P_{\ell}(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}'), \tag{3.2}
\end{aligned}$$

from which we conclude that

$$W_{\ell} = \iint d\Omega d\Omega' \Psi(\hat{\mathbf{n}}) \Psi(\hat{\mathbf{n}}') P_{\ell}(\hat{\mathbf{n}} \cdot \hat{\mathbf{n}}'), \tag{3.3}$$

and

$$\tilde{C}_{\ell}^{TT} = W_{\ell} C_{\ell}^{TT}, \tag{3.4}$$

where W_{ℓ} is the window function, C_{ℓ}^{TT} is the true TT power spectrum of the sky and \tilde{C}_{ℓ}^{TT} is the power spectrum as measured by the instrument [227]. Calibrating the absolute response of the detectors using the CMB dipole corresponds to applying a multiplicative detector gain such that W_{ℓ} is unity by definition at $\ell = 1$. Beam calibration then amounts to estimating $\Psi(\hat{\mathbf{n}})$ for all values of $\hat{\mathbf{n}}$ so that the window function, and therefore the absolute calibration, can be extended to all angles.

The coefficients in the spherical harmonic decomposition of the beam function,

$$b_{\ell m} = \int d\Omega \Psi(\hat{\mathbf{n}}) Y_{\ell m}(\hat{\mathbf{n}}), \tag{3.5}$$

encapsulate the spatial beam response in multipole space. After determining the coefficients, the beam window function can be found from [228]:

$$W_{\ell} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |b_{\ell m}|^2. \tag{3.6}$$

For *Planck* HFI beam analysis, the beam response, $\Psi(\hat{\mathbf{n}})$, is first established. The spherical harmonic transform of the beams then determines $b_{\ell m}$ and therefore W_ℓ [207].

To a very good approximation, all N detectors within a given frequency band are observing the same microwave sky. Because of this, intra-frequency calibration tests are routinely performed for the HFI cosmology channels. This is done by finding calibration coefficients, α_d , to minimize the residual

$$\mathcal{R}_\ell = \sum_{\ell=\ell_{\min}}^{\ell_{\max}} \left(\bar{C}_\ell - \frac{1}{N} \sum_{d=1}^N \alpha_d C_{\ell,d} \right), \quad (3.7)$$

where

$$\bar{C}_\ell \equiv \frac{1}{N} \sum_{d=1}^N C_{\ell,d} \quad (3.8)$$

is the band average power spectrum, and $C_{\ell,d}$ are the temperature power spectra obtained from individual detector maps. By varying the set of α_d within a frequency band, the residual can be minimized over some range of angular scales. This method will reveal systematic effects with an angular scale dependence.

These intra-frequency calibrations provide a relative scaling among detectors in a given frequency band. All that remains in order to correctly debias the power spectrum estimate is to understand the band-averaged beam window function \bar{W}_ℓ . However, that process is best done by estimating $\Psi_d(\hat{\mathbf{n}})$, the angular sensitivity of individual detectors.³

3.3.1 Beam Definitions

Planck's scan strategy, time-response deconvolution errors, and filtering procedures generate the need for three slightly different definitions of the function $\Psi(\hat{\mathbf{n}})$. These are referred to as the physical beam, the scanning beam, and the effective beam. The physical beam represents the true optical response of the instrument. It can be determined from an unfiltered

³In reality, the process is reversed. Individual detector beam window functions are determined from planet observations before the intra-frequency calibration.

timestream while observing a point source in the limit of slow (relative to time response) scanning. The physical beam estimate is mainly used to compare with models describing the in-flight alignment of optical elements.

The interplay between bolometer transfer function residuals, filtering, and scan rate produces signal variations on arcminute scales [213, 207]. As all data are processed using the same procedure, the scanning beam represents the point-source response that is applicable to the processed timeline and is therefore necessary for accurate photometry. See [229] for the most up-to-date description of the *Planck* HFI time-domain processing. Unlike the physical beam, the scanning beam can change with revision to the official time-domain processing. This means that a scanning beam appropriate for the 2013 data release need not be applicable for future releases.

The effective beam represents the scanning beam averaged on the whole sky, taking the scan strategy into account.⁴ In the limit of a rotationally symmetric beam, $\Psi(\hat{\mathbf{n}}) = \Psi(\theta, \phi) = \Psi(\theta)$, the effective beam is identical to the scanning beam. In reality, the amount of symmetrization due to scanning varies over the sky. This requires careful accounting for accurate cosmological analysis (see discussion on cross-linking in Section 2.10). For *Planck*, cross linking is almost non-existent in the ecliptic but maximal at the celestial poles. Because of this variation in coverage, the effective beam also varies on the sky.

As the CMB is decomposed from a large fraction of the sky, the beam that is appropriate to debias an angular power spectrum according to Equation 3.4 is formed by a scan-weighted linear composition of effective beams centered on all pixels in the *Planck* maps that are used in cosmological analysis. The effective beam window function therefore also depends on the sky weighting used for cosmological analysis. The problem of efficiently calculating an effective beam applicable for cosmological analysis of *Planck* data is described in the literature (see for example [207, 230, 231]). However, in order to calculate an effective beam,

⁴Phrased differently, the effective beam is simply a full-sky average of the scanning beam, properly weighted by the observation region.

the scanning beam has to be determined. The first step invariably involves looking at bright point sources such as planets.

3.3.2 Planet Observations

The outer planets entered *Planck*'s field of view a few times during the lifetime of the experiment. As the planets lie close to the ecliptic, the raster-like scanning of the instrument is roughly perpendicular to their proper motion. This results in a characteristic 2 arcmin striping of the planet data. Table 3.2 describes some features of the *Planck* HFI planet crossings. Of the five outer planets, Mars, Jupiter, and Saturn are most useful for characterizing the *Planck* HFI beams. Although Jupiter saturates detectors above 143 GHz, it can still be used to characterize the near sidelobe response of the beams. With the exception of modest SNR, Mars is almost ideal for characterizing the main beam. Its relatively small angular size, lack of rings, large moons, and rare events, reduce features that are difficult to simulate. Similarities between the peculiar motion of Mars and *Planck*'s scanning strategy caused the planet to stay in the field of view of the satellite for a prolonged time during the first observation, resulting in significantly greater sample density compared to most of the other planet observations.

Assuming constant beam response, distinct point source observations can be combined to increase the signal to noise ratio in beam reconstruction. Such observation merging also improves the sampling of the beam shape. However merging of observations potentially induces systematics in beam determination that, if not done carefully, can negate improvement in coverage. Here, time-domain simulations of the beam reconstruction become crucial. Systematics induced by stacking observations include normalization and pointing errors, discussed in greater detail in Section 3.5.

Table 3.2: Some properties of the *Planck* planet observations as they pertain to channel 143-1b. Diameter represents the planet diameter as viewed from L_2 , averaged over all observations. Sample density refers to the average sample density, accounting for flagging, within a 40 arcmin wide field of view centered on the planet. The Signal-to-Noise Ratio, SNR, is the ratio between a fit to the peak signal registering in the timeline and the noise Root Mean Square (RMS).

Obs	Diameter	Sample density	SNR
	[arcsec]	[samples/arcmin 2]	
Mars ^{*,†}	1–3	8	16 / 12 / 12
Jupiter ^{‡,○,§}	1–5	40	9 / 13 / 14 / 9 / 11
Saturn ^{‡,○,§}	1–4	17	11 / 14 / 10 / 14
Uranus [*]	1–5	3.5	12 / 12 / 13 / 12 / 9
Neptune [*]	1–4	2.3	10 / 13 / 10 / 13

^{*}Astronomical background confusion. [†]Time-variation. [‡]Rings.

[○]Moons. [§]Rare events.

3.4 Beam Reconstruction

The scanning beam describes the data-processed response to a point source excitation normalized to unity at peak responsivity. In addition to filtering, the time-domain processing is further complicated by the presence of cosmic rays, electromagnetic interference from cryocoolers, and ADC non-linearity. For example, approximately 15–20% of the Time-Ordered Data (TOD) are flagged (removed) because of cosmic rays and the transients that they produce [138]. These effects make it difficult to construct an optimal time-domain processing algorithm.

Because of the satellite’s raster scan, data with only a few arcmin angular separation, can be separated temporally by a similar number of hours. With no observation time dedicated to beam characterization, all information about the beams come from raster scans with planets more-or-less moving in the ecliptic, perpendicular to the primary scanning direction. The *Planck* scan strategy and bolometer time-response therefore drive the optimal time-domain processing. Even in the limit of perfectly Gaussian physical beams, due to filtering, the

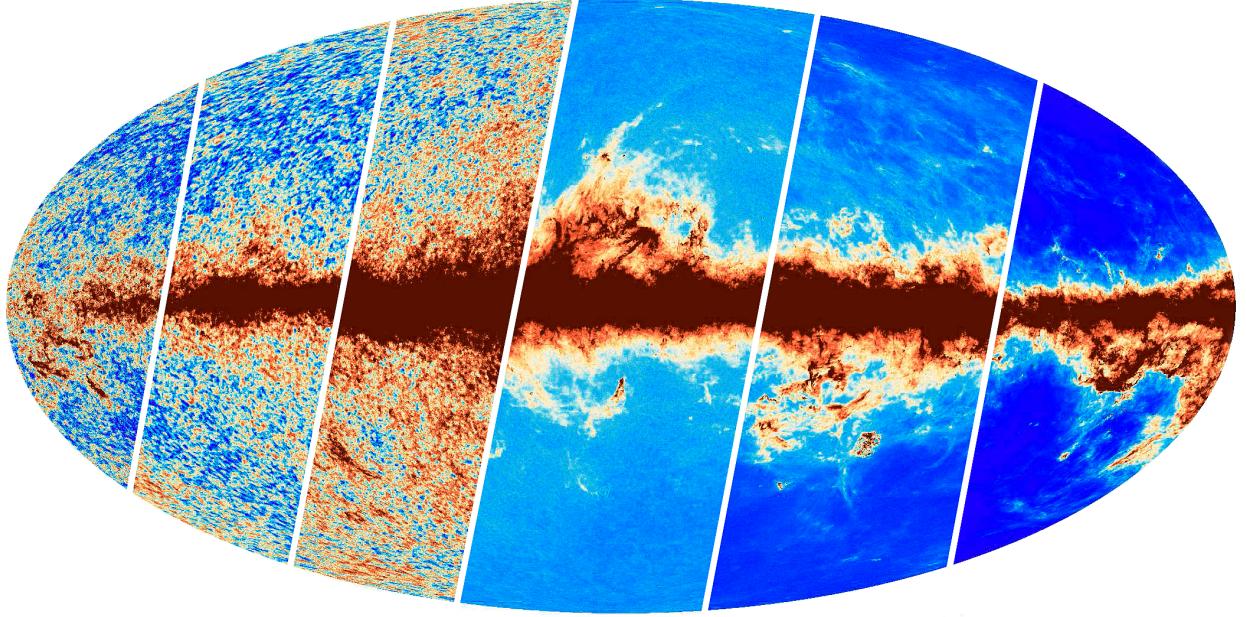


Figure 3.3: From left to right, 100, 143, 217, 353, 545, and 857 GHz maps. The color scale is linear in signal amplitude. Placement of diagonal lines is chosen to emphasize the 353 GHz map. From this, one can surmise that three lowest frequency bands are most likely to constrain the CMB signal amplitude. Maps extracted from [138].

Planck scanning beams are not azimuthally symmetric. The beams must be treated as two-dimensional objects in analysis. When probing a sufficiently small patch on the sky, a region of the unit sphere can be accurately described by a Cartesian coordinate system. As the HFI beams are compact, all analysis is performed assuming the flat sky approximation.⁵

Sparse sampling and the requirement of reconstruction at arbitrary resolution bring about the need for a functional form. The remainder of this section is devoted to the two-dimensional Gauss-Hermite functions, an orthogonal function basis on \mathbb{R}^2 , which can be used to describe the scanning beam.

⁵In the following text, it is safe to assume that the cross-scan direction is parallel with the x -axis and the co-scan direction parallel with the y -axis.

3.4.1 The Gauss-Hermite Function Basis on \mathbb{R}^2

Many CMB experiments are designed to have angular sensitivity that can be described by an azimuthally symmetric two-dimensional Gaussian function

$$\Psi(\mathbf{x}) \propto \exp(-\mathbf{x}^2/2\sigma^2), \quad (3.9)$$

where σ represents the width of the beam. Optical aberrations will lead to asymmetries in the angular sensitivity which can often be captured by assuming that the Gaussian beam width is different along the two axis of a Cartesian coordinate system centered on the peak response

$$\Psi(x, y) = \exp\left(-\frac{1}{2}[x^2/\sigma_x^2 + y^2/\sigma_y^2]\right). \quad (3.10)$$

This is referred to as an elliptical Gaussian function. A natural extension of this beam parametrization is found in the Gauss-Hermite (GH) functions. This functional basis has been used for beam modeling by other CMB experiments [232, 233], with Huffenberger et al. describing the methodology in relation to *Planck* HFI [234]. The two-dimensional Gauss-Hermite functions are defined here as

$$\psi_{n_1, n_2}(x, y) = \frac{H_{n_1}(x/\sigma_x)H_{n_2}(y/\sigma_y)}{\sqrt{2^{n_1+n_2}n_1!n_2!}} \exp\left(-\frac{1}{2}\left[\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2}\right]\right), \quad (3.11)$$

with the orthogonality relation

$$\langle \psi_{n_1, n_2} \psi_{n'_1, n'_2} \rangle = 2\pi\sigma_x\sigma_y\delta_{n_1 n'_1}\delta_{n_2 n'_2}. \quad (3.12)$$

In this notation, the often quoted beam Full Width at Half Maximum (FWHM) is approximately $\theta_{\text{FWHM}} = \sqrt{8\ln(2)\sigma_x\sigma_y}$. Beams can be modeled using an ensemble of these functions

$$\Psi(x, y) = \sum_{(n_1+n_2) \leq N_{\text{max}}} s_{n_1, n_2} \psi_{n_1, n_2}(x, y), \quad (3.13)$$

where s_{n_1, n_2} is the decomposition coefficient of that Gauss-Hermite mode. The maximum number of functions used in the Gauss-Hermite decomposition along any single dimension, N_{\max} , is called the decomposition order parameter. This number should be chosen based on the resolution of the raw maps which in turn depends on the spatial coverage. For the *Planck* HFI beam analysis, N_{\max} has spanned 8–16, corresponding to 45–153 unique two-dimensional functions (see discussion on mapping of indices in [234]).⁶ A higher value for the decomposition order parameter leads to increase in systematic induced variance. For results presented in this chapter, we set $N_{\max} = 9$. This corresponds to 55 Gauss-Hermite coefficients.

In addition to the Gauss-Hermite coefficients, this beam description uses the five parameters of an elliptical Gaussian fit, corresponding to beam centroids, rotation angle, and the beam widths. The Gauss-Hermite functions are exponentially suppressed outside $\sqrt{2\sigma_x\sigma_y N_{\max}}$, see for example [232]. Figure 3.4 shows plots of the first few Gauss-Hermite functions of an elliptical Gaussian beam for a typical 100 GHz channel. It is important to note that ψ_{n_1, n_2} is an odd function if either n_1 or n_2 are odd. Odd functions do not contribute to the solid angle of the beam.

The first step in the Gauss-Hermite reconstruction involves estimating the best-fit elliptical Gaussian parameters. They are the beam widths, σ_x and σ_y , the beam center c_x and c_y , and the rotation angle, ϕ , of the beam coordinate system relative to some predefined coordinate system. Once these are established, the Gauss-Hermite coefficients are calculated in a straightforward manner (see Appendix C).

The Gauss-Hermite functions are useful for beam reconstruction for a number of reasons. The GH parameters can be decomposed quickly in the time-domain, generating a reasonable parametrization of the beam in a few seconds.⁷ The parametrization therefore lends itself easily to simulations of the beam reconstruction process. Furthermore, because the Gauss-Hermite functions are eigenfunctions to the Fourier transform operator, the beam window

⁶The calculation time scales like N_{\max}^2 .

⁷As measured in some pseudo-standardized way on a 3 GHz core.

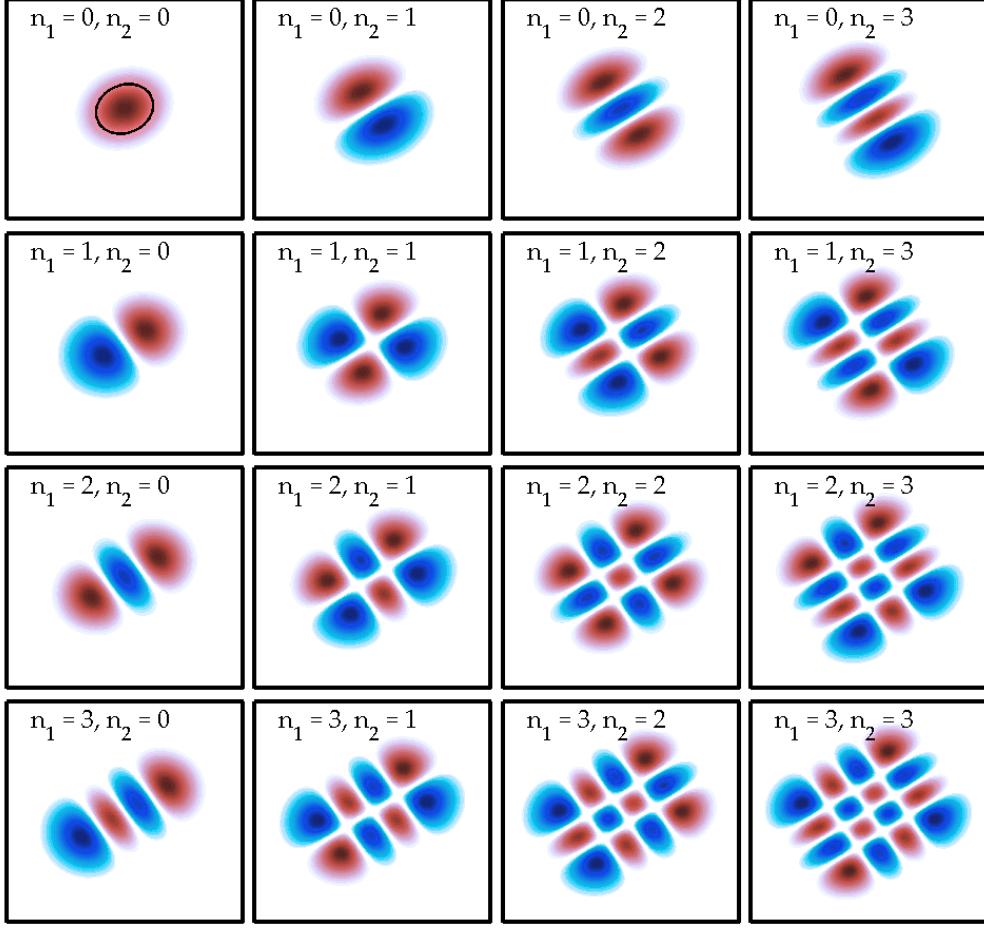


Figure 3.4: The first few Gauss-Hermite functions associated with an elliptical Gaussian with $\theta_{\text{FWHM}} = 9.2$ arcmin and a tilt angle $\phi = 33$ deg. The maps are 20 arcmin on each side and the FWHM of the elliptical Gaussian is shown in the top left corner.

function corresponding to a known set of GH parameters can be easily calculated using a semi-analytical expressions. The expression is

$$W_\ell = \frac{(\sigma_x \sigma_y)^2}{2\pi} \int_0^{2\pi} d\phi \sum_{n_1, n_2} s_{n_1, n_2}^2 \times \left| \frac{e^{-i(\ell' \cos(\phi) c_x + \frac{n_1 \pi}{2})} H_{n_1}(\ell' \cos(\phi) \sigma_x)}{e^{\ell'^2 \cos^2(\phi) \sigma_x^2 / 2} \sqrt{2^{n_1} n_1!}} \times \frac{e^{-i(\ell' \sin(\phi) c_y + \frac{n_2 \pi}{2})} H_{n_2}(\ell' \sin(\phi) \sigma_y)}{e^{\ell'^2 \sin^2(\phi) \sigma_y^2 / 2} \sqrt{2^{n_2} n_2!}} \right|^2, \quad (3.14)$$

where a numerical integral has to be performed over the angle ϕ for each Gauss-Hermite parameter and $\ell' \equiv \ell + 0.5$. The full derivation is presented in Appendix C. Members of the

Planck collaboration have used this expression to quickly estimate the Gauss-Hermite beam window function.

The Gauss-Hermite function basis does have its limitations. The parametrization loses its flexibility once the elliptical Gaussian basis and N_{\max} have been defined. As we limit the number of functions, the shapes that we can describe with the GH function basis are also restricted. Compare this with a cubic spline interpolation which sets less stringent constraints on the beam shape [235]. Because of these constraints, the GH parametrization does not fully capture some features in the beams, such as the acausal ringing caused by the filtering of timelines. Increasing the number of GH functions to further capture secondary features leads to unwanted noise fitting. To understand just how useful the Gauss-Hermite pipeline is for *Planck* HFI beam reconstruction, we perform simulations.

3.5 Time-Domain Simulations

A range of non-idealities affect the scanning beam reconstruction. These include, but are not limited to, pointing errors, time-variations, noise, residual glitches, incomplete astrophysical background subtractions, intricate frequency dependence, and destriping. Figure 3.5 describes the approximate order of simulation steps.

Analytic expressions for first order contributions of various beam non-idealities to temperature and polarization spectra are discussed in the literature (see for example [236, 237, 238]). Such estimates are only tractable for generalized scan strategies. Proper time-domain simulations offer a tailored view of systematics induced by complicated time-domain processing and beam reconstruction techniques.

3.5.1 Methodology

There is no single well-defined method for performing cosmological analysis using CMB temperature and polarization maps. Experimental groups will employ varying analysis

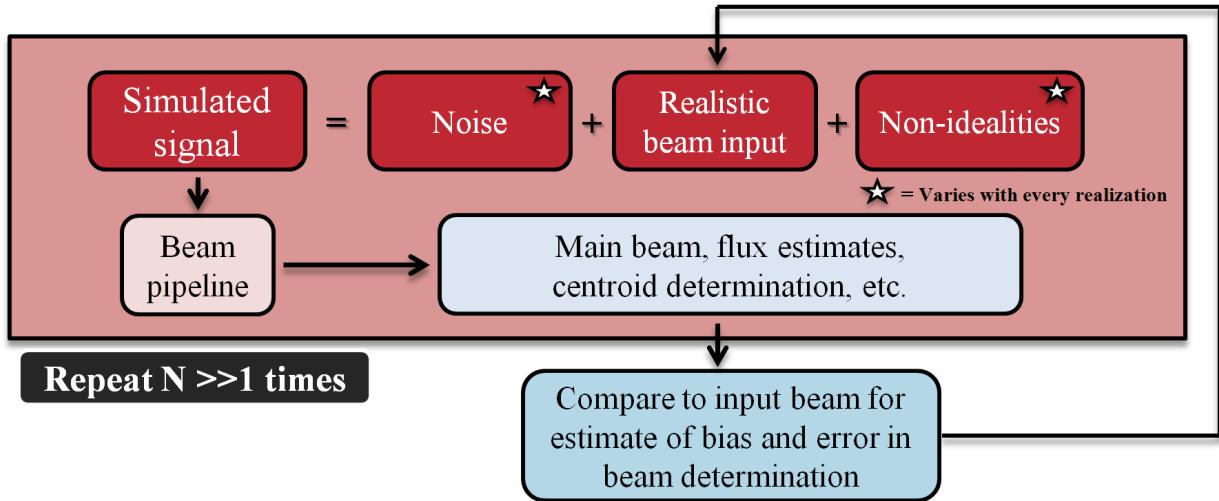


Figure 3.5: A flow diagram describing the rough simulation methodology. The simulated signal is generated by a combination of a realistic beam input, which takes pointing information as input, an estimate for the noise, and some set of non-idealities, which could include pointing errors or a residual astrophysical background. The simulated timestream is fed into the beam reconstruction pipeline, from which one can estimate the fidelity of beam reconstruction by comparing against the input beam. The simulations are repeated a number of times to negate sample variance.

methodologies. Generally, however, map-making algorithms will convert time-ordered data into a pixelized map in the HEALPix data format [4]. The spherical harmonic transform of those maps is then computed, which allows one to calculate an angular power spectrum. Such spectra will be biased due to cut-sky, filtering, and spatial response effects. The act of estimating the spherical harmonic transform from the windowed sky, is commonly referred to as a pseudo-power spectrum estimate, or pseudo- C_ℓ [239, 240].

In the end, cosmological parameters are constrained by finding how their variations result in power spectra which are consistent with measured values. In some cases, and for some ranges of multipole, beam modeling errors constitute a significant fraction of the errors on the power spectra. For that reason, the task of determining beam reconstruction errors with high fidelity becomes pressing. In lieu of defining an exact process, I will offer the following guidelines:

- The beam simulation pipeline should incorporate the same data processing algorithms as those used on real data.
- Simulations should include all known sources of that error that can be reasonably simulated and injected into the data processing algorithm.
- Using data cuts (jackknives), beam products should be checked for internal consistency against simulation results. This can be done by splitting the data used to derive real beam products into two disjoint datasets before deriving a beam estimate. The two resultant beams should be checked for consistency against an identical process performed on simulated data.

All of these steps should be completed before publishing cosmological analysis.

3.5.2 Quantitative Description of Simulation Framework

Time-domain simulations can assume various levels of complexity. For an almost ideal polarimeter, plagued only by cross-polarization, a given timestream sample can be written as

$$d_i = g_i [T(\hat{\mathbf{n}}) + \gamma_i(Q(\hat{\mathbf{n}}) \cos(2\psi_i) + U(\hat{\mathbf{n}}) \sin(2\psi_i))] \quad (3.15)$$

where T , Q , and U are the Stokes parameters on a location of the sky specified by the unit vector $\hat{\mathbf{n}}_i$, g_i is the time-varying bolometer gain, ψ_i is the polarization angle of the detector at the time when that sample was acquired, γ_i represents cross polar leakage, with

$$\gamma_i = \frac{1 - \epsilon_i}{1 + \epsilon_i}, \quad (3.16)$$

and ϵ_i corresponding to the fractional response of a polarized detector to orthogonally polarized light. Polarimeters are generally affected by non-idealities such as non-negligible detector response function, detector noise, variable gains, involved frequency and angular

response functions, and so forth. A more truthful description of signal timestream therefore follows

$$d(t) = K * \left(n_i + g_i \int d\nu A_e(\nu) F(\nu) \int d\hat{\mathbf{n}}' \Psi(\hat{\mathbf{n}}' - \hat{\mathbf{n}}_i, \nu) \right. \\ \left. \times [T(\hat{\mathbf{n}}_i) + \gamma_i(Q(\hat{\mathbf{n}}_i) \cos(2\psi_i) + U(\hat{\mathbf{n}}_i) \sin(2\psi_i))] \right) + \tilde{n}_i, \quad (3.17)$$

where $K*$ represents a convolution with the detector time response, n_i is the noise, which we assume is uncorrelated with signal, $A_e(\nu)$ represents the effective area of the telescope, $F(\nu)$ is the spectral responsivity, and \tilde{n}_i represents noise terms that are not convolved by the detector response, including readout noise. Full incorporation of all these factors in a time-domain simulation pipeline is considered intractable. To make simulations manageable, we assume that our input beam represents a frequency averaged scanning beam. We also assume that the point sources are unpolarized.

Although the planets represent the brightest microwave sources on the sky, we cannot neglect effects from the CMB and foregrounds such as the Galactic cirrus emission. We therefore need to simulate planet crossings on realistic backgrounds, having some prior on the shape and amplitude of those components from our measurements of the sky without the wanderers in the field of view. In this case, a simple incarnation of a simulated timestream takes the form

$$s_i = C_p \Psi(\hat{\mathbf{n}}_i) + T^{\text{CMB}}(\hat{\mathbf{n}}_i) + \mathcal{F}^{\text{PSM}}(\hat{\mathbf{n}}_i), \quad (3.18)$$

where T^{CMB} and \mathcal{F}^{PSM} are the beam convolved CMB and foreground signal derived using the *Planck* Sky Model (PSM) [129], Ψ represents the input scanning beam assumed for these simulations, and C_p is the peak intensity of the planet in the same units as T^{CMB} and \mathcal{F}^{PSM} . This process assumes that the frequency dependence suggested by Equation 3.17 can be properly encapsulated by the sum of three independent terms which have been convolved with the same input beam. Both CMB and foregrounds are removed from the timeline during beam reconstruction. We do this using data that show the same region of the sky, but

with the planet missing. The process is straightforward for *Planck* HFI, because the sky is observed roughly five times during the lifetime of the experiment. Assuming perfect removal of the CMB and any other astrophysical signal, the simulated signal finally reduces to

$$s_i = C_p \Psi(\hat{\mathbf{n}}_i). \quad (3.19)$$

The name of the game is then to use the beam modeling pipeline to estimate the input beam, $\Psi(\hat{\mathbf{n}})$, in the presence of non-idealities. The beam reconstruction fidelity is then fully defined by the residual $\Delta\Psi(\hat{\mathbf{n}}) \equiv \Psi(\hat{\mathbf{n}}) - \tilde{\Psi}(\hat{\mathbf{n}})$ where $\tilde{\Psi}$ is the reconstructed beam response. Similarly, we calculate the resultant error in window function using Equation 3.3. The stated window function error, σ_ℓ , and bias, γ_ℓ , are

$$\sigma_\ell = \frac{1}{N} \sum_i^N (\bar{W}_\ell - \tilde{W}_\ell^i), \quad (3.20)$$

$$\gamma_\ell = W_\ell - \frac{1}{N} \sum_i^N \tilde{W}_\ell^i, \quad (3.21)$$

where \tilde{W}_ℓ^i represents the window function reconstructed from realization i , \bar{W}_ℓ is the ensemble average of all those realizations, and W_ℓ is the true window function of the input beam.

3.5.3 Non-Idealities

With Equation 3.19 defining a method to simulate an ideal signal timeline, the next step involves adding error inducing effects. Note that the pointing timeline, $\hat{\mathbf{n}}_i$, is extracted from the real timeline.⁸ Of course, the first ingredient that comes to mind is detector noise.

⁸The real timeline represents our best estimate for the actual pointing timeline. Pointing errors can be added to this timeline.

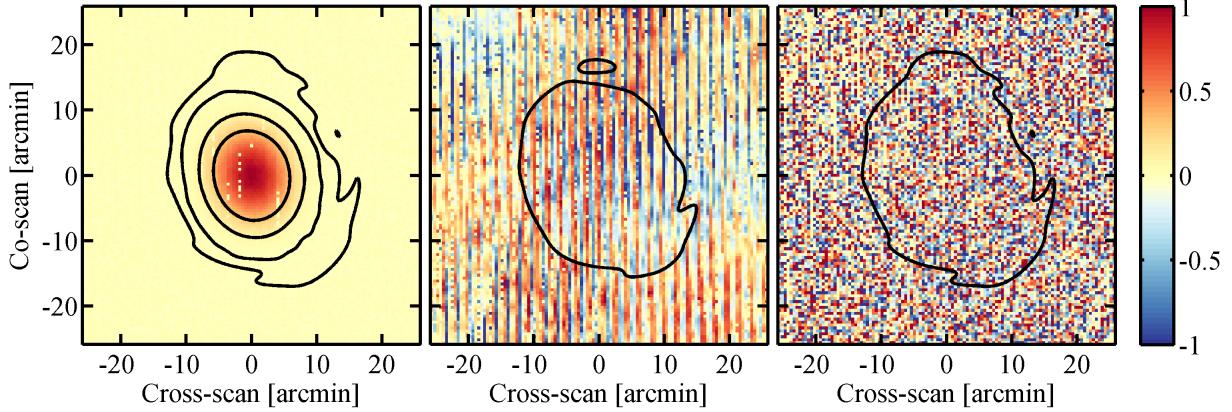


Figure 3.6: *Left*: A simulated 100-1a beam map combined from the four Saturn observations. With contours from the official data derived DX11 beam map overlaid. Beginning from the center, the contours represent -5, -10, -20, and -30 dB regions. *Middle*: The astrophysical background obtained by combining four Saturn observations. The color scale has been inflated by a factor of 300 with respect to the beam map for increased contrast. The contour line shows the region of the beam where the beam amplitude is comparable to the astrophysical background. *Right*: The combined noise map inflated by a factor of 500 with respect to the main beam for increased contrast. The contour shows the region of the beam where the beam amplitude is comparable to the noise RMS.

Noise

Noise can contribute significantly to the beam reconstruction error. For *Planck* HFI, the band averaged Noise Equivalent Temperature (NET) is in the range of $50\text{--}90\,\mu\text{K}\sqrt{\text{s}}$ for 100–217 GHz [138]. During a single planet observation, a detector will spend no more than 5 seconds inside a one FWHM diameter disk centered on the planet. Because of this restricted observation time, noise can limit our ability to probe the shoulder of the beams where the signal to noise ratio becomes low. We use a planet signal subtracted timeline as a baseline for noise estimate. As expected, the noise timelines generated this way do in many cases show clear signal of 4 K cooler lines [138]. It is crucial that noise realizations be generated from data, as the amplitude of 4 K cooler lines, and residuals thereof, vary with time.

These noise Power Spectral Densities (PSDs) are subsequently used as templates for Gaussian noise realizations which are added to the simulated timestream. Comparison with officially vetted noise PSDs shows that the resultant noise RMS generated by our algorithm

is consistent at the 10% level for all detectors. In the presence of noise, the generated signal becomes

$$s(\hat{\mathbf{n}}_i) = C_p \Psi(\hat{\mathbf{n}}_i) + n_i, \quad (3.22)$$

where we assume that the noise contribution, n_i , is uncorrelated with the signal amplitude so that $n(\hat{\mathbf{n}}_i) = n_i$. The simulations inject noise from 1 Hz up to the 90 Hz Nyquist frequency, where the noise has begun to roll off due to filtering. The simulations do not inject low frequency noise below 1 Hz; neither does the subsequent analysis subtract low frequency noise components from the timelines. The noise injected in these simulations is assumed to be uncorrelated among detectors.

Figure 3.6 shows a decomposition of signal, astrophysical background, and noise obtained by combining four Saturn observations. The contour of the beam where the noise RMS becomes comparable to the signal from the planet observations is shown for comparison. We did not include low frequency noise in the simulations used to generate these diagrams, hence the trivial morphology of the noise map. In contrast, the astrophysical background map clearly indicates complex morphology. Any error in background estimates will retain a fraction of this structure.

To understand noise contributions to the beam reconstruction error we can run simulations with noise as the only injected non-ideality. Figure 3.7 shows the bin-by-bin (1σ) standard deviation on the scanning beam window function as derived from the first observations of the three main beam calibrators, Mars, Jupiter, and Saturn. This particular simulation used 500 realizations of each observation, making sample variance of the underlying distribution negligible. Even using Saturn, the noise error can affect the beam window function at the 1% level for $\ell \sim 2000$. The Jupiter observations are clearly not noise limited. As the beams for the 2013 *Planck* data release were derived using only the first two Mars observations, we conclude that the beam determination was likely noise limited. As the 2014 release uses both Saturn and Jupiter observations [229], it is clear that statistical noise now plays a less dominant role in the beam reconstruction error.

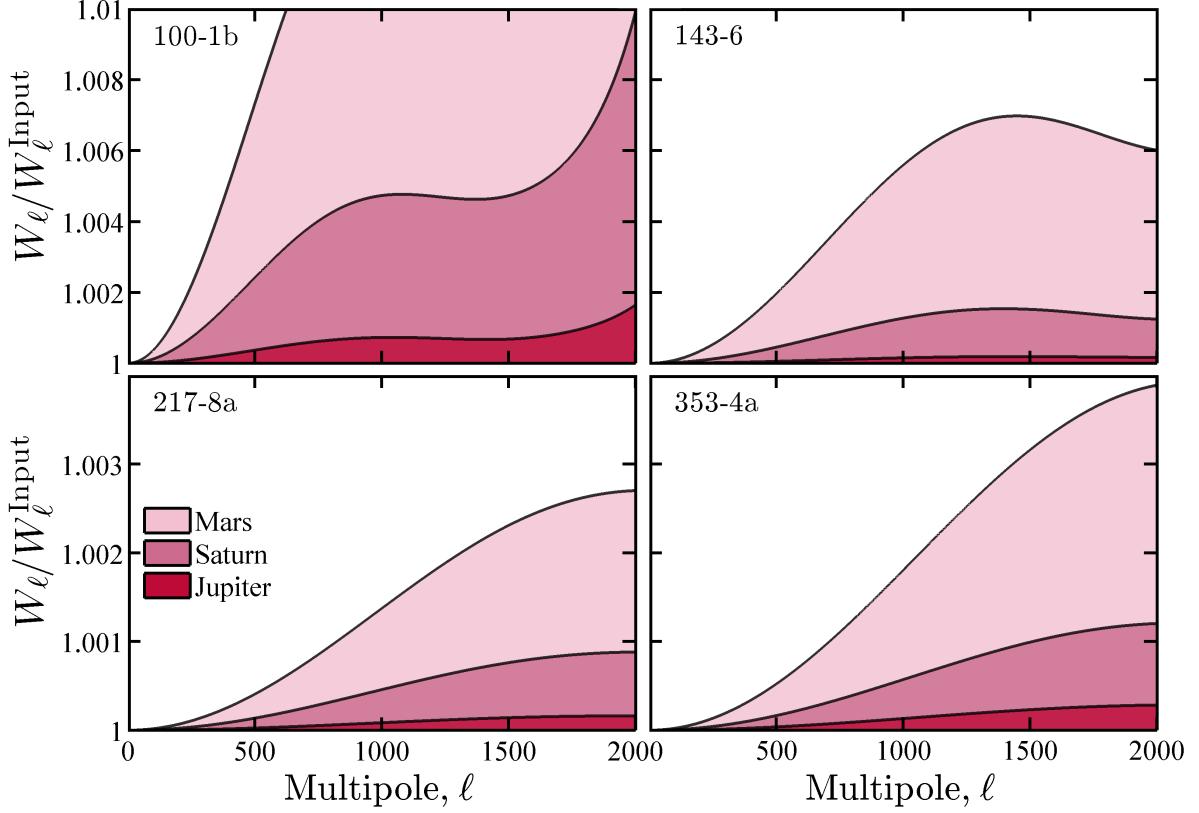


Figure 3.7: The 1σ statistical error in scanning beam window function reconstruction (compared to the input beam) based on one observation of Mars, Jupiter, and Saturn. Presented here are the results for four detectors, 100-1b, 143-3a, 217-8a, and 353-4a. Note the change in vertical scale between rows. The effect of noise is clearly strongest in beam reconstruction using Mars observations. It would be almost negligible if we were to use Jupiter for main beam reconstruction.

Pointing Error

The satellite pointing reconstruction uses data from two time-delay integration star cameras that are mounted to the side of the service module. The satellite also houses gyroscopes which provide ancillary pointing information. On the time-scales of days, pointing reconstruction is verified by comparing beam centroids derived from observing bright point sources with locations predicted by the JPL ephemerides database.

The five outer planets and some radio sources are primarily used to verify low frequency pointing stability. Initially, this process revealed a strong correlation with temperatures on the service module in the proximity of the star cameras. It was concluded that thermoelastic

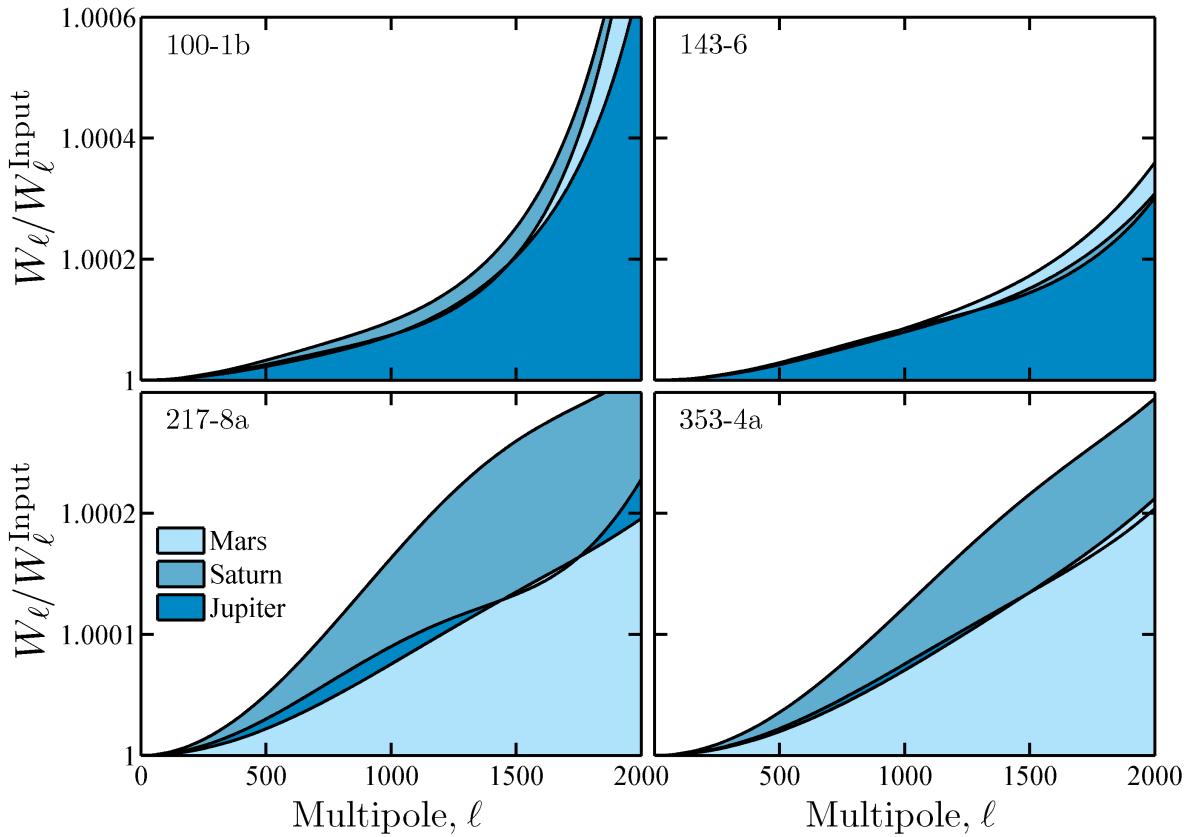


Figure 3.8: Beam reconstruction error for detectors 100-1b, 143-6, 217-8a, and 353-4a due to random pointing jitter. Note the change in scale between the two rows. We see that the errors are more or less consistent between observations.

deformations, caused by cycling of hydrogen sorption coolers, affected the star cameras, resulting in few arcsec errors in the pointing which were subsequently corrected.

When generating simulated timelines, the real pointing timelines are used to sample a fiducial input beam. Any pointing error is then added to the pointing solutions. This allows us to accurately simulate realized sampling densities as well as any other features in the real pointing timelines. Similarly, samples that are flagged for various reasons in the real data timelines are also flagged in simulations. For example, if there is an unusually high density of flagged data close to the beam center, the simulations will incorporate such an effect.

In our notation, a pointing error causes our assumed pointing $\hat{\eta}_i$ to be offset from the true

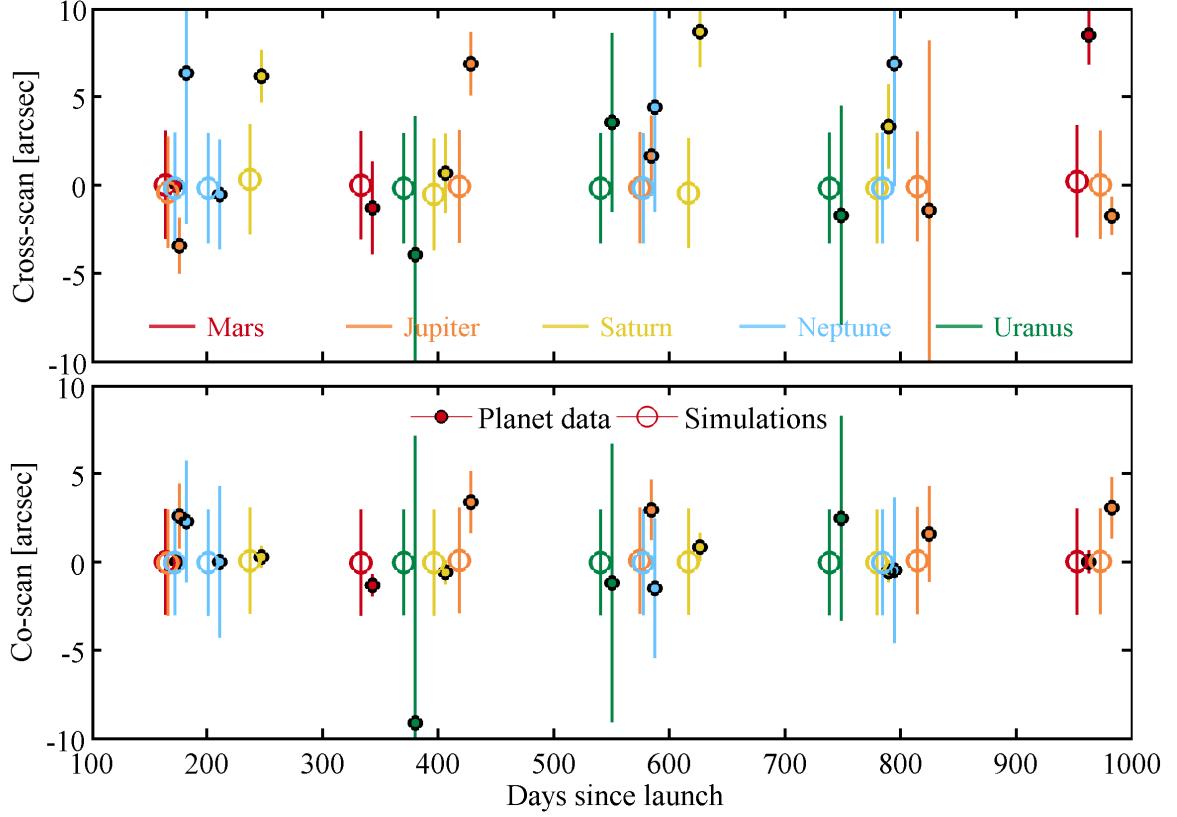


Figure 3.9: This figure shows the co- and cross-scan centroid offset and corresponding error averaged over all HFI detectors as derived from DX11 data. Simulation results are also shown for comparison. The vertical lines represent the standard deviation in the ensemble. The results suggest few arcsec pointing errors on month time scales. Clearly cross-scan errors are larger. This is understandable given *Planck*'s raster-like scan strategy. Different colors represent different planets.

pointing $\hat{\mathbf{n}}_i$ by a small amount $\delta\hat{\mathbf{n}}_i$, i.e. $\hat{\mathbf{n}}_i = \hat{\mathbf{n}}_i + \delta\hat{\mathbf{n}}_i$, with the corresponding signal offset

$$\Delta s(\hat{\mathbf{n}}_i) = s(\hat{\mathbf{n}}_i) - s(\hat{\mathbf{n}}_i + \delta\hat{\mathbf{n}}_i) \approx -\nabla s(\hat{\mathbf{n}}_i) \cdot \delta\hat{\mathbf{n}}_i. \quad (3.23)$$

Clearly, pointing errors will broaden beam estimates.

Figure 3.8 shows beam reconstruction from simulations that include 2.5 arcsec RMS pointing error as the only injected non-ideality. The injected pointing error is consistent with that assumed for the 2013 papers [207]. Taken at face value, pointing jitters cause a 1σ shift of 0.04% at $\ell = 2000$ at 143 GHz. This is clearly a small effect.

Traditionally, beam centroids have been derived using parameters of an elliptical Gaussian fit. To probe the fidelity with which we can reconstruct the true centroids of those observations, we simulate all twenty *Planck* observations of the outer planets. Each planet observation is simulated a thousand times. We then fit an elliptical Gaussian beam to the simulated data to determine the beam centroid. Figure 3.9 shows the cross- and co-scan centroid errors and biases, averaged over all frequency bands, derived from simulations of the planet observations and compares those to centroids derived from real data.

The simulations do not show variations in the pointing of the order suggested by the observation averaged centroids. It is possible that correlated errors in the pointing realizations would improve the consistency. Data derived results suggest that the assumed 2 arcsec co-scan pointing errors are conservative.

Coverage Effects

Figure 3.10 shows a coverage map for both a single Jupiter observation and one which is obtained by combining four Saturn observations. The color of each pixel in those maps represents the number of times a sample was drawn from that location. The pointing data have been corrected for the peculiar motion of Jupiter so that a planet centered map can be constructed. Note how the number of data samples falling within a single pixel varies by a factor of two in the 600 arcmin² field of view. If these beam reconstruction simulations were taking place before the satellite acquired data, we would have to assume that the planet center could fall at an arbitrary location on that coverage map.⁹

The coverage maps shown in Figure 3.10 are made less uniform by a number of effects, including cosmic rays, flagging algorithms, and scan non-uniformity. Note that two beam contours are shown in each panel, one of which has been translated by 9 arcmin along both co- and cross-scan directions. The number of samples falling within the 20 dB contours differs significantly for the two cases. Beam reconstruction algorithms can be sensitive to

⁹This is similar to the concept of cosmic variance.

such variations. However, now we know how the planet center falls with respect to *Planck*'s anisotropic coverage. Such variations in coverage should therefore not be included in an a posteriori analysis.

However, there are counterarguments. A number of data processing steps are currently not included in any beam simulation pipelines. For example, the *Planck* data are heavily flagged. It is likely that some data samples are flagged when they shouldn't be, and conversely, some data are not flagged when they should. This will affect coverage. Flagging algorithms are not probed in any simulations and shifting the beam center with respect to the coverage map does in some sense probe sensitivity to such effects. Other data pre-processing steps are also not simulated. Finally, this procedure probes the covariance between the elliptical Gaussian parameters and the decomposed Gauss-Hermite coefficients. We therefore choose to include this effect in our simulations.

Astrophysical Background Residuals

The beam reconstruction uses data that combine signals from planets as well as the CMB and any other astrophysical background. Estimates for the latter are subtracted from timelines prior to beam reconstruction. Inaccurate background estimates can therefore affect the beam reconstruction. Figure 3.6 shows the astrophysical background obtained by combining four Saturn observations. Striping and sub-degree scale anisotropies are visible. Signal errors will most likely have similar morphologies. We simulate inaccurate background estimates by assuming some fraction of the background remains after subtraction.

Input Beam

Simulation results depend on the input beam. Ideally, the input is almost identical to the true beam response of the detector being simulated. Obviously, our ability to choose an input beam is limited by our understanding of the real beam response. Additionally, simulations require an input beam response function that can accept arbitrary pointing information. A simple

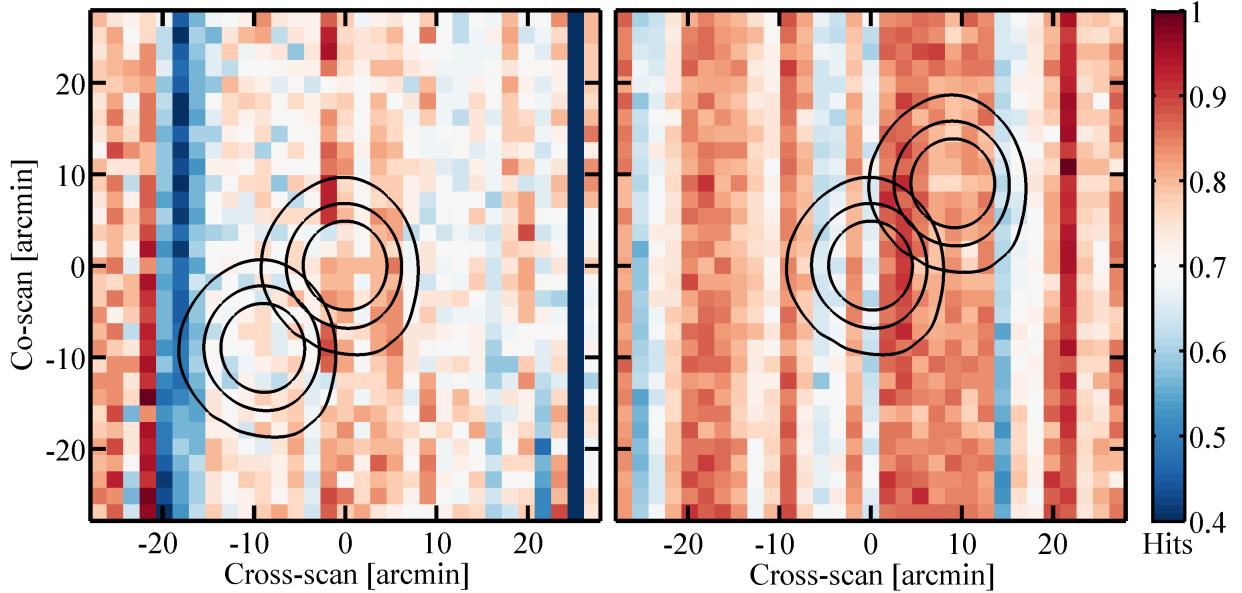


Figure 3.10: *Left:* The outline of the DX11 beam product for channel 143-7 overlaid on a coverage map from the first Jupiter observation. The three contour lines represent the -5, -10, and -20 dB curves. Another version of the same beam is shown translated by 9 arcmin in both co- and cross-scan direction. The coverage is normalized to unity, with 55 samples falling on the pixel that is most often hit. The pixels are 1.8 arcmin wide. *Right:* The outlines of the DX11 beam product for 143-7 overlaid on a coverage map which this time is obtained by stacking all four Saturn observations. There are 196 samples populating the best covered pixel.

elliptical Gaussian input beam is more easily described by a multivariate beam model than a beam that is the result of complicated interplay between optical components and subject to non-idealities. An input beam cannot be described using a function basis that is similar to the basis used for reconstruction, as this can result in an unrealistic advantage. Possible parametrization methods include Gaussian or elliptical Gaussian input beam, Gauss-Hermite input beam, B-spline beam, and high resolution rasterized maps with nearest neighbor lookup algorithm.

For the simulation of Gauss-Hermite beam reconstruction, we create a beam product which merges all available Jupiter and Saturn observations. The data are then interpolated using a cubic spline which captures both the high spatial frequency information in our beam as well as the near sidelobes which extend many FWHM out from the beam center.

Time-varying Planet Flux

The mm-wavelength luminosity of Mars is modeled by calculating the intensity of light reflected from regions of the planet with varying albedo [241, 242]. As the planet rotates around its axes with a period of nearly 25 hours, time-varying flux could affect the beam reconstruction. We have simulated Mars observations with and without time-varying flux, and found a small effect that we can correct. At the moment, Mars observations are not used to derive beam products. The author is not aware of models that predict time-variations in the flux of Jupiter or Saturn on hour timescales.

Cosmic Rays – Glitches

Planck HFI is constantly bombarded by Galactic cosmic rays which are modulated by variation in the Sun’s activity.¹⁰ These cosmic rays affect a significant portion of the *Planck* TOD. In fact, at the beginning of the mission, about 98% of the data were touched by a deglitching algorithm which flags, removes, and corrects data due to cosmic rays [229] prior to time response deconvolution.¹¹ This number came down over time as the Galactic cosmic ray incidence fell, yet more than 90% of the *Planck* HFI data have been affected by this algorithm. Glitches and their temporal variations can affect beam reconstruction. Unfortunately, the simulation of glitches and subsequent removal thereof is only possible within the official *Planck* data processing pipeline, as the algorithm is highly specific.

Other Effects

A large number of non-idealities should be incorporated in an ideal simulation pipeline. We have added the most significant of these effects, but some still remain. These include signal non-linearity corrections, gain variations, and destriping and despiking errors, all of which can affect beam reconstruction. For example, low frequency noise is removed from timelines by a

¹⁰The daily fraction of data that are flagged (omitted) as a result of cosmic rays fell from 20% to 13% over the 30 months of observation [138]. This trend is highly correlated with indicators of the Sun’s activity.

¹¹We sometimes use the word despiking to mean the same thing as deglitching.

destriping algorithm. We find that turning this algorithm off, will have a significant effect on the derived beam product. Any variations in the destriping algorithm could therefore affect the final product.

The frequency dependence of the beam and effective area are not probed in these simulations, nor are they fully understood [207]. As Equation 3.17 indicates, the effective area, $A_e(\nu)$, and beam response, $\Psi(\hat{\mathbf{n}}, \nu)$, are frequency dependent. Since the beam response is calibrated using planets, which have a Rayleigh-Jeans like spectrum to a good approximation, the beam response appropriate for CMB observations might be slightly different. Physical optics simulations have suggested that this effect could cause up to 1% error in beam window function at $\ell = 2000$ for the 217 GHz band [229]. This effect adds significant uncertainty to our beam reconstruction.

Reconstruction Bias

Unlike beam reconstruction error, we find that the amplitude and sign of the bias in beam reconstruction (see Equation 3.21) depends strongly on the input beam used for simulations. We iterate the beam input to the simulations with the scanning beam derived from the planetary observations. During this process, we have observed a range of beam reconstruction errors, normally, however, these bias functions are comparable in amplitude to the 1σ error estimates shown in Figure 3.12. As we do not know the exact shape and functional form of the input beam, we choose not to correct for bias when performing beam reconstruction. Comparison of different beam reconstruction algorithms suggests that differential reconstruction bias is small (see Section 3.6.1).

3.5.4 Error Budget

Having probed various systematics, we can construct a window function error budget. This quantifies the relative contribution of different systematics to the total beam reconstruction

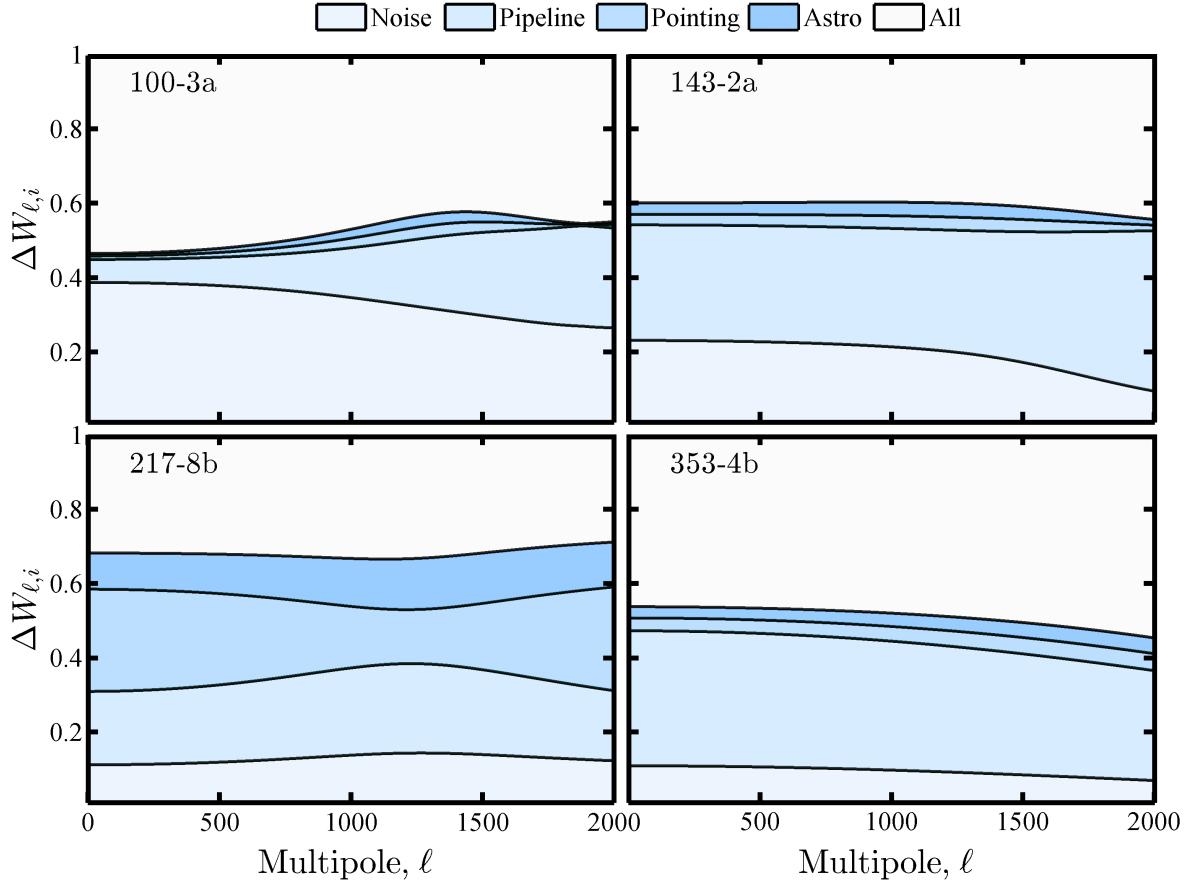


Figure 3.11: The fractional contribution of various non-idealities to the total window function error budget as derived for a stacked beam product using all available Jupiter and Saturn observations; shown here for four detectors. The white region represents the relative window function error when all non-idealities are added at the same time. If they were completely uncorrelated, the white region would begin at 0.5 for all multipoles.

error. In the limit of uncorrelated errors, the total window function error follows

$$\Delta W_\ell = \sum_j^{n_{\text{err}}} \sigma_{\ell,j}, \quad (3.24)$$

where $\sigma_{\ell,j}$ represents the beam window function error from one of the n_{err} types of systematics probed by simulations. Note that the window function is proportional to the square of the spherical harmonic transform of the beam response. The above expression therefore corresponds to a quadrature sum of errors in the decomposition coefficients. To see how the errors correlate, we can run simulations which incorporate subsets of errors.

Figure 3.11 shows the window function error budget of the Gauss-Hermite beams as derived from a stacked beam product that includes the first four observations of both Saturn and Jupiter. The simulations used to derive these results include 1000 realizations for every non-ideality, as well as 1000 realizations with all effects added at the same time. The total length of the timestreams generated this way corresponds to approximately 20 days of integration time per detector. We simulate noise, coverage, pointing, and error in the subtraction of astrophysical background.

The error budget shows how the detector (statistical) noise affects the beam reconstruction error primarily at the lower frequencies. For 217 GHz and up, other errors are larger. It is interesting that the error induced by coverage effects dominates in many cases. Pointing error and incorrect subtraction of astrophysical background also affect the derived beam window function. The white region in Figure 3.11 shows the relative size of the reconstruction error when all non-idealities are included at the same time. In the limit where all errors are uncorrelated, the white region would cover half of each graph with a lower boundary at 0.5. The error correlation structure is noteworthy. For example, we see that the errors interfere constructively for channel 353-4b such that the white region covers more than 50% of the area at high multipoles.

3.5.5 Principal Component Analysis

A complete discussion of the *Planck* HFI cosmological analysis is found in [49, 64]. The following excerpt describes how the beam window function error estimates are propagated through the cosmological error analysis.

We performed Principal Component Analysis (PCA) on the Gauss-Hermite covariance matrix to see whether reconstruction errors can be captured by a subset of the GH parameters. This analysis does not suggest a significant reduction in matrix rank.¹² Contrary to what is

¹²Although principal component analysis does not effectively compactify the real space variance of the Gauss-Hermite simulations, such parametrization can be useful for error analysis in physical optics simulations where one needs to generate an ensemble of beam realizations to compare against mirror deformations.

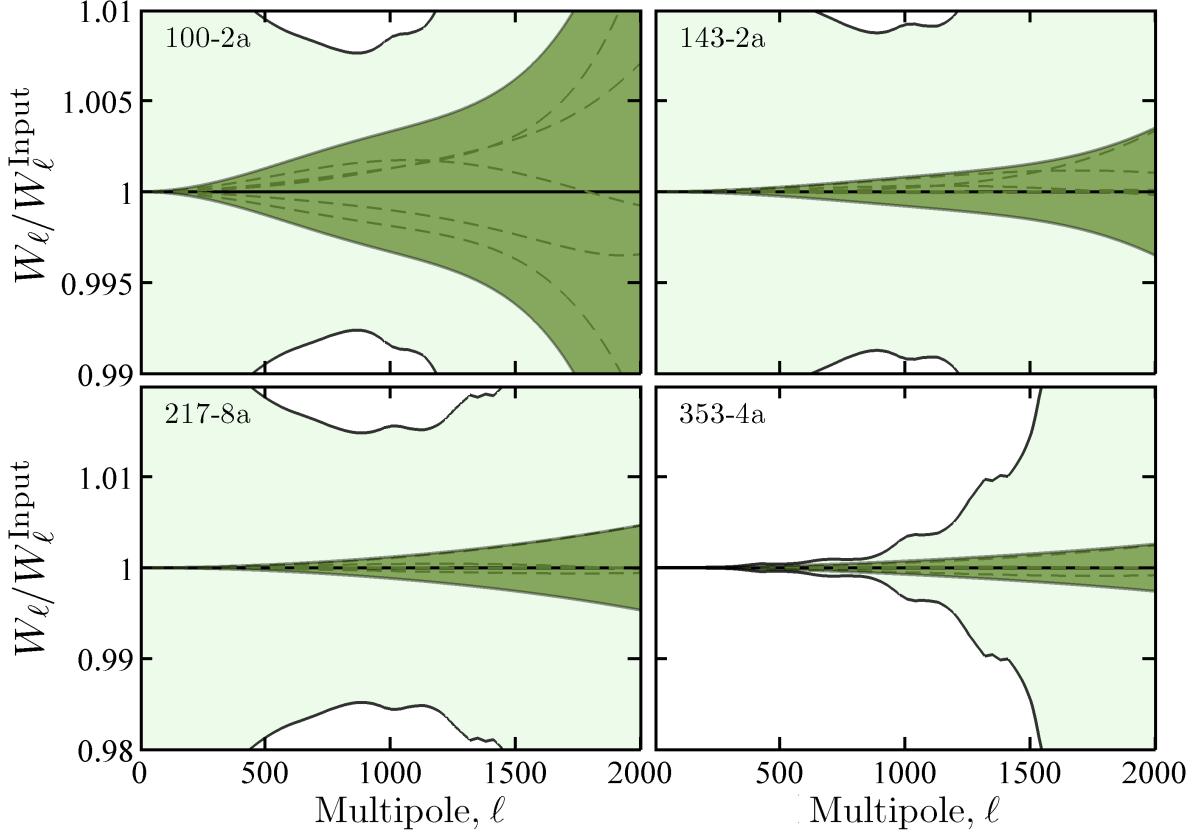


Figure 3.12: Total beam reconstruction error for detectors 100-2a, 143-2a, 217-8a, and 353-4a as derived from all available observations of Jupiter and Saturn. Green region is the 1σ error envelope. The first five eigenmodes, scaled with their eigenvalue, are shown with green dashed lines. The light green region represents the combined error due to noise and cosmic variance calculated using the standard Fisher matrix approximation [133] while setting bin width and sky fraction that are consistent with [49, 64]. Channel 353-4a only shows the noise contribution.

seen with the GH coefficients, the window function error covariance matrices, established by simulations, are smooth. We find that the first few eigenmodes capture a significant fraction of the error (see detailed discussion in Appendix A.6 of [64] and Appendix G of [207]).

The beam window function error can be written as a sum of orthonormal error eigenmodes, $V_{\ell,k}$, multiplied by their respective eigenvalue

$$\Delta W_\ell = \sum_{k=1}^{n_{\text{MC}}} \lambda_k V_{\ell,k}, \quad (3.25)$$

where n_{MC} corresponds to the number of Monte-Carlo realizations. Because of low dimensionality, the first few eigenmodes of the covariance matrix capture a large fraction of the variance. For that reason, the eigenmodes used for cosmological analysis are truncated to a small number. We write

$$\Delta W_\ell \approx \sum_{k=1}^{n_m} \lambda_k V_{\ell,k} \quad \text{with} \quad \sum_{k=1}^{n_m} \lambda_k \gg \sum_{k=n_m+1}^{n_{MC}} \lambda_k, \quad (3.26)$$

where $n_m = 5$ represents the number of beam window function error modes used in cosmological analysis.

With these tools, we can perform statistical tests using the total window function error derived from the Gauss-Hermite pipeline. Figure 3.12 shows the total beam window function error and eigenmodes for four detectors as derived using all available Jupiter and Saturn data. These are based on 1000 realizations per detector, using all four non-idealities described in Section 3.5.4. The 1σ error envelope is represented by the brown region. The first five eigenmodes, scaled with their eigenvalue, are shown by the dashed black lines. Note that the error envelope is obtained by adding the first five eigenmodes in quadrature.

The *Planck* HFI analysis incorporates errors from beam reconstruction simulations into a pseudo-spectrum likelihood approach.¹³ The beam error eigenmodes are used to construct a generalized beam eigenmode appropriate for the 100, 143, and 217 GHz auto- and cross-spectra. By propagating the beam errors through the entire cosmological analysis pipeline, we can determine how the errors affect the final cosmological parameter estimates.

3.6 Data Derived Beams

In the first years of beam analysis, it was found that beams derived from Saturn and Jupiter observations differed at the percent level in solid angle from ones derived from Mars

¹³Many likelihood pipelines exist within the *Planck* HFI collaboration. The one that was used to derive the 2013 cosmological results is referred to as **CamSpec** and is maintained by George Efstathiou's group at Cambridge.

observations alone. We thought that this could be due to poorly understood phenomena associated with the Jovian planets. At the same time, the first Mars observation was thought to provide greatest beam reconstruction fidelity due to high sample density. Finally, no beam reconstruction algorithm which effectively combined many planet observations into a single beam product appropriate for cosmological analysis had undergone sufficient characterization. It was therefore decided that the beam product for the 2013 *Planck* release would only use data from the first two Mars observations [207].

The final goal, however, had always been to use all of the planet observations to constrain beams. In the 2014 data release, the official beams are derived using a hybrid beam model which combines Saturn and Jupiter observations using a B-spline parametrization for the main beam and an azimuthally symmetric power law added at large angular separation [207, 229].¹⁴ As Saturn is significantly brighter than Mars, little is gained in signal to noise ratio by adding Mars to the merged beam product. Currently, Mars observations are only used for validation purposes.

The act of merging planet observations requires accurate centering and normalization. We incorporate this into the Gauss-Hermite beam reconstruction algorithm. The simulations allow us to estimate the accuracy with which we can perform these operations. In the following subsections, we look at Gauss-Hermite beams reconstructed from real data. We show that these products are internally consistent and compare them with the B-spline results.

3.6.1 Pipeline Comparison

Comparative studies of the B-spline and Gauss-Hermite beams help spot outliers and improve the quality of both algorithms. When analyzing the same data, the two beam reconstruction

¹⁴The planets are not perfect point sources, especially Jupiter. We correct for the finite size of the planets in ℓ -space [207]. This leads to an approximately 0.1% correction of the CMB channel window functions at $\ell = 2000$.

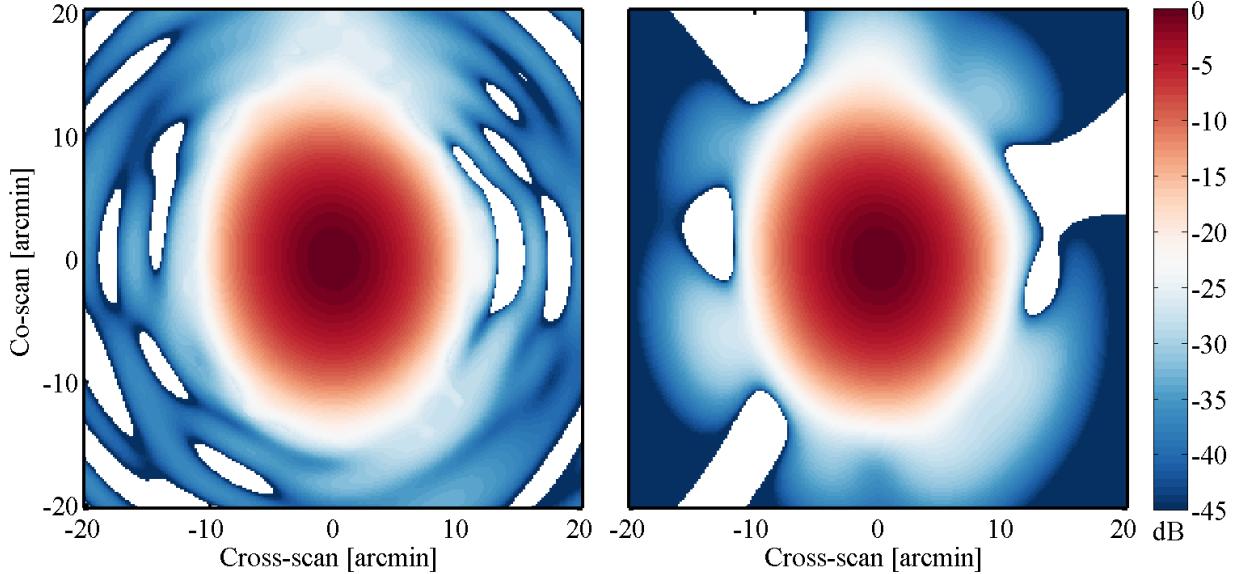


Figure 3.13: The channel 100-2b scanning beam derived from a combination of Jupiter and Saturn observations using a B-Spline hybrid basis (left) and a Gauss-Hermite beam (right). The beams are shown here on a log scale to accentuate differences in the shapes. Note that the B-Spline beam has more variations on small angular scales.

methods should give similar answers. With simulations available for both pipelines, the consistency tests on data derived products become more robust.

When comparing beam models, simulations suggest that the B-spline parametrization captures the main beam response with greater fidelity than the Gauss-Hermite algorithm. Based on those results, and since the B-spline pipeline was fully integrated with the official HFI data processing infrastructure, it was decided that this method should be used to derive the *Planck* HFI beam parameters. The B-spline pipeline as it pertains to *Planck* HFI beams is described in [207, 235].

The *Planck* hybrid beam model is described in [229]. The hybrid model incorporates planet data extending out to 100 arcmin from the beam center, while the Gauss-Hermite pipeline, with $N_{\max} = 9$, can only extend out to 20 arcmin (see Section 3.4). For this reason, we add a Ruze envelope to the Gauss-Hermite beam to capture structure in the near sidelobes (see Section 3.6.3). Figure 3.13 shows a scanning beam map of channel 100-2b, reconstructed

Table 3.3: The mean value of DX11 scanning beam parameters for each HFI band. The beam full-width at half maximum, θ_{FWHM} , is derived from an elliptical Gaussian fit. Two values for the beam solid angle are reported; the first is the solid angle obtained from the B-spline hybrid model, Ω_{BSH} , the latter, Ω_{GHR} , is obtained by combining a Gauss-Hermite beam with a Ruze envelope for the near sidelobes (see Section 3.6.3). The ellipticity, ϵ , defined as the ratio of the larger beam width and the smaller beam width, is determined from an elliptical Gaussian fit to data (out to 30 arcmin) that combine Jupiter and Saturn observations.

	θ_{FWHM} [arcmin]	Ω_{BSH} [arcmin 2]	Ω_{GHR} [arcmin 2]	ϵ
100 GHz	9.55 ± 0.03	105.0 ± 1.0	104.9 ± 1.0	1.21 ± 0.01
143 GHz	7.11 ± 0.11	58.6 ± 1.9	58.7 ± 1.8	1.05 ± 0.02
217 GHz	4.88 ± 0.05	27.1 ± 0.6	27.3 ± 0.7	1.21 ± 0.03
353 GHz	4.65 ± 0.05	24.9 ± 0.6	24.9 ± 1.2	1.19 ± 0.03
545 GHz	4.68 ± 0.34	25.5 ± 3.0	26.0 ± 3.2	1.18 ± 0.12
857 GHz	4.29 ± 0.03	23.0 ± 0.8	27.8 ± 1.0	1.40 ± 0.11

from combined observations of Jupiter and Saturn. The beam maps suggest the B-spline allows greater spatial frequency reconstruction.

Comparison of all HFI beams reveals good consistency between the two beam products. For example, the band average solid angles for the two products, agree to within 0.2%, 0.15%, 0.3%, at 100, 143, and 217 GHz respectively. Table 3.3 describes some properties of the two products as derived on data appropriate for the 2014 cosmology release (DX11). The B-spline product has a slightly larger solid angle in all cases.

It is instructive to look at azimuthally averaged beam profiles. Despite the scan-strategy, the data offer tremendous signal-to-noise ratio. Figure 3.14 shows the azimuthally averaged data obtained by combining all available Jupiter and Saturn observations. Blue points represent average values within a bin. Light-blue colored points emerge at the noise floor, representing negative values. The blue curve in these plots is the beam profile used for the 2013 cosmological analysis (DX9). Note that the beam data were apodized, beginning at roughly 12 arcmin. The 2014 beam profiles (DX11) are shown in orange. We also include the

predicted beam profile derived from physical optics simulations in light-blue. Finally, a hybrid beam using Gauss-Hermite functions for the main beam and a Ruze envelope at large angular separations is shown with a red dashed line. The Gauss-Hermite main beam, sans Ruze envelope, is shown with a solid red line. Section 3.6.3 describes the Ruze parametrization. Clearly, the DX11 and hybridized Gauss-Hermite profiles agree well.

Two beams with identical solid angles can have very different window functions. Comparison of the beams in multipole space is therefore warranted. Figure 3.15 compares the window function ratios of a few channels that are representative of the average consistency within a frequency band. The error envelope for the Gauss-Hermite pipeline is shown. The B-spline simulations appropriate for DX11 data are not yet available, precluding quantitative consistency checks. We note that the two window functions appear to be consistent at approximately the 1σ error envelope of the Gauss-Hermite beam. The low- ℓ window function difference is intriguing. This is likely caused by differences in the estimates of the near sidelobes. Neither simulation pipeline incorporates the near sidelobe reconstruction. We expect that the low- ℓ reconstruction error is underestimated.

3.6.2 Consistency Tests

We can use the statistical tools presented in Section 3.5.5 to perform consistency tests of the GH beams using subsets of planet observations included in the final beam product. Comparison of data-derived and simulation results then provides a quantitative statement about beam consistency. The test involves calculating a probability to exceed for the following reduced chi-squared statistic

$$\chi^2 \equiv \frac{1}{n_{\text{ev}}} \sum_{k=0}^{n_{\text{ev}}} \left(\hat{\lambda}_k / \lambda_k \right)^2, \quad (3.27)$$

where $\hat{\lambda}_k$ are the best-fit values in the n_{ev} element eigenvector decomposition required to construct the discrepancy between the two observation subsets, and λ_k are the eigenvalues derived from simulations. We split the planet observations into two independent subsets, each

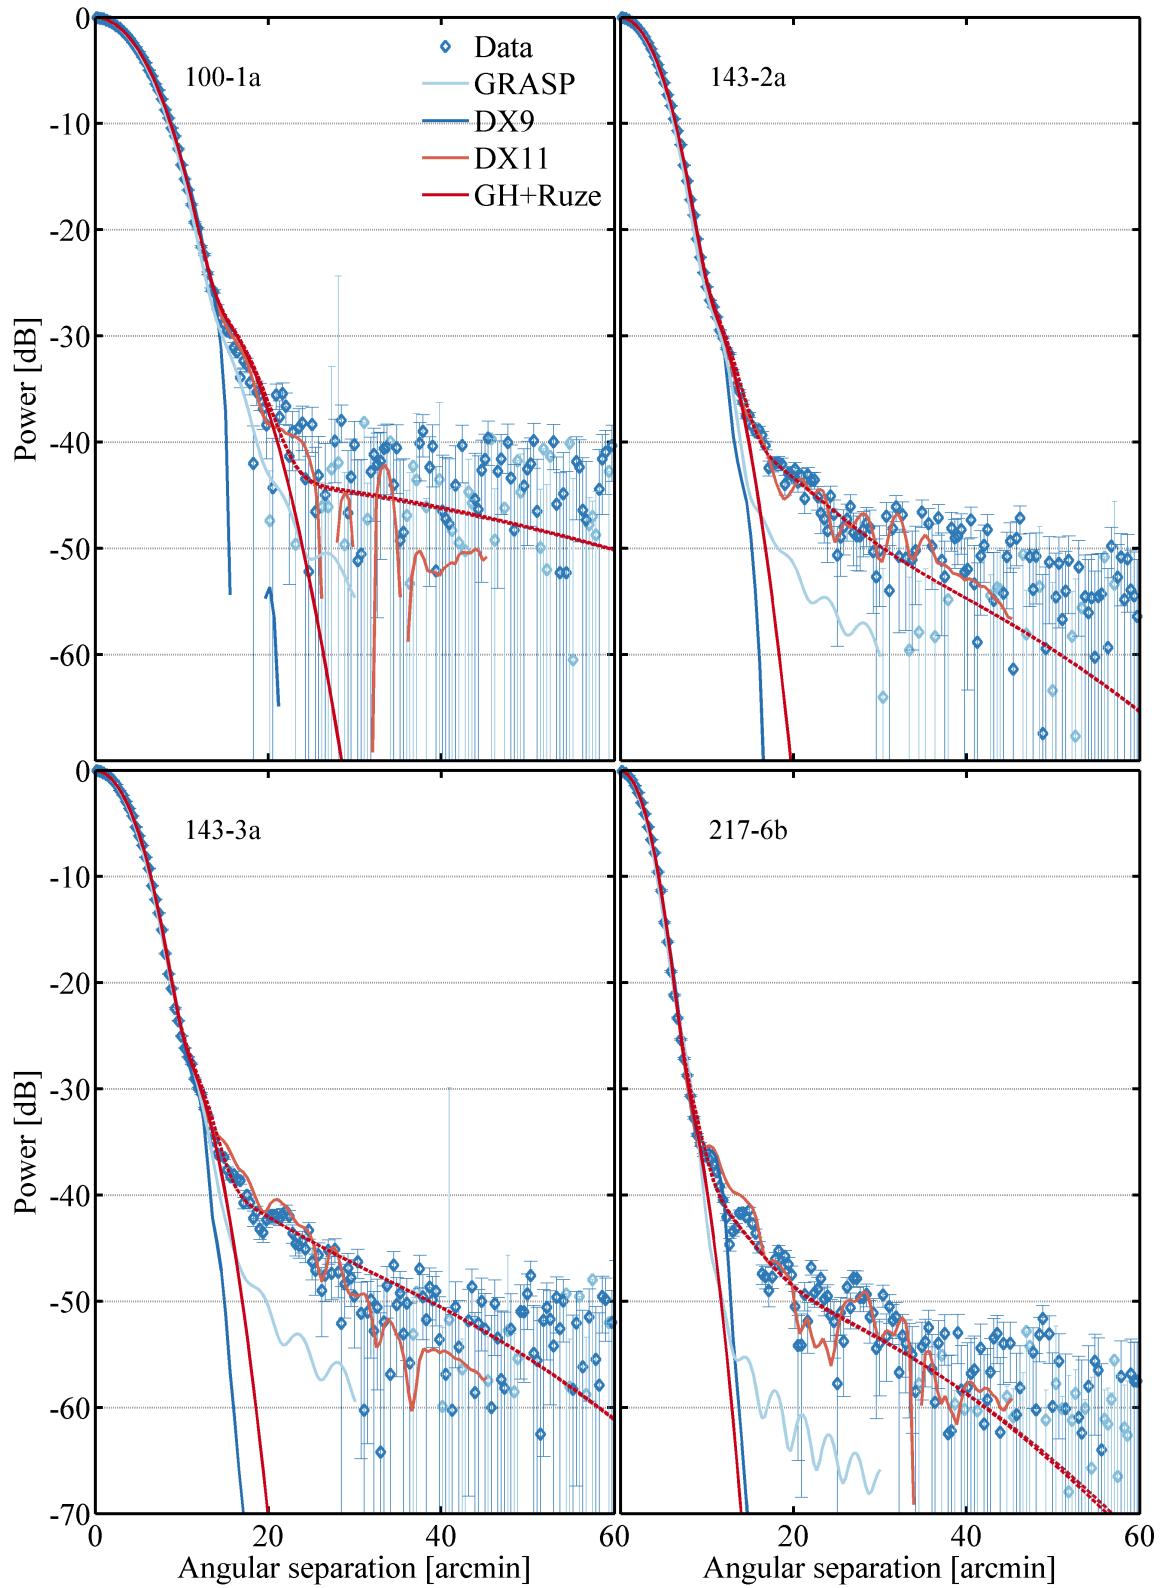


Figure 3.14: The azimuthally averaged beam profiles for four detectors spanning 100–353 GHz. See description of plots in Section 3.6.

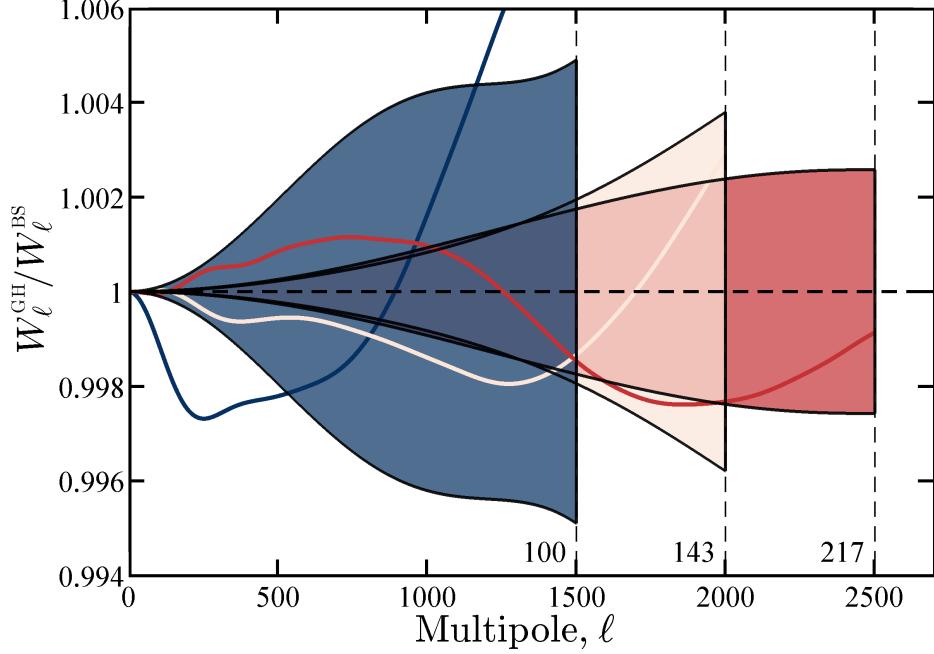


Figure 3.15: Comparison of Gauss-Hermite (numerator) and B-spline (denominator) beam window functions derived using all available Jupiter and Saturn data. The three channels shown here, 100-3a, 143-2b, and 217-6b, are representative of the overall agreement between the two pipelines. The colored region represents the 1σ error envelope for the Gauss-Hermite pipeline derived using simulations.

incorporating four planet observations. The first subset uses the first and third Jupiter and Saturn observations while the second subset combines the second and fourth observations. For every realization, i , we calculate the ratio

$$\Delta W_{\ell,i} \equiv \frac{W_{\ell,i}^{13}}{W_{\ell,i}^{24}}, \quad (3.28)$$

where $W_{\ell,i}^{13}$ and $W_{\ell,i}^{24}$ are the window functions derived from the two independent subsets. After performing principal component analysis on an ensemble of ratios generated by simulations, we decompose the data derived window function ratio, $W_{\ell}^{13}/W_{\ell}^{24}$ using the corresponding eigenfunctions, $V_{\ell,k}$. The best fit decomposition minimizes the residual, \mathcal{D}_{ℓ} , defined as

$$\mathcal{D}_{\ell} \equiv \left(\frac{W_{\ell}^{13}}{W_{\ell}^{24}} - \sum_{k=1}^{n_m} \hat{\lambda}_k V_{\ell,k} \right)^2, \quad (3.29)$$

Table 3.4: The band average consistency test results. Here, $\bar{\chi}^2$ and χ_M^2 represent the band average and band median values, while PTE corresponds to the probability to exceed for the band median value. The three bands used for cosmology pass the consistency test. The results suggest that the beam product generated this way are internally consistent.

	$\bar{\chi}^2$	χ_M^2	PTE
100 GHz	2.1	1.5	22%
143 GHz	1.2	0.3	57%
217 GHz	2.3	1.2	28%

over all angular scales used in cosmological analysis. Table 3.4 lists the band average result of this consistency test for the three frequencies used in cosmological analysis. The results suggest that the Gauss-Hermite beams are internally consistent.

3.6.3 Dimpling Lobes and Ruze Envelopes

A clear example of the connection between the reflector surface finish and the optical response can be seen in the sub-mm channels. The high signal-to-noise-ratio, offered in part by the multi-moded nature of the beams, allows us to probe further down the near-sidelobe to reveal a hexagonal pattern of dimples (see Figure 3.16).

Prior to launch, deformations of *Planck* reflectors were measured using interferometric techniques [218]. The results of these measurements are shown in Figure 3.17. It should be noted that only the secondary was cooled to its operating temperature of about 50 K during these measurements. The higher resolution measurement of the secondary clearly indicates 20 μm deformations on 10 mm scales. A hexagonal pattern, corresponding to the carbon-fiber reinforced plastic structure supporting the front-facing surface, can also be seen in the interferogram of the secondary; these cells have a characteristic length of approximately 60 mm. Only coarse resolution measurements exist for the primary. Those indicate deformations of similar amplitude, approximately 20 μm , over the scales probed. The two-dimensional Fourier transform of these deformation maps reveal a hexagonal structure

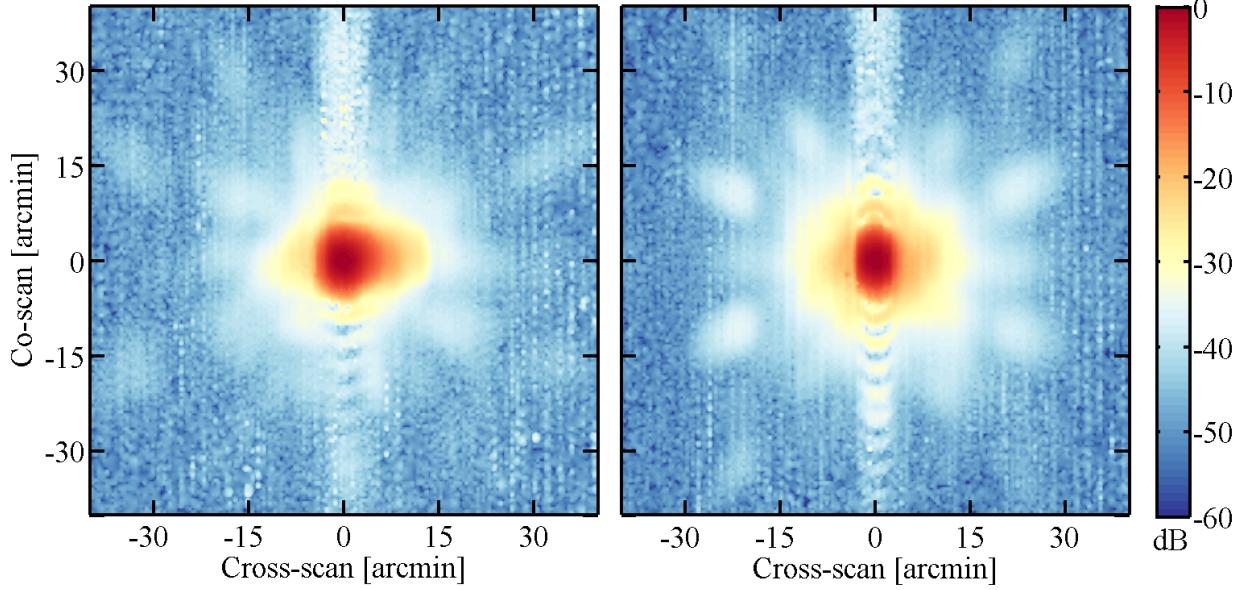


Figure 3.16: *Left:* A beam map of channel 545-1 obtained by combining all available observations of Jupiter and Saturn. *Right:* A beam map of 857-1 obtained through the same process. The non-Gaussian shape of these multi-moded beams is clear. At least three sets of dimpling lobes are also visible. Finally, detector time-response filtering residuals are visible on both sides of planet crossing. Note that data from Jupiter have been excluded in regions where the detector is saturated.

similar to the ones that are observed in Figure 3.16. This is a wonderful example of Fourier optics.

Gaussian surface errors on reflector elements redistribute power from the main beam to larger angles. Errors with large spatial correlation lengths (large dimples), compared to the wavelength of reflected light, will distribute energy to relatively small angular scales. Conversely, small dimples will distribute energy over large angular scales. The beam shoulder generated this way is often called a Ruze-envelope. Ruze derived an expression for loss in antenna gain due to uncorrelated surface errors with a Gaussian distributions of zero mean deformations spanning a range of physical scales. The expression appropriate for the beam response is [243]

$$\Psi(\theta, \phi) = \Psi_0(\theta, \phi) e^{-\overline{\delta^2}} + \left(\frac{2\pi l}{\lambda} \right)^2 e^{-\overline{\delta^2}} \sum_{n=1}^{\infty} \frac{\overline{\delta^2}^n}{n \times n!} e^{-(\pi l \sin(\theta)/\lambda)^2/n}, \quad (3.30)$$

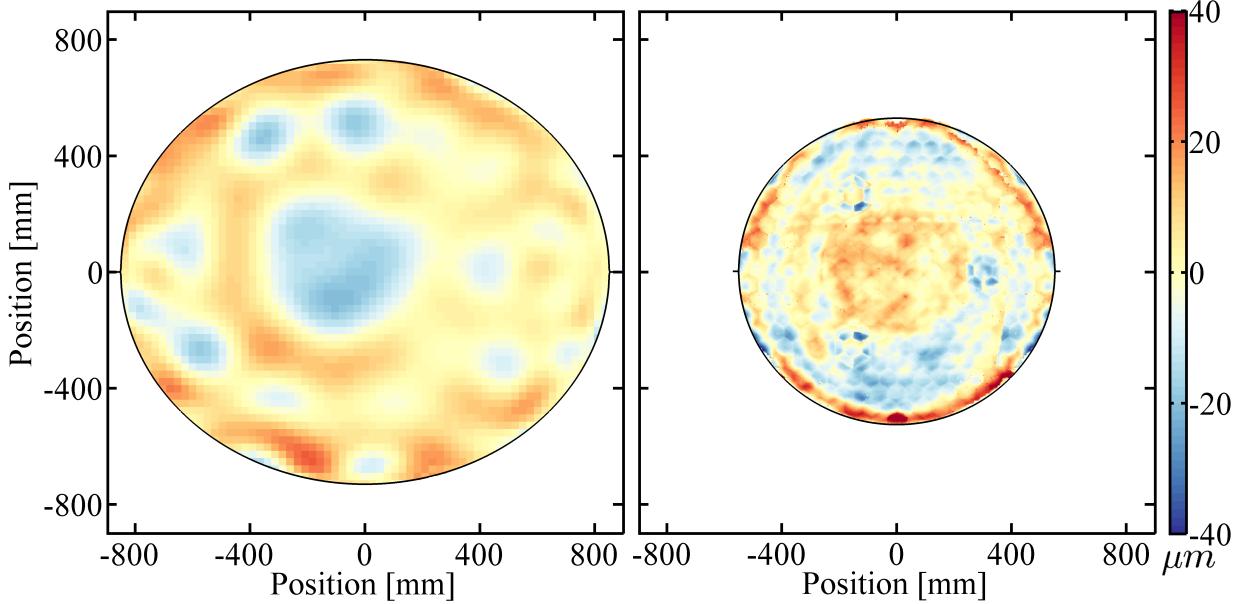


Figure 3.17: *Left:* The *Planck* primary mirror deformations from a photogrammetric measurements with 26 mm resolution. *Right:* The *Planck* secondary mirror deformations observed when cooled to 50 K, corresponding to operational temperatures. Obtained at a maximum resolution of 3–4 mm. Note that in both cases, these measurements are projected onto a plane that is perpendicular to the head-on viewing angle, coincident with the rim of the mirrors. More detail is provided in [218].

where l is the correlation length of the surface deformation, λ is the wavelength, $\Psi_0(\theta, \phi)$ is the ideal beam shape and $\overline{\delta^2}$ represents the variance of the phase errors. The equation is applicable in the limit when $D/(2l) \gg 1$, where D is the diameter of the optical element. According to Equation 3.30, loss in forward gain is mainly determined by the amplitude of the RMS error and not correlation lengths.

Although the above formula serves as an excellent parametrization of near-sidelobes, we have been unable to identify a set of correlation lengths and surface RMS errors that generate near-sidelobe structure that is consistent with our measurement for all frequency bands (see Figure 3.14). This is probably related to variation in reflector illumination and distribution of surface errors.

3.6.4 Bolometer Non-Linearity and Saturation

Of all the planets that *Planck* observes, Jupiter is by far the brightest (see Table 3.2). The flux from Jupiter is so great that it saturates the detectors. Different types of signal processing non-linearities affect the *Planck* data [138, 207]. Bolometric non-linearity corrections are described in [138]. Even with signal non-linearity correction applied, above certain signal amplitudes, the bolometer becomes saturated. A saturated detector will display zero or highly non-linear response to flux variation. As Jupiter is still a point source compared to the HFI beams, it is useful for characterizing the beam response. Because of saturation, data obtained close to the planet center are not fit for use. However, below the saturation threshold, there is little reason to throw away data, and the large signal to noise ratio makes Jupiter ideal for characterizing the near-sidelobes response. We developed a method to normalize Jupiter observations by fitting the response in the beam shoulders to Saturn observations. Since Saturn does not saturate the bolometers, we can estimate the main beam and near sidelobes. By normalizing the shoulders of Jupiter observations to Saturn, we extend the beam response below the noise floor in Saturn timelines. The algorithm is as follows:

- Create a cubic spline interpolation function of two planet observations, one of which potentially suffers from non-linearity.
- Use these two interpolation functions to recreate signal timelines within a region where neither observation suffers from non-linearity. The interpolant function that describes the non-linearity-prone observation is still valid in an annular region centered on the planet.
- Find a normalization constant, C , that minimizes the squared residual of the two observations.

With the last step, one can write

$$\tilde{\Psi}(x, y) = \tilde{\Psi}_{\text{Sat}}(x, y) + C\tilde{\Psi}_{\text{Jup}}(x, y), \quad (3.31)$$

where $\tilde{\Psi}(x, y)$ represents an approximation for the beam response using both Saturn and Jupiter data. This method indicates detector saturation for Jupiter observations above 217 GHz. Analysis of the first three Jupiter observations suggested 7, 9, and 70% clipping at 353, 545, 857 GHz, respectively. This suggests that the maximum filtered signal amplitude for these beam maps is between 7–70% lower than it should be, if the detector hadn't saturated. Independent analysis on the same dataset performed by Brendan Crill suggested 12, 12, and 66% clipping at 353, 545, 857 GHz respectively. A similar algorithm is used to merge Jupiter and Saturn observations. This method also allows the estimation of flux from Jupiter, despite detector saturation.

3.7 Validation of Absolute Calibration using Planets

Planck constrains the flux of the outer planets at mm and sub-mm wavelengths. For planetary science, this helps inform radiative transfer modeling, which in turn constrains atmospheric densities and chemical abundances. Planet observations can also be used to cross-calibrate *Planck* HFI with other CMB experiments. Along those lines, point source flux reconstruction offers one of the only viable checks on the solid angle of the far sidelobes. Finally, sufficiently precise planet models can be used to bracket the absolute calibration of the instrument.

All planet observations are part of the standard raster scan of the satellite; generally, no observation time is specially assigned to planets [223]. As *Planck* is a space based observatory, Earth's atmospheric signal does not add confusion to these measurements. In the analysis presented here, we probe an extensive stretch of frequencies using a consistent method. The results suggests that the absolute calibration of the *Planck* HFI and *WMAP* is consistent at 100 GHz.

3.7.1 Planet Brightness Measurements

It is difficult to make definitive statements about annual variations in planet flux. Some measurements of long-term trends exist [244, 245, 232]; however, measurements of fluxes in the HFI frequency range are mostly sporadic. Assuming constant brightness, planet observations allow cross calibration of different experiments. By comparing Mars brightness derived by *Planck* HFI and *WMAP*, one can probe systematics in beam modeling and the consistency of absolute calibration. Planet flux as measured by *WMAP* and ACT is described in [246] and [247, 248] respectively. Many other CMB observatories have also reported planet fluxes [249, 250].

Assuming planetary modeling is robust, measurements of planet flux can be used to calibrate the absolute responsivity of the HFI; this would replace calibration using the CMB orbital dipole. However, all models of planet flux in the mm and sub-mm range quote approximately $\gtrsim 5\%$ errors, which is worse than what can be achieved from the CMB dipole below 545 GHz [224]. At and above 545 GHz, however, the CMB dipole signal is too weak.

Shortly after the first data were acquired, comparison of Uranus and Neptune flux measurements with models of the brightness temperatures suggested that the absolute calibration of sub-mm channels, as determined from cross-calibration with FIRAS measurements, was inaccurate. As instruments onboard *Planck*'s sister experiment, the *Herschel* satellite, base their absolute calibration on observations of Uranus and Neptune, it was decided that the absolute calibration of the *Planck*'s sub-mm channels would be linked to those planet models. The absolute calibration is obtained by scaling measurements of Uranus and Neptune flux so that they are consistent with ESA models [251] (see further discussion in [224, 226] and in Section 3.7.5).

Accurate calibration through planet flux measurements allows us to verify foreground models such as the *Planck* sky model [129]. Such models include contributions from sources such as the cosmic infrared background and Galactic foregrounds. As accurate foreground

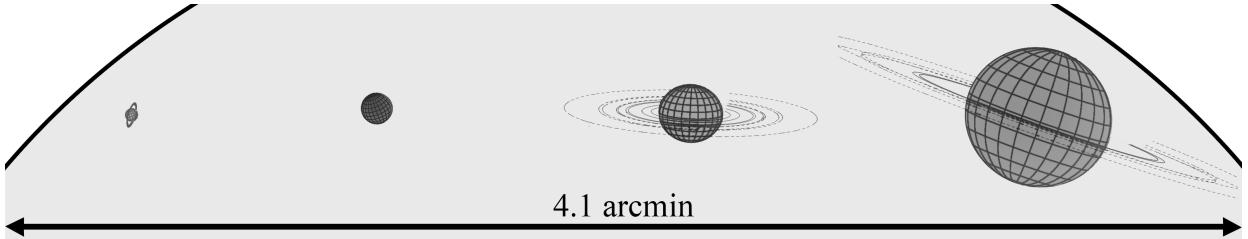


Figure 3.18: Composite showing four planets as viewed from the center of Earth around the turn of year 2010. The view from L₂ would be very similar. From left to right, the figure shows Uranus, Mars, Saturn, and Jupiter. The planets are drawn on the same scale with Jupiter approximately 40 arcsec in diameter. A fraction of a typical 143 GHz beam is shown by the black line and the grey region, the diameter of the circle corresponds to the FWHM of the beam. Planet diagrams extracted from [252].

estimates are required for cosmological analysis, high-fidelity planet flux reconstruction is valuable.

Planet flux can be estimated in different ways. The most common approach is referred to as aperture photometry, whereby pixelized maps are numerically integrated out to some cutoff radius. This method can be imprecise, especially when coverage is poor. As an example, determination of constant map offset is complicated by non-negligible contributions from near sidelobes. Two related methods are commonly used: one involves extraction of the peak planet flux using an elliptical Gaussian fit (or similar) to the data; the other extracts the peak using a global fit of the beam Point Spread Function (PSF). With an estimate for the peak signal, the flux from the planet can be derived.

We describe planet flux reconstructed from *Planck* HFI observations using Gauss-Hermite estimates of peak brightness. Comparison with both aperture photometry and PSF fitting methods suggest overall good agreement. For validation, we show predictions from ESA models [242, 251]. We offer minimal interpretation of these results and refer the reader to [224] as well as the 2014 collection of *Planck* HFI papers. Weiland et al. present a detailed analysis of planet brightness as measured by *WMAP* [246]. Such analysis could be repeated with these data.

A detailed mathematical description of the method used to derive these results, including error analysis, is presented in Appendix D.

3.7.2 Mars

The Martian orbital period corresponds to little less than two Earth years. The planet rotates around its axis approximately once every 24.6 hours. The 25.2 deg axial tilt is comparable to Earth, but the relatively large eccentricity makes the southern hemisphere experience greater seasonal variations. The perceived brightness temperature is highly dependent on viewing location as the Martian surface is far from homogeneous. Diurnal heating and cooling of the atmosphere play a significant role. Finally, dynamical factors such as dust storms can affect the planet's albedo.

A number of models predicting Mars brightness temperature exist [253, 254]. We have primarily considered the models of Lellouch et al. [242] and Weiland et al. [246]. The Lellouch model has been used by the *Herschel* science collaboration. It incorporates surface and sub-surface temperatures taken from the European Mars General Circulation Model [255] but otherwise employs methodologies similar to [253, 254].¹⁵ The Weiland model is an alternative version of a model that was originally constructed by Edward Wright [241]. This updated version is used by the WMAP team and incorporates L_2 viewing angles and extends the spectral coverage down to WMAP frequencies.

Detectors on the HFI focal plane observe the same point on the sky within the span of a week. For that reason, accurate beam reconstruction can depend on predictions about rotational variations. Using the Lellouch model as a prior, we detect rotational variations in Mars brightness with high significance. We correct for this in analysis, scaling measured values to coincide with the time at which channel 100-1a observes the planet head on. For example, this rescaling changes the standard deviation in measured brightness temperature of the second Mars observation at 217 GHz from 3.1 to 2.0 K_{CMB} .

¹⁵According to Lellouch, the model is validated by spatially resolved temperature measurements of the surface.

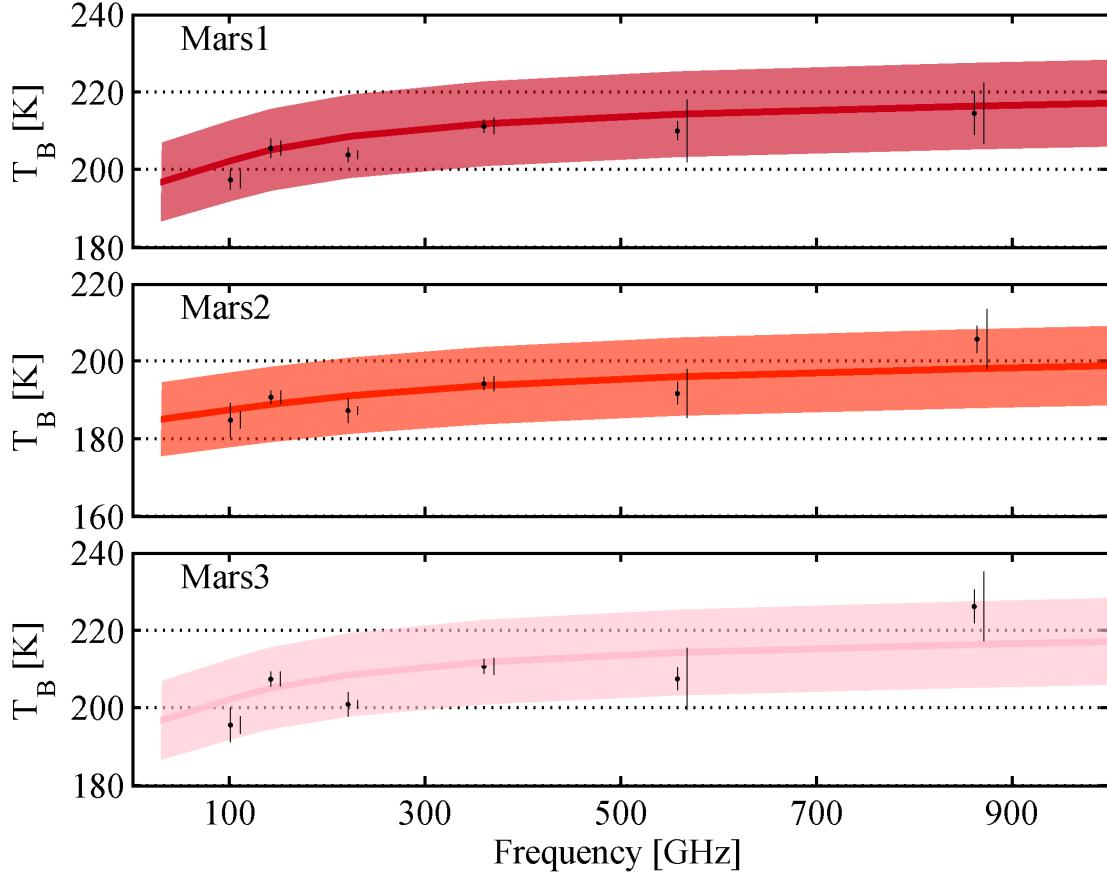


Figure 3.19: The measured brightness temperature for the three *Planck* HFI observations of Mars (see Table 3.2). Black points represent band averages and the vertical lines standard deviation in the ensemble. The vertical lines that have been shifted to the right of the points represent the Monte-Carlo derived error estimate (see Section D.2.1). The predictions of the Lellouch model (line), scaled by ζ_w , and the corresponding 5% errors (colored region) are also shown.

Figure 3.19 compares the measured brightness temperature of Mars with predictions by the Lellouch model. The model output has been convolved with spectral response of the *Planck* detectors for accurate comparison. All brightness temperature estimates are reported at the formal *Planck* HFI band centers, corresponding to 100, 143, 217, 353, 545, and 857 GHz. Error bars indicate the standard deviation of measurements within a band, not including Bessel's correction. The Monte-Carlo derived errors are shown for comparison (see Appendix D.2.1).

Because of seasonal variations, we compare different observations separately. Measure-

ments are consistent with model predictions and their stated 5% modeling error. However, we observe variations between frequencies that are common to all three observations. We do not observe such systematic behavior in simulations of the flux extraction pipeline using the Lellouch model as input. This could suggest a systematic intrinsic to the flux reconstruction.

We find that a global rescaling of the Lellouch model predictions by $\zeta_P = 0.973$ minimizes the residual between the model predictions and the three HFI measurements of Mars brightness temperature. A similar recalibration of model predictions was performed in *WMAP* analysis of the Weiland model [246]. Having access to both the Weiland and the Lellouch models, we performed the same recalibration analysis using *WMAP* results. This suggests a rescaling of the Wright model predictions of $\zeta_N = 0.953$, consistent with reports in [246]. Using the same method, we calibrate the Lellouch model predictions to *WMAP* results. This procedure suggest a $\zeta_W = 0.968$ rescaling of the Lellouch model minimizes the residual between model predictions and *WMAP* results. The ratio of scaling factors, $\zeta_P/\zeta_W = 1.005$, suggests that *WMAP* and *Planck* HFI are consistent in their absolute calibration at the percent level.

Finally, we note fluctuations in brightness temperature estimates which are common to all three Mars observations. This could point to a systematic in the flux reconstruction or a feature in the planet spectrum. We note that a similar jump at 143 GHz, is also observed in Jupiter and Saturn measurements (see Sections 3.7.3 and 3.7.4).

3.7.3 Jupiter

There exists a rich literature on the mm and sub-mm flux densities of the Jovian planets [256, 257]. Measurements in these wavelengths help constrain chemical abundances and atmospheric models. According to [258], broad emission from ammonia (NH_3) inversion lines is the main source of opacity for Jupiter and Saturn at short cm and mm wavelengths.

Jupiter's size and proximity make it the brightest planet on the sky as seen by *Planck*. Although some detectors are driven to saturation, the planet's brightness can still be estimated. Applying the non-linearity corrections described in Section 3.6.4, we derive the planet's flux

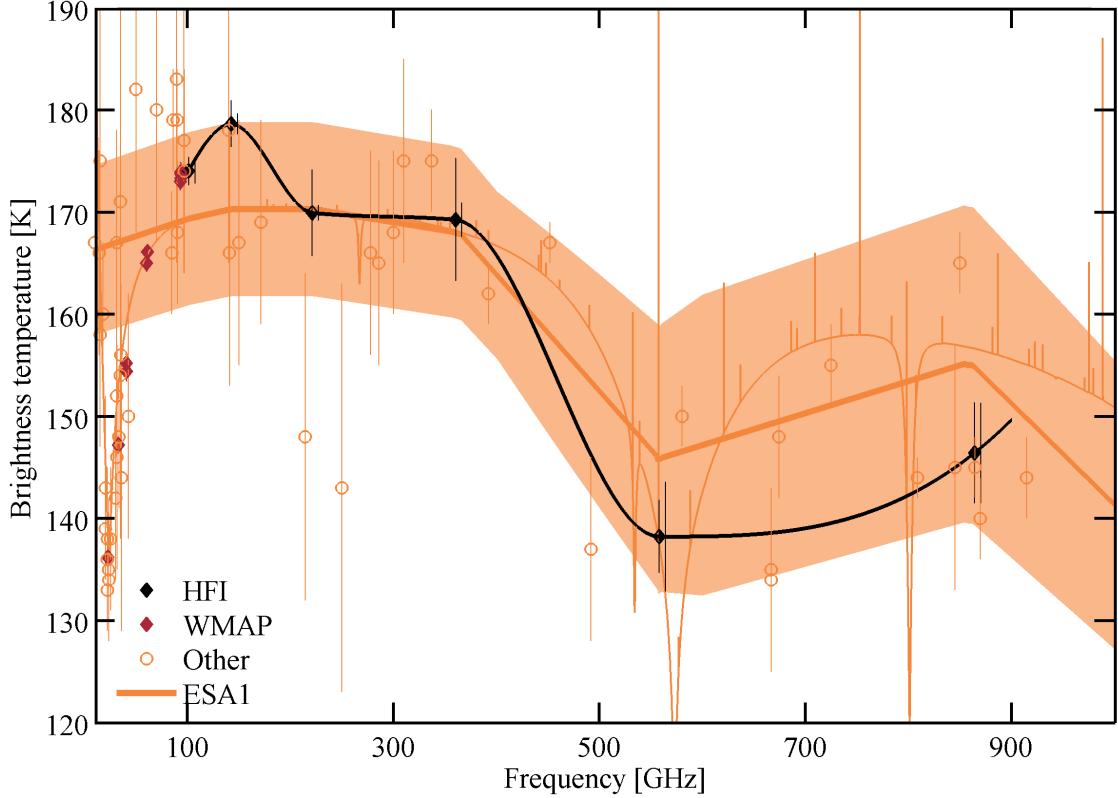


Figure 3.20: The measured brightness temperature for the five *Planck* HFI observations of Jupiter. Black points represent band averages and the vertical lines standard deviation in the ensemble. The vertical lines that have been shifted to the right of the points represents the Monte-Carlo derived error estimate (see Section D.2.1). The ESA1 model has been convolved with the *Planck* HFI spectral bandpass in a way that makes the model, as presented here by the orange region, only applicable at the HFI band center frequencies. The modeling error is inflated from 5% to 10% between at 353 and 545 GHz [259]. The raw output of the model is shown with a thin orange line. Note the ammonia line at 570 GHz.

at all frequencies. No simulations have been performed to characterize the fidelity with which the non-linearity correction can be determined. Because of this, the flux reconstruction error is inflated with increasing non-linearity correction. Figure 3.20 shows the flux at all frequencies for the 5 observations available. We do not account for any contamination in flux determination due to Jupiter's rings and moons.

As seen from L_2 , Jupiter is approximately 40 arcsec in diameter. In performing background subtraction of signal timelines, we account for planet occulting effects by nulling signal

estimates which fall inside the planet disk. We apply this correction in all planet flux reconstruction. It has a small, but non-negligible effect.

Jupiter is an ideal candidate to transfer *WMAP*'s dipole calibration to another instrument [232]. Assuming flux stationarity, we use our estimate for the observation averaged brightness temperature to compare with *WMAP* brightness estimates at 94 GHz. The seasonal averaged Jupiter brightness temperature at 94 GHz, as reported by *WMAP*, is $T_b^W = 174.6 \pm 0.9$ K. The HFI measurement at 100 GHz is $\tilde{T}_b^P = 174.0 \pm 1.4$, with the 1σ error estimate derived from the observed distribution. Using the ESA1 model to scale HFI predictions down to 94 GHz, we obtain $T_b^P = 173.8 \pm 1.4$. The ratio is $T_b^P/T_b^W = 0.995 \pm 0.010$.

Finally, we note the large variations in measured brightness temperature at 143–353 GHz. These are not seen in corresponding simulations. We believe these large signal variations are caused by increasing detector non-linearity and the onset of saturation.

3.7.4 Saturn

Saturn's flux determination is complicated by the presence of extended rings. This spectacular structure, primarily composed of ice [260], appears to have formed during a collision event [261]. The structure is composed of multiple annular objects which extend out from the planet's center. During the four *Planck* observations of the planet, Saturn's ring inclination angle spanned 3 to 13 deg as viewed from L_2 . Using the methodology presented in Weiland et al., one can fit a model which incorporates contributions from both the planet disk and the rings. Here, we do not present the two-component analysis, but we note that the data suggest $T_b^{\text{disk}} = 16 \pm 3$ K, which is consistent with *WMAP* analysis.

Figure 3.21 shows the observation averaged Saturn brightness temperature as a function of frequency, assuming a single component model, and compares them with the ESA2 model [259]. As we do not simulate variations due to the ring inclination angle, it is understandable that the Monte-Carlo derived error is smaller. We observe an interesting similarity in

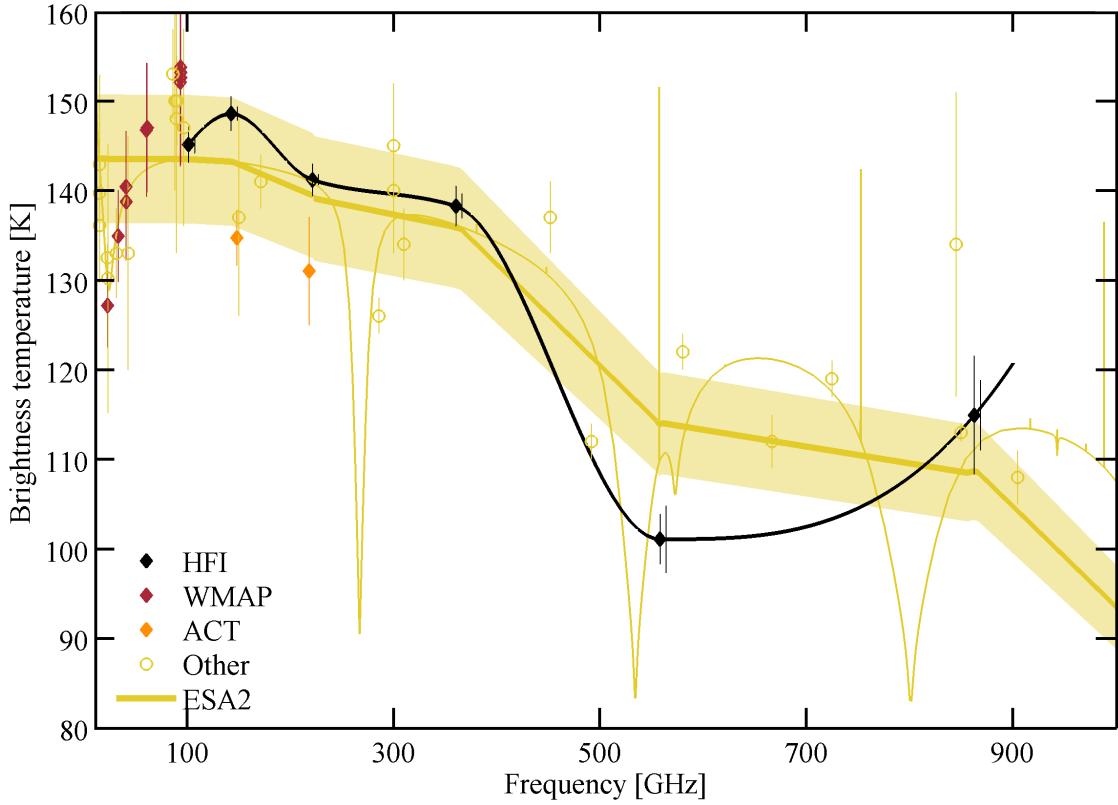


Figure 3.21: The average measured brightness temperature for the four *Planck* HFI observations of Saturn compared against the ESA2 model output after convolving with the *Planck* HFI spectral bandpass. The model does not account for contributions from Saturn’s rings. The three prominent absorption lines, including the one at 530 GHz, are phosphine (PH_3) absorption lines [256]. WMAP [246] and ACT [247] measurements are included for comparison.

our measurement of the Jupiter and Saturn brightness temperatures. A minor jump in temperature at 143 GHz and a significant dip at 545 GHz. We do not know if the 143 GHz jump is physical or if represents a systematic effect. Note that this feature is not observed in Uranus and Neptune (see Section 3.7.5). The dip near 545 GHz, observed in both Jupiter and Saturn, is understood, and is due to absorption features from both PH_3 and NH_3 . It is worth noting that the 545 GHz observations of Saturn are subject to significant color corrections (see Appendix D.3).

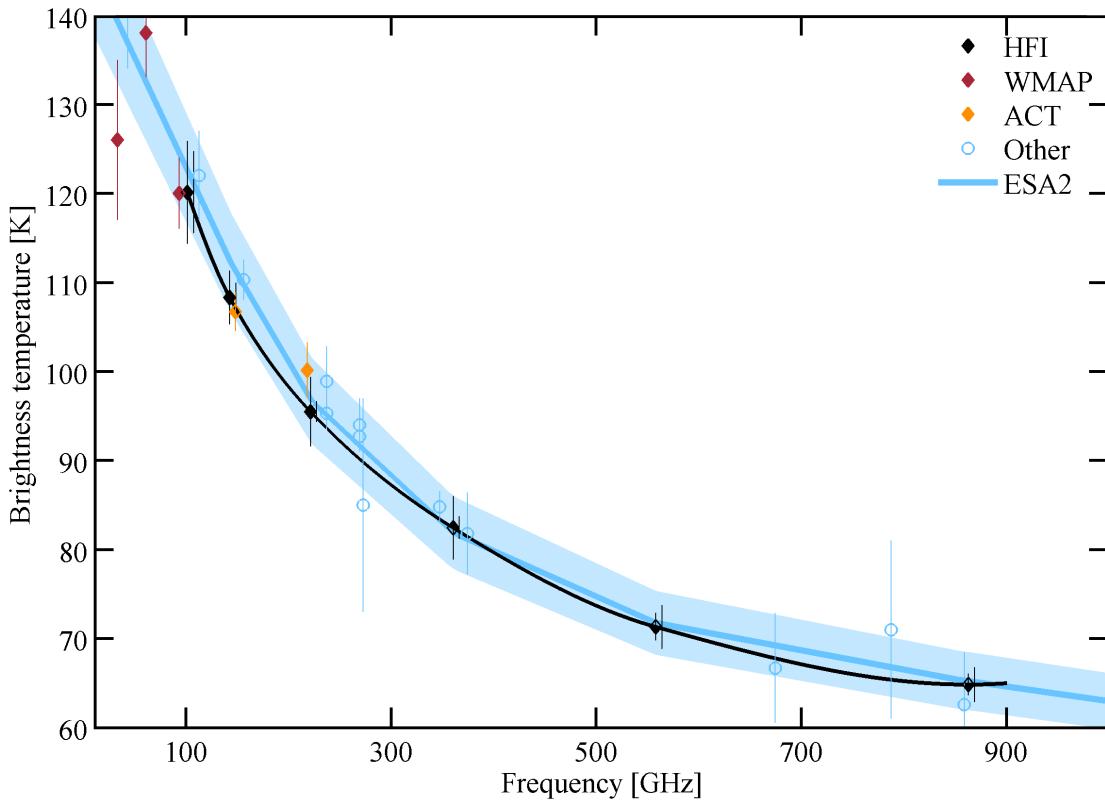


Figure 3.22: The measured brightness temperature for the four *Planck* HFI observations of Uranus compared to predictions of the ESA3 model after convolving with the spectral response of the *Planck* HFI. Other data shown include results from WMAP [246] and ACT [262].

3.7.5 Uranus and Neptune

The dimmest of the Jovian planets, Uranus and Neptune are often used as calibrators for CMB experiments probing relatively small angular scales. The near millimeter brightness temperature of Uranus and Neptune is discussed in [263]. For lack of a better way to determine absolute calibration at those frequencies, the sub-mm channels have been calibrated such that flux estimates from these two planets agree with predictions from ESA models. This calibration was performed by Guilaine Lagache et al. [224]. The method involves aperture photometry. For that reason, we do not expect identical results; those two algorithms will have different intrinsic error and bias.

The two planets are quite dim at 100 and 143 GHz (see Table 3.2). Unlike Mars, Jupiter,

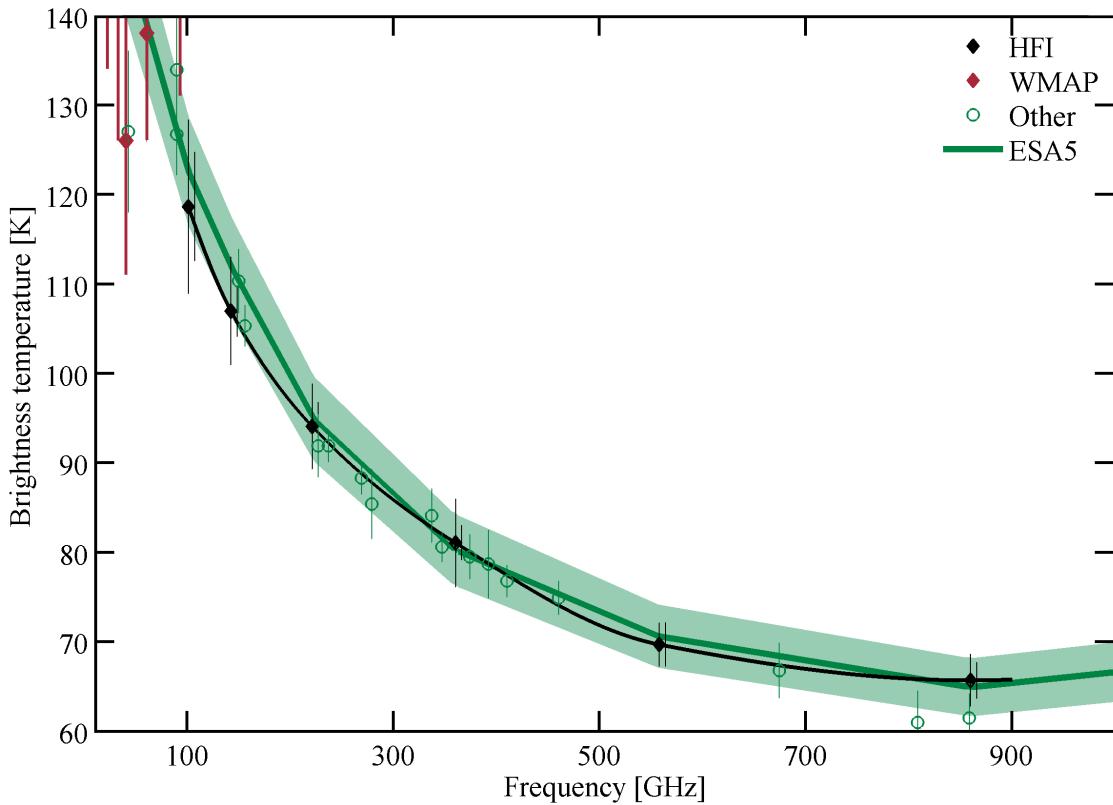


Figure 3.23: The measured brightness temperature for the four *Planck* HFI observations of Neptune compared to the convolved version of the ESA5 Neptune model. Other data shown include results from *WMAP* [246].

and Saturn, where instrument calibration dominates uncertainty, statistical error is significant in flux determination of Uranus and Neptune. Figures 3.22 and 3.23 show a comparison of the predicted brightness temperature of these planets with the most up to date ESA models at the time of writing [259], ESA2 for Uranus and ESA5 for Neptune. The models appear consistent with these measurements, even in frequency bands which are not directly calibrated to these observations.

A model of Uranus flux, attributed to Griffin and Orton [263], has recently been incorporated into the set of available ESA models [251]. This model is discussed in ACT analysis of Uranus flux [262]. The Griffin and Orton model of Uranus is referred to as ESA5 by scientists working on the calibration of SPIRE, an instrument on the *Herschel* satellite [251, 259].

Using Monte-Carlo derived errors, we see similar agreement between the *Planck* HFI measurements and the ESA2 and ESA5 models. The χ^2 per degree of freedom, as calculated using the six band averaged HFI measurements, is 1.4 and 1.5 for ESA2 and ESA5, respectively. However, the agreement is frequency dependent. Only considering 100–217 GHz, the χ^2 per degree of freedom becomes 2.7 and 0.8 for ESA2 and ESA5, respectively. Conversely, at 353–857 GHz, the χ^2 per degree of freedom becomes 0.1 and 2.1. The ESA5 model is therefore favored at low HFI frequencies and the ESA2 model at high frequencies. Note that the ESA2 model was used to calibrate the absolute gain of the sub-millimeter channels, it is therefore expected to agree quite well with these results at 545 and 857 GHz.

3.7.6 Summary

Figure 3.24 shows the difference between measured spectral flux density and that which is predicted by the models that are presented in Sections 3.7.2–3.7.5. Note that, unlike the brightness temperature, this ratio is linearly dependent on parameters that commonly affect the calibration, such as the beam solid angle and gain. The dashed horizontal lines represent the 5% model errors and the colored region is the standard deviation within the measurement derived from data. The flux jump at 143 GHz is common to Mars, Jupiter, and Saturn. Overall, the measurements are in good agreement with model predictions.

Uranus and Neptune are remarkably consistent with model predictions over all six *Planck* HFI frequency bands. The measured flux from Mars appears systematically lower than the Lellouch model by approximately three percent (see Section 3.7.2). This result is consistent with the rescaling determined in *WMAP* analysis of the Wright model for Mars flux presented in [241, 246]. We conclude that measurements of Mars flux density made by *WMAP* and *Planck* HFI are in excellent agreement. Comparison of Jupiter measurements suggest similar agreement.

Planck observed each of the outer planets approximately four times during a two-year

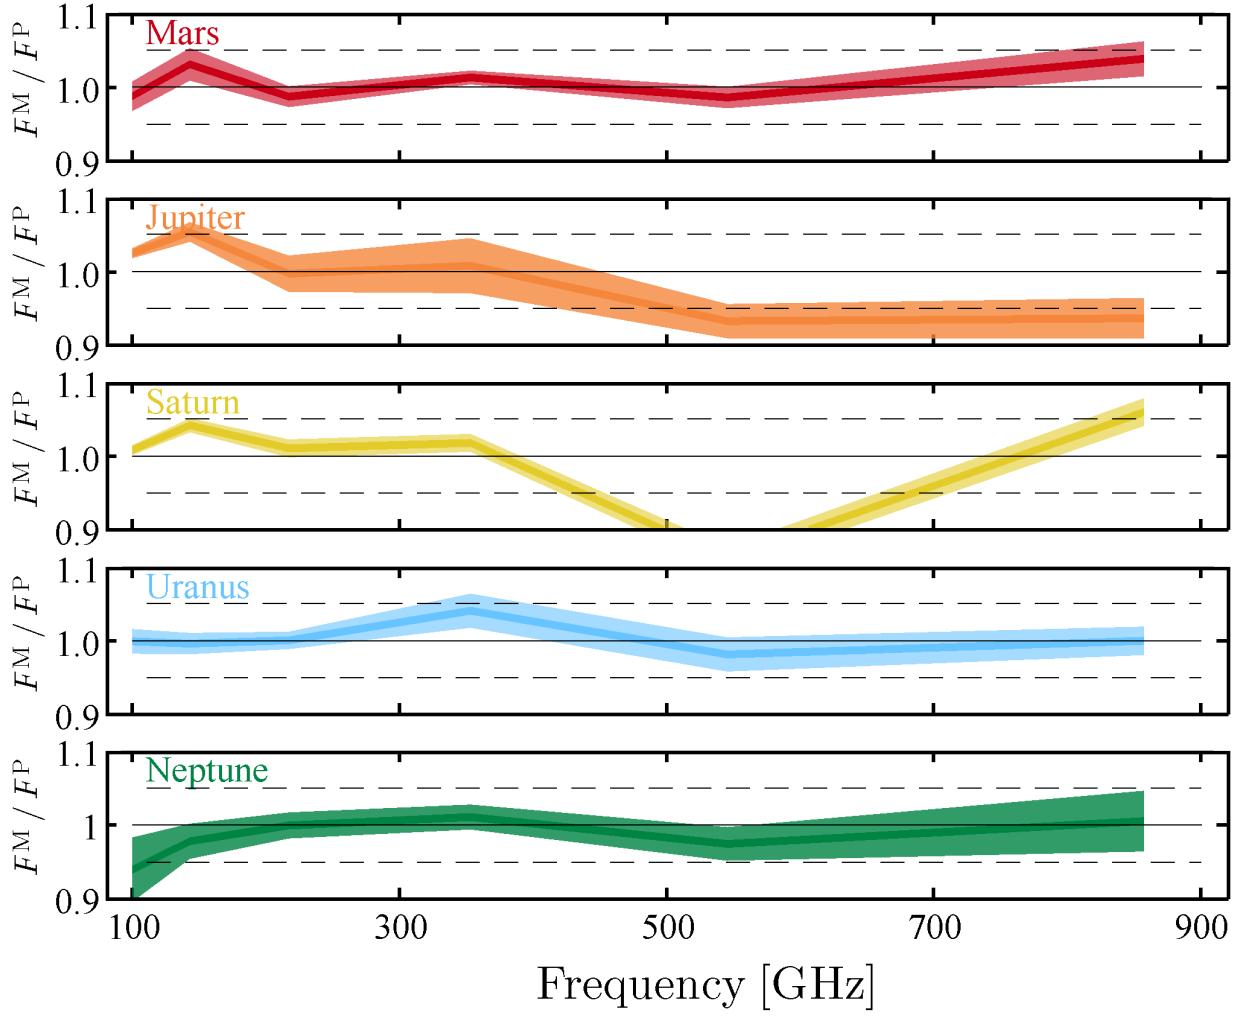


Figure 3.24: The average ratio of measured flux (numerator) and model predicted flux (denominator) as a function of frequency shown for all five outer planets. The model predicted Mars spectral flux density is scaled by $\zeta_P = 0.973$. Dashed horizontal lines represent 5% errors envelopes. Overall, models agree well with the *Planck* HFI measurements.

period. Although we have not examined this in detail, apart from Mars, there is limited, $\lesssim 2\sigma$, evidence for seasonal variation in planet flux.

3.7.7 Tabulated Results

The band averaged flux for each planet observation analyzed using this pipeline is presented in Tables D.2–D.6 of Appendix D.

Chapter 4

Calibrations and Simulations of the SPIDER Optical System

The SPIDER telescopes have been extensively calibrated. The detectors, described in Chapter 2, have been screened and characterized both at JPL Microdevices Laboratory and in Caltech using test cryostats large enough to cool down individual telescopes. Further characterization with telescopes mounted inside the flight cryostat has taken place both at Princeton University and the Columbia Scientific Ballooning Facility (CSBF) in Palestine, Texas. In this section I will describe the phenomenology of SPIDER’s far sidelobe response, discuss simulations that constrain *B*-mode contamination from polarized sidelobes, and compile band averaged main beam statistics.

4.1 Baffles and Far Sidelobes

Sidelobes are generated by unwanted radiation that propagates through the optical system all the way to the receivers, either through unlikely reflections, or low-level diffraction. Light that propagates through these sidelobes can produce a spurious signal. Sidelobe pickup is constrained through careful design of optical elements, including baffles.

The high packing density of the SPIDER baffles sets considerable constraints on the length

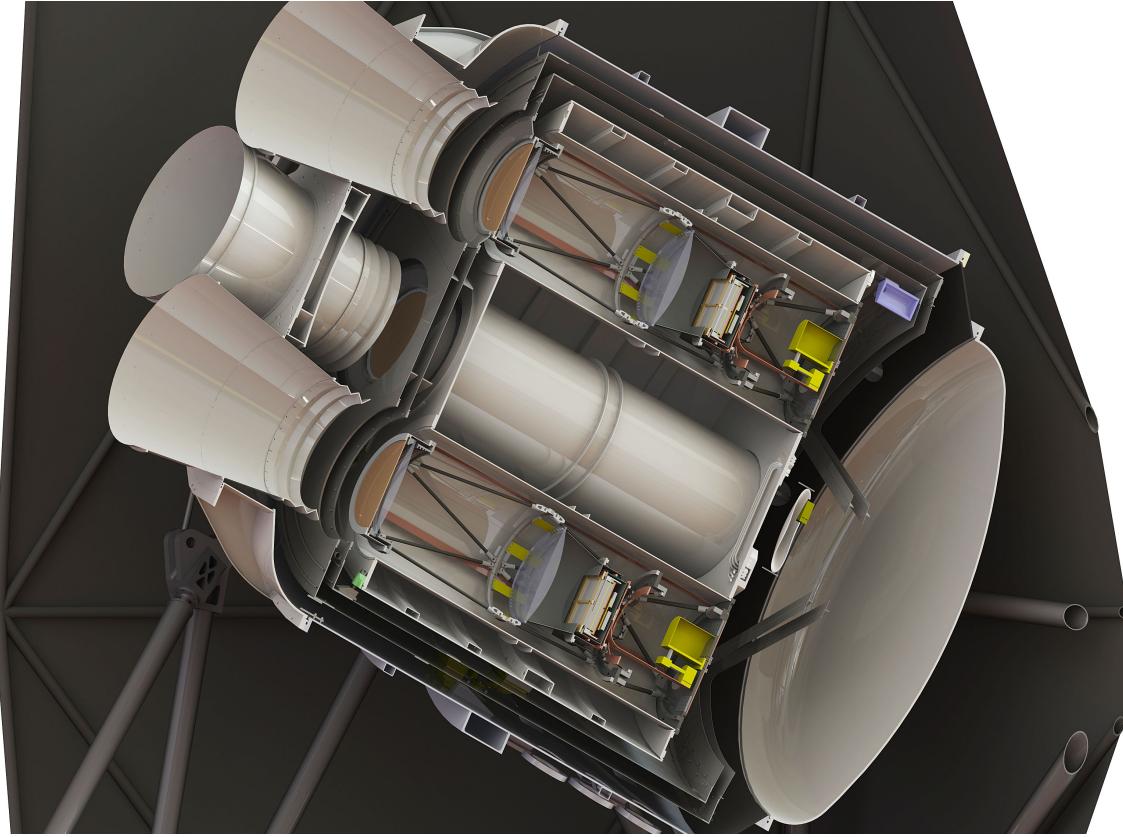


Figure 4.1: A section view of the payload displaying sunshields, gondola, baffles, and the internals of the cryostat. The window assembly, located just below the baffle accordion segment, is omitted from this figure. The half-wave plates and VCS filters are not shown. The payload is tilted at 40 deg elevation, corresponding to the upper region of the elevation range. The high packing density of baffles is evident from this figure.

and opening angles of the baffles. Figure 4.1 shows a section view of the payload, with the gondola sunshields in the background. An additional sunshield wing, not shown in Figure 4.1, is mounted on the port side of the sunshield frame to prevent illumination of the inside of the sunshields, as the boresight reaches the maximum extent of its azimuthal scan in the direction of the Sun, corresponding to 70 deg.

4.1.1 Baffle Design and Sidelobe Phenomenology

SPIDER’s reflective baffles were designed by Johanna Nagy at Case Western Reserve University. The baffles are conically shaped with an opening angle of 9.25 deg, with the base stretching

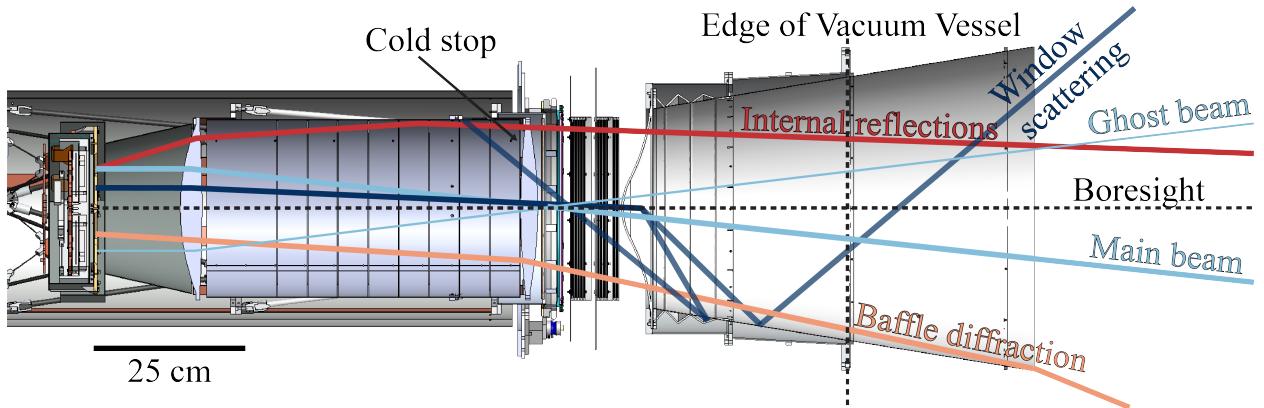


Figure 4.2: A cross section through the SPIDER telescope truncated just below the focal plane. Three rays represent different sources of sidelobes. Red represents reflections off the cooled and blackened optics sleeve. Somewhat surprisingly, it was found that the cooled optics sleeve is quite reflective in the mm-wavelength range. Blue represents window scattering; light is absorbed and then remitted at oblique angles with respect to the boresight. Orange represents diffraction off the edge of the baffle. This light is expected to be very polarized. The edge of the vacuum vessel is shown with a vertical dashed line. Note that the window deformation is exaggerated in this figure.

34 cm inside the cryostat where it meets the vacuum window and extending roughly 31 cm out from the top surface of the cryostat (see Figures 2.19 and 4.2). One of the baffle elements is shaped like an accordion in order to reduce large angle scattering at the window. Although suggested ray tracing simulations, testing has not yet shown conclusively that this improves the baffle performance.

Given the SPIDER optical system, sidelobe pickup is predominantly generated through three distinct mechanisms (see illustration in Figure 4.2). These are:

- Reflections on cooled optics sleeve – shown in red: Light enters the window at a large incident angle with respect to a detector's principal ray, then reflects on the cooled optics sleeve before generating a signal in the phased antenna array. Although the cooled optics sleeve is blackened with a mixture of Styccast and Stainless Steel powder, lab measurements suggest that the thin, $\lesssim 1$ mm, layer of this material is as much as 80% reflective at mm-wavelengths. Blackened baffle rings have been mounted internal to the optics sleeve to reduce its effective reflectivity by increasing the number of bounces

required to propagate to the focal plane. In addition, we have lined the inside of the sleeve with Eccosorb HR-10.

- Window scattering – shown in dark blue: In the time reversed sense, rays emitted from the focal plane are scattered by the 4.2 mm thick UHMW window to oblique angles with respect to the boresight, sometimes bouncing off the reflective baffle in the process. This effect can be reduced by making the baffle more absorptive, however, this would increase detector loading and therefore, noise amplitude. We have opted to make the baffle reflective at the risk of increasing susceptibility to large angle scattering. Lab measurements suggest that the UHMW window absorbs about 0.7% of in-band radiation.
- Diffraction on the baffle – shown in pink: SPIDER employs aggressive edge tapering, such that rays emitted from the focal plane cannot hit the baffle without first diffracting. Yet, lab measurements suggest approximately 5% of the beam terminates on a completely absorptive baffle. This is due to diffraction on the various optical elements in the telescope; note that ray tracing does not include diffraction. If the edge taper were reduced, eventually some rays emitted from the FPU could hit the baffle without diffracting first. Figure 4.3 compares measurements of the sidelobe response of a single 150 GHz detector to both the polarized BICEP1 profile [264] and the target for SPIDER’s science goal described in [124]. Our measurements still lack the signal to noise ratio required to constrain the sidelobe response at the level set by SPIDER’s science goal.

4.1.2 Spurious Polarization from Baffle Diffraction

At early stages in the development of the SPIDER experiment we generated an ensemble of simulations to understand how the Galaxy could produce a false B -mode signal through diffraction on the baffle edges (see pink ray in Figure 4.2). This work is discussed in more detail

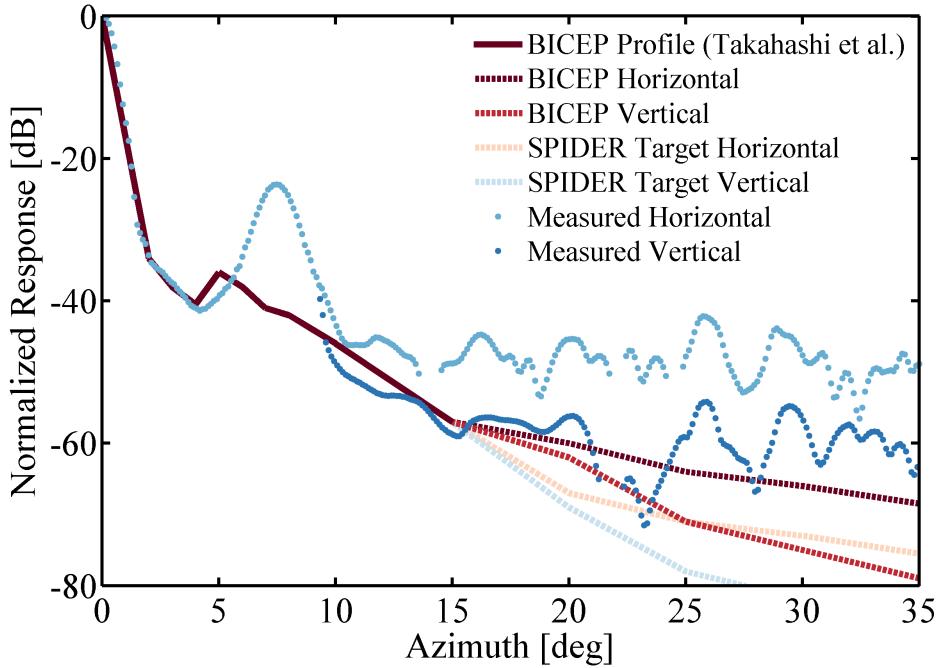


Figure 4.3: The sidelobe profile for a single 150 GHz detector when illuminated by a polarized source shown with blue and navy blue points. Note that the -23 dB peak at 7.5 deg corresponds to a ghost beam (see Section 4.2.2). The figure also shows the polarized BICEP1 profile (red and dark red) and SPIDER’s target profile (pink and light blue), described in [264, 124]. We do not currently have sufficient signal to noise ratio to constrain the SPIDER sidelobe down to the limit set by our science goal.

in Section 2.3 of [124]. In order to establish mission requirements for the SPIDER sidelobe response, a polarized sidelobe profile was constructed and convolved with a temperature map of Galactic emission at 150 GHz discussed in O’Dea et al. [130]. Figure 4.4 shows a temperature map convolved by a beam profile which captures the difference between the two polarizations of the BICEP1 beam shown in Figure 4.3.

Because of its complex morphology and proximity to SPIDER’s observing region, the Galaxy is expected to represent the most significant source of contamination on the scales of interest. Of secondary interests are illumination of the optics by the balloon and Earth’s limb. Such effects were not probed in these simulations.

The SPIDER main beam and near (2 – 12 deg) sidelobe profiles are informed by a physical-optics model of the ideal optical chain internal to the cryostat (see Section 4.2) which includes

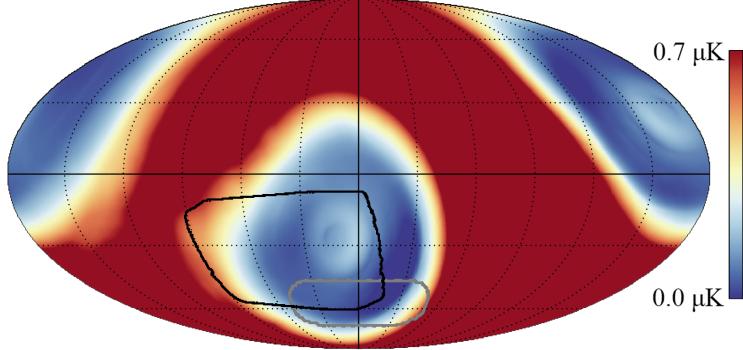


Figure 4.4: The boundaries of the SPIDER (black) and BICEP1 (grey) observation regions overlaid on a 150 GH temperature map smoothed with a BICEP1-like polarized sidelobe profile identical to the one shown in Figure 4.3. The Galactic plane is clearly bleeding into the sidelobe convolved map – see red regions which bleed into the outskirts of the observation regions.

both lenses and the half-wave plate. Accurate modeling of the SPIDER far (> 12 deg) sidelobes, however, requires a complete simulation of the optical system including sunshield, baffle, gondola, and balloon. This is a theoretically and computationally difficult task, which will likely never be performed. Instead, we adopt a model of the SPIDER far sidelobe profile as a power law out to 50 deg, where the beam is truncated. This model is consistent with measurements of BICEP1 polarized sidelobe response [264]. The B -mode power spectrum derived from the BICEP1-like sidelobe model is shown in Figure 4.5. Compared to the primordial $r = 0.03$ B -mode signal, it is over 12 times fainter at $\ell = 100$, while it is significantly brighter at large angular scales ($\ell \lesssim 25$). The BICEP1 baffle, which was designed to provide rejection at the level of $r \sim 0.1$ while observing at relatively high Galactic latitudes, does not provide sufficient rejection of the foreground signal for SPIDER’s large-scale measurements. Our simulations suggest that SPIDER requires an additional -7 – 10 dB of attenuation beyond 12 deg off axis in order to reduce the systematic contribution to a level near or below the $r = 0.03$ B -mode spectrum at all scales of interest. The contamination is then over an order of magnitude fainter than the cosmological signal at $\ell \sim 30$, and entirely negligible at the $\ell \sim 80$ peak.

Although SPIDER beam sidelobe measurements still lack the signal to noise ratio needed

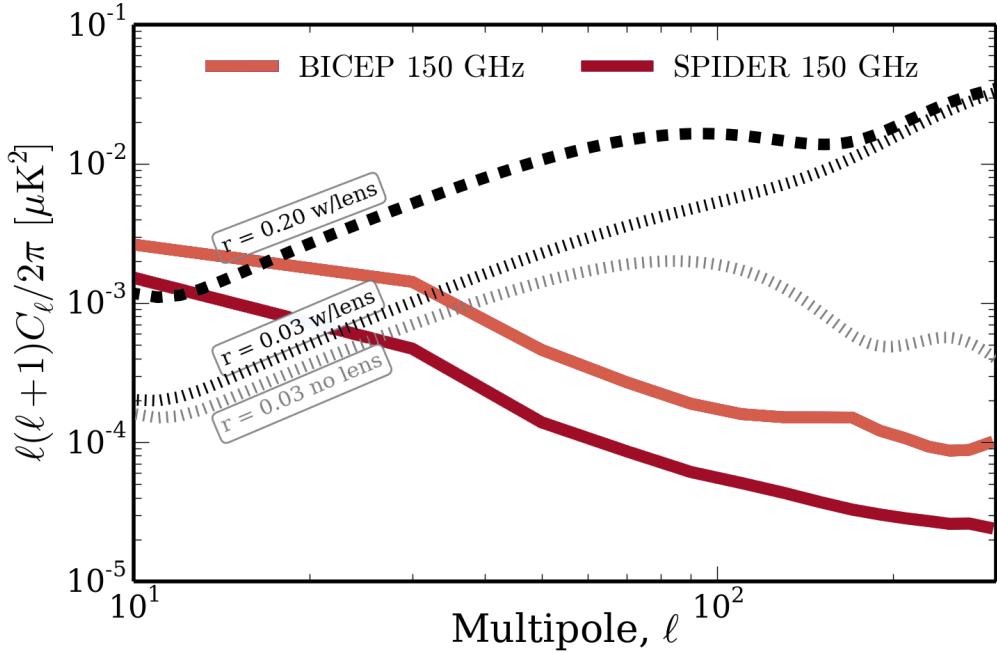


Figure 4.5: The B-mode power spectra generated by Galactic pickup from BICEP1-like polarized sidelobes shown in Figure 4.3. We caution against over interpretation of the difference in power observed in the SPIDER and BICEP1 regions. The two spectra agree to within a factor of 2–3 within the entire multipole range shown.

to demonstrate that its baffle will provide the required factor of ten improvement over the BICEP1 measurement, it appears to be well within reach. For example, geometrical theory of diffraction calculations indicate that the addition of another low-angle diffraction edge to the baseline baffle provides the extra attenuation required [265]. For this reason, we are considering adding a tapered angle on the edge of the baffles to reduce the polarization fraction of this type of sidelobe pickup [265].

4.2 Beam Maps

All six SPIDER telescopes have been characterized at Caltech. Some fraction of all focal planes was also probed during integration in Texas. At Caltech, the beam maps were acquired using a custom-made beam mapping turret that accommodates the Caltech test cryostat. A hot source, modulated at around 14 Hz, is placed 29 m away from the telescope window. Defining

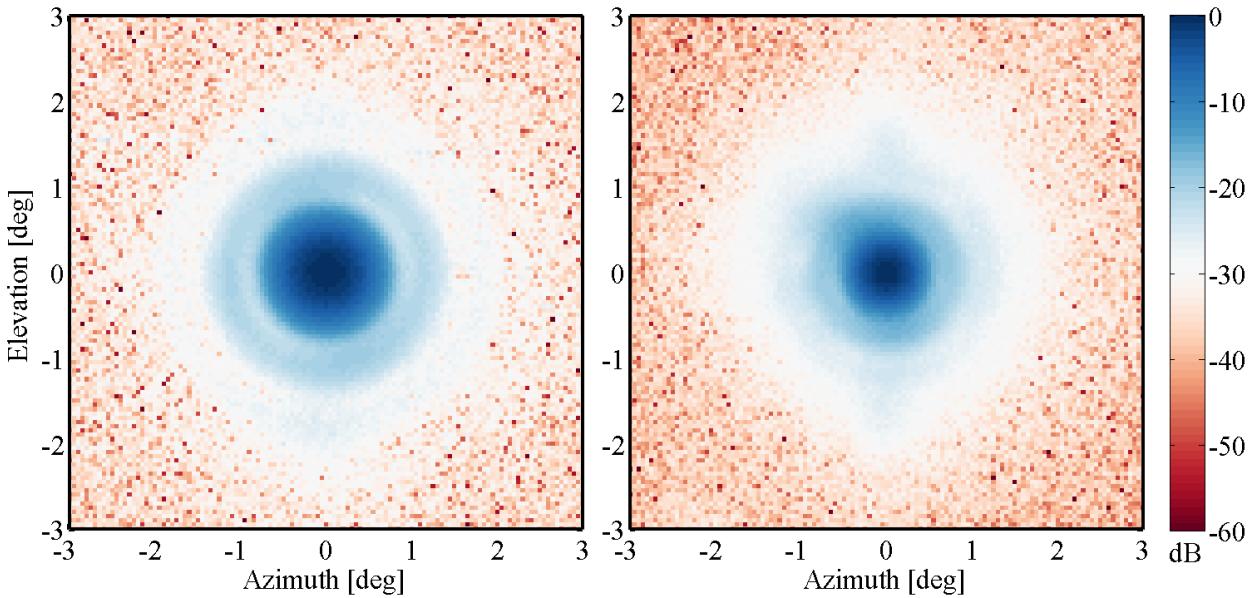


Figure 4.6: Beam maps of X6 (left) and X3 (right) obtained from stacking hundreds of individual detector beam maps obtained at Caltech. A couple of sinc-like rings can be seen in both frequencies along with detector cross talk ghosts above and below the main beam, corresponding to the detector plate scale, approximately 1.6 deg. Note, that all detectors scan the hot-source under approximately the same angle during the beam map. Beam symmetrization from cross-linking, sky-rotation, and rotation of half-wave plates will improve the symmetry of the effective beam.

the far field as D^2/λ , which corresponds to 36 and 54 m for 94 and 150 GHz respectively, it is clear that these beam maps are *almost* in the far-field. The turret then scans the room, effectively mapping a 14×14 deg patch. This process results in beam maps with 25–30 dB signal to noise ratio at 150 GHz using a pixel width of 5 arcmin. By combining hundreds of individual detector beam maps, features that are common to all detectors are revealed at the -35–40 dB level. Band averaged beam window functions will be compared to the effective window function derived from calibrating the SPIDER power spectra on degree scale temperature anisotropies as measured by *WMAP* and *Planck* HFI. Ghosting from internal reflections and fiducial beams due to inductive detector crosstalk in the SQUID multiplexing system are of some concern. The analysis presented here will help quantify effects from such non-idealities.

SPIDER’s optical elements are expected to produce a sinc-like beam that is truncated on an aggressive edge taper which is cooled to 2 K (see Figure 4.2).¹ The sinc-function can be adequately described by a symmetric Gaussian inside the region bounded by its first minimum. To see that, note that equating the Taylor expansions of $\text{sinc}(\rho x) \equiv \sin(\rho x)/\rho x$ and $\exp(-x^2/2\sigma^2)$, where we keep the first two terms for both expansions and assume that $x \ll 1$, gives $\sigma = \sqrt{3}\rho$. Although nothing is preventing us from using an asymmetric sinc-function to describe SPIDER’s main beams we have chosen to use the asymmetric Gaussian function instead, also known as an elliptical Gaussian function. See Appendix C for further information about the mathematical properties of such functions.

4.2.1 Stacking Beams

An extensive discussion on SPIDER main beam shapes, non-idealities, and impact on cosmological analysis can be found in [197]. Here, we briefly discuss the band averaged beam shape and our ability to probe the near-sidelobes using ground based measurements. The data used for this analysis was generated by the Caltech group.²

A simple stacking algorithm combines individual beam maps using centroids derived by fitting elliptical Gaussian functions to the main beam. The stacked and re-normalized beam map, $\Psi(\theta, \phi)$ can be expressed as

$$\Psi(\theta, \phi) = \frac{1}{N} \sum_{i=0}^N \psi_i(\theta, \phi), \quad (4.1)$$

where $\psi_i(\theta, \phi)$ are the individual beam maps, translated such that the beam centroid is at the origin. This allows us to probe features that are common to all detectors within a given telescope and effectively push through the noise floor of an individual detector beam map. Using this method, features that are not spatially coherent among all detectors, such as ghosts and crosstalk effects, are integrated down. Figure 4.6 shows a beam map acquired

¹The sinc-function corresponds to the Fourier transform of a square aperture [266].

²Many thanks to Rebecca S. Tucker.

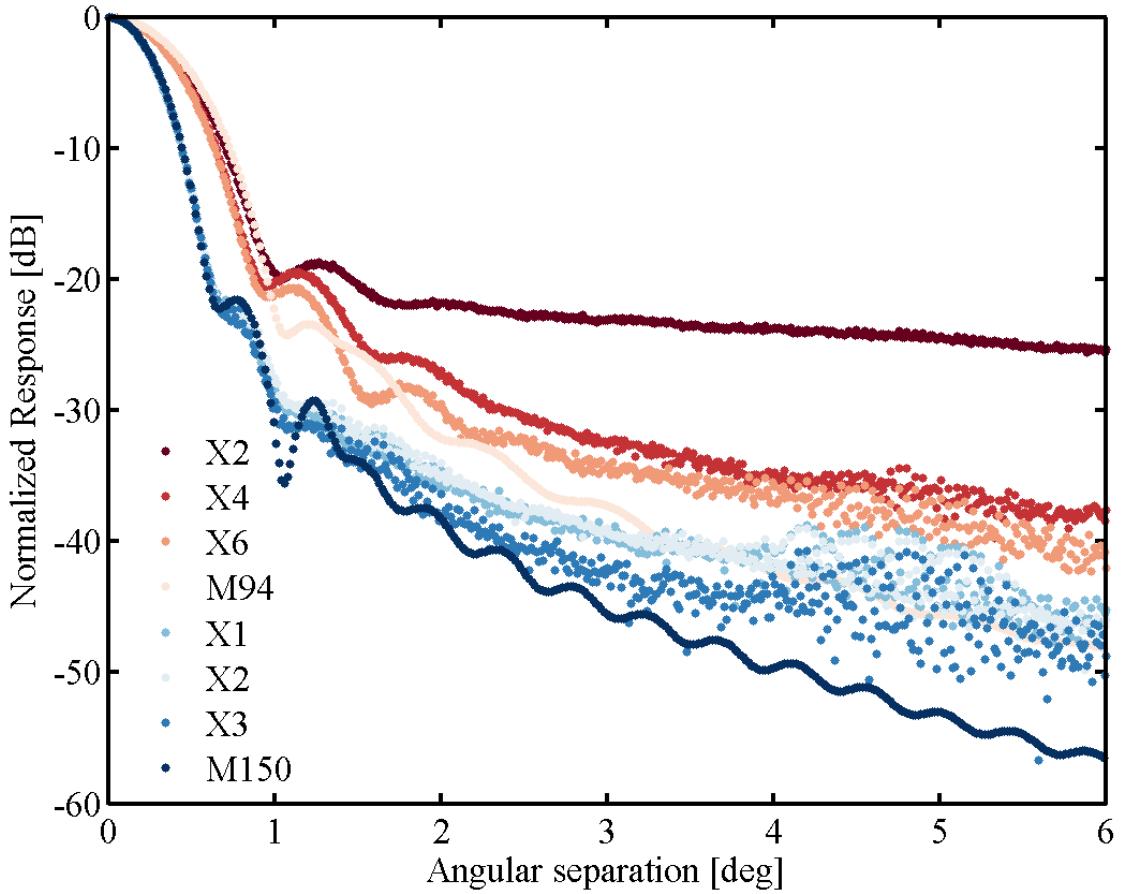


Figure 4.7: The azimuthally averaged beam profiles for the stacked beam product of all six telescopes, X1–X6. Also shown are the results of physical-optics simulations performed by Marc Runyan. With the exception of X2, the noise floor is hit around -35 – 45 dB.

by applying the stacking method on beams from the X6 and X3 telescopes. We observe smeared sinc-like features, which are in qualitative agreement with results from ray tracing and physical-optics simulations. We expect some smearing of the sinc-function, not only because of finite bandwidth, but also because of uncertainties in determination of centroids. Crosstalk beams can be observed as well, approximately 1.6 deg away from the beam center, along a direction corresponding to a detector mux column. These are due to inductive coupling in the time-domain multiplexing system which causes correlation between adjacent detectors. Figure 4.7 shows the band averaged beam profiles for all six SPIDER telescopes. We have also included predictions from physical-optics modeling, shown as M94 and M150 for 94 and 150 GHz, respectively. Table 4.1 summarizes some beam properties obtained by

Table 4.1: Telescope averaged beam properties as well as predictions from physical-optics modeling, referred to as M94 and M150 for 94 and 150 GHz, respectively. Here, θ_{FWHM} refers to the beam full width at half maximum, while Ω_{FWHM} is the beam solid angle of a Gaussian beam with the corresponding beam width. We can also estimate the beam solid angle by numerically integrating the stacked beam maps, denoted Ω , which represents the beam solid angle derived by integrating the beam out to $4 \times \theta_{\text{FWHM}}$. We chose to exclude X2 in that calculation due to low signal to noise ratio. We expect significant discrepancies between Ω_{FWHM} and Ω as the beam is not fully captured by an elliptical Gaussian function. Note that $\theta_{\text{FWHM}} = \sqrt{8 \ln(2) \times \sigma_x \sigma_y}$, where σ_x and σ_y are the two beam width parameters of an elliptical Gaussian model. Finally, ϵ represents the beam ellipticity defined as the ratio between the major and minor axis of the elliptical Gaussian fit. Reported errors correspond to one standard deviation in an ensemble drawn from individual beam map properties.

	θ_{FWHM} [arcmin]	Ω_{FWHM} [deg 2]	Ω [deg 2]	ϵ
X2	47.3 ± 2.0	0.704 ± 0.062	N/A	1.008 ± 0.006
X4	43.3 ± 0.7	0.590 ± 0.020	0.662 ± 0.023	1.032 ± 0.011
X6	42.6 ± 0.7	0.571 ± 0.018	0.611 ± 0.020	1.017 ± 0.009
M94	48.7	0.747	0.776	1.002
X1	30.1 ± 1.3	0.285 ± 0.025	0.306 ± 0.027	1.026 ± 0.016
X3	29.3 ± 0.5	0.270 ± 0.010	0.288 ± 0.010	1.014 ± 0.014
X5	29.7 ± 0.5	0.277 ± 0.009	0.289 ± 0.009	1.002 ± 0.009
M150	29.6	0.275	0.289	1.007

stacking beam maps within a given telescope, including errors derived from single detector beams within that ensemble. It is instructive to compare this table with the average beam properties derived from individual beam maps [197]. Overall, the agreement between the two methods is quite good. Note that X2 appears to have a significantly larger θ_{FWHM} than the other two 94 GHz telescopes. This is most likely due to noise bias in the beam map. The noise floor is very clear in Figure 4.7. Beam maps produced so far allow us to push the band averaged sidelobes down to -60 dB for a subset of detectors (see Figure 4.3). This is not sufficient to rule out the presence of systematic inducing polarized sidelobes caused by scattering on our baffle (see Section 4.1.2). It is, however, promising to see that the

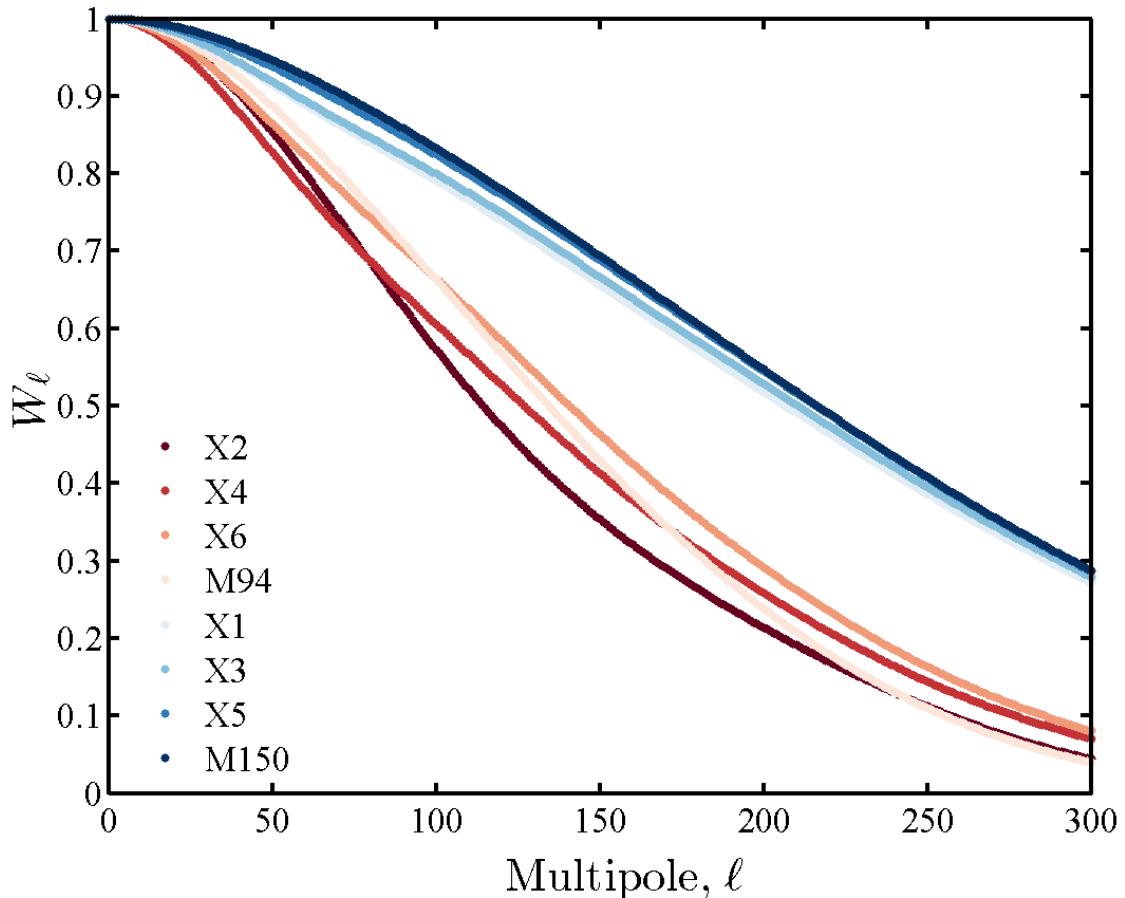


Figure 4.8: The squared Legendre transformation of the beam profiles shown in Figure 4.7. Note that SPIDER’s sensitivity to any angular power spectrum is below 10% and 30% for $\ell = 300$ at 94 and 150 GHz respectively.

sidelobe power continues to fall over the entire region probed. Beam maps performed either in Princeton or in the field will hopefully allow us to further constrain sidelobe amplitudes.

It is interesting to propagate the beams into multipole space. We do this to first order by taking the Legendre transform of the azimuthally symmetric beam profiles. If the azimuthally averaged beam profile is expressed by the function $b(\theta)$, the corresponding Legendre transformation becomes

$$b_\ell = 2\pi \int_1^{-1} P_\ell^0(\cos \theta) b(\theta) d(\cos \theta), \quad (4.2)$$

where P_ℓ^0 represents the associated Legendre polynomial of order zero. The corresponding

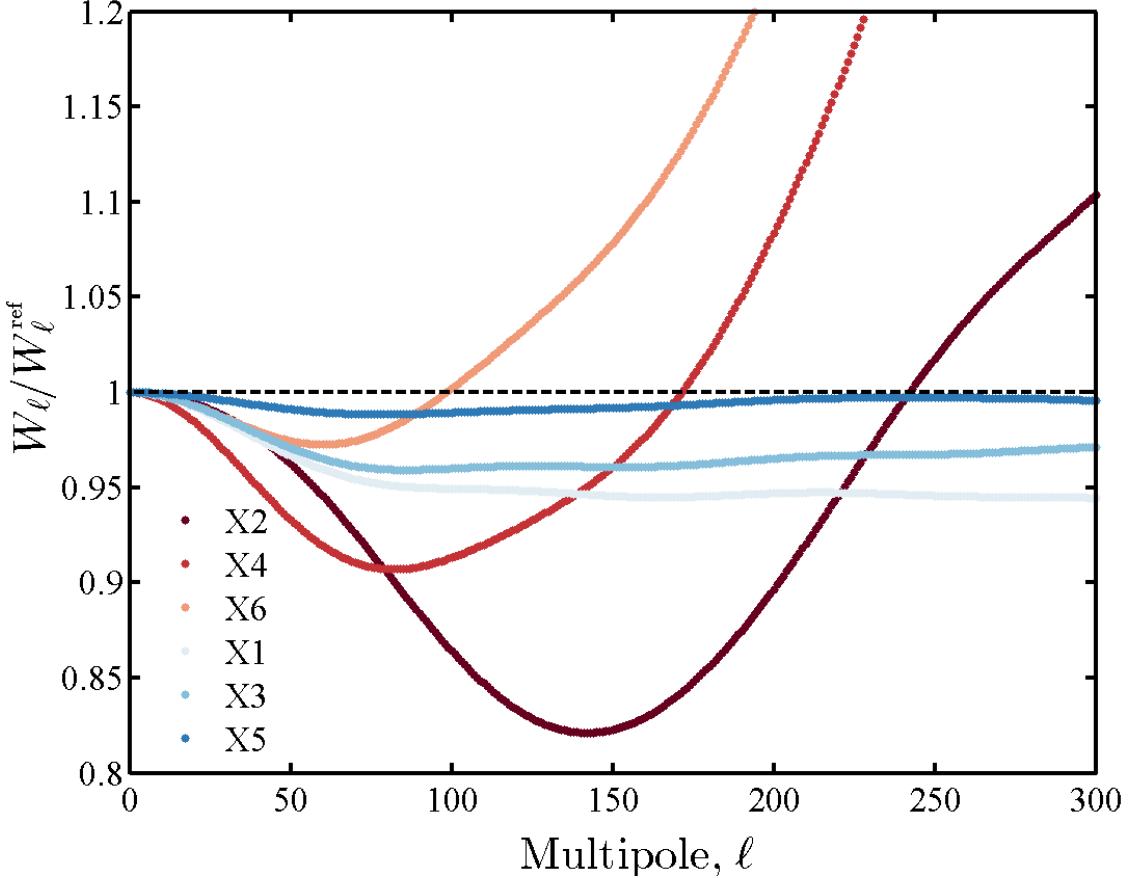


Figure 4.9: The ratio of stacked beam window functions compared to the ones derived from physical-optics simulations, W_ℓ^{ref} . A line falling on unity represents perfect agreement with predictions from simulations.

window function is then proportional to the square of b_ℓ . Neglecting any normalization factor, we write $W_\ell = b_\ell^2$.

Calibrations of large sky surveys such as *WMAP* or *Planck* HFI are based on the CMB dipole signal. Figure 4.8 shows the window function corresponding to the beam profiles shown in Figure 4.7 while Figure 4.9 shows the corresponding ratio, with predictions from physical-optics simulations in the denominator. The agreement with physical-optics modeling is surprisingly good at 150 GHz, with the three 150 GHz telescopes consistent at the 2% level out to $\ell = 300$. At 94 GHz, physical-optics modeling does not agree that well with realized window functions. We also note that there is significant spread among the three telescopes. Part of this is likely due to inadequate signal to noise ratio in the beam maps,

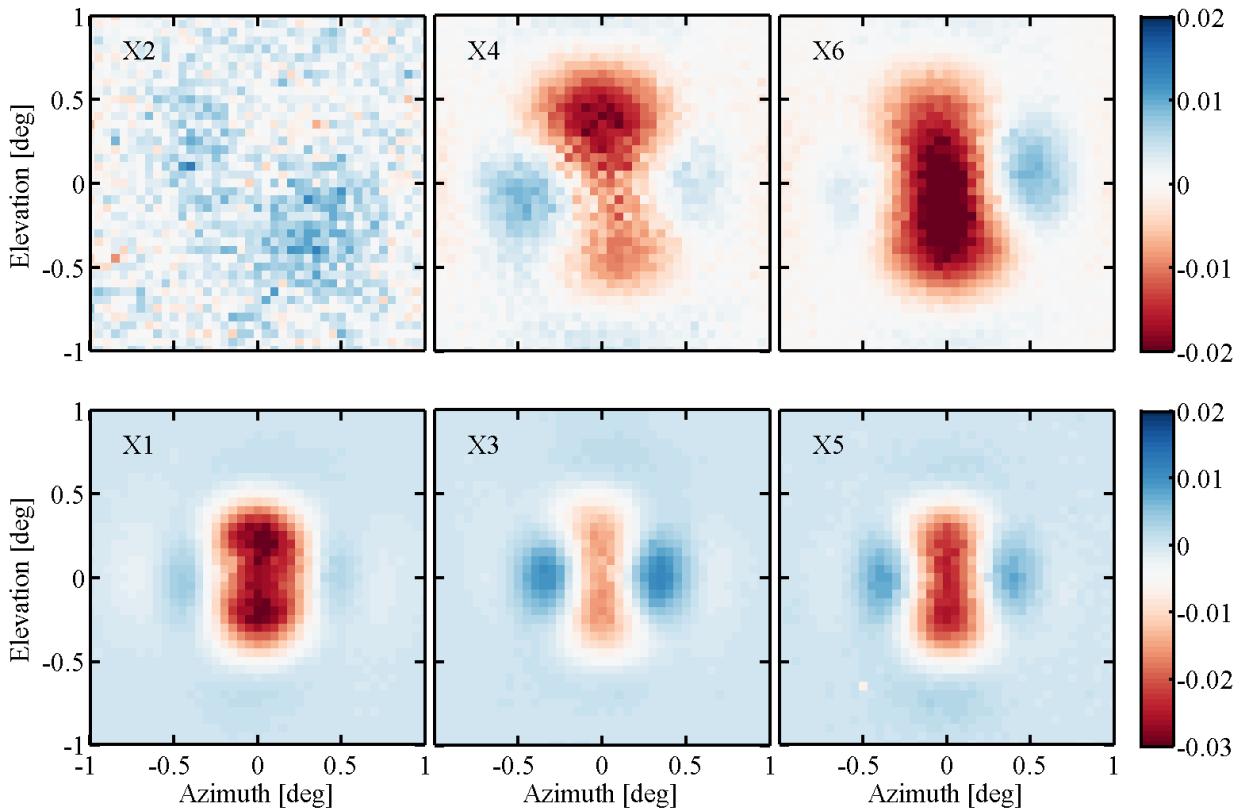


Figure 4.10: Telescope averaged difference beams obtained by stacking hundreds of individual beam difference maps. Shown here on a linear scale, normalized relative to peak. All but X2 show similar features. Note the slightly different color scales used for the two frequencies.

especially for X2. We have been unable to compare the dimensions of optical components used in the physical-optics simulations to the realized values. SPIDER will calibrate on degree scale temperature anisotropies using data from *Planck* HFI. Such a procedure zeros the beam error (on average) over this range, but it will also generate an effective beam window function which will be compared to these measured beam window functions. Any discrepancy between the two results will have to be understood.

Using this stacking algorithm we can also look at A/B polarization difference maps. Figure 4.10 shows these for all telescopes as derived using beam maps that were acquired at Caltech. In this case, the stacking algorithm translates both A and B polarization by the same amount so as not to introduce spurious alignment issues. We have yet to determine if

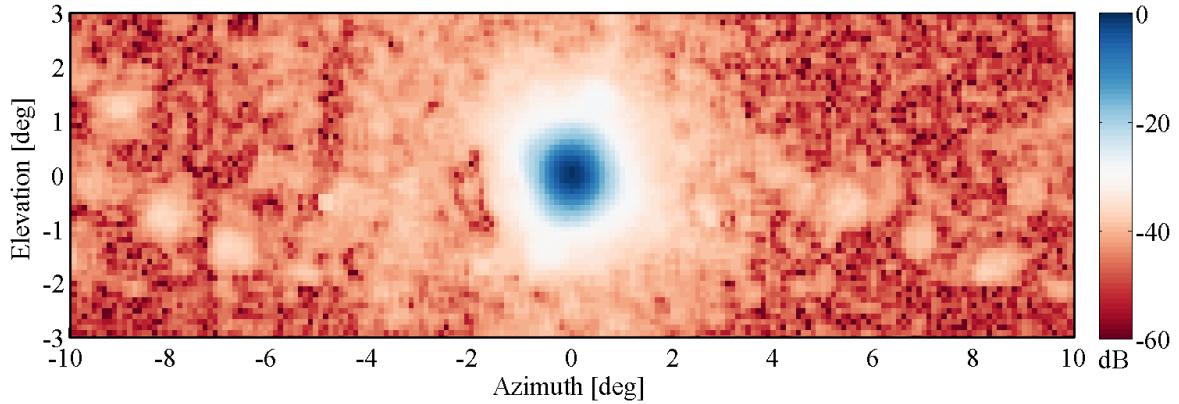


Figure 4.11: An inverse variance weighted stacked beam map from about a hundred detectors on the X3 focal plane. Ghosts are visible as rows of diagonal blobs on either side of the main beam. This beam map indicates the presence of cross talk beams, located diagonally about 2 deg from the main beam.

the lack of structure in the X2 difference maps is due to inadequate signal to noise ratio or simply lack of structure in the A/B difference maps.

4.2.2 Optical Ghosting

Optical ghosting is caused by internal reflections whereby a fraction of the outward propagating beam is cast down towards the focal plane before eventually propagating back out through the telescope window. Ghosting is an unavoidable feature of simple refracting telescopes.³ Optical components have Anti Reflection (AR) coating to minimize these effects, but some residual reflections are unavoidable, especially for an instrument with a 25% bandpass. An example of a ghost beam is depicted in Figure 4.2.

In the SPIDER optical system, the primary cause of ghosting is due to the half-wave plate (see Section 2.7). We have observed that the ghosting amplitude is dependent on the level of delamination on the HWP's anti-reflection coat. Changes to the AR coating procedure are thought to have reduced the level of ghosting down to approximately 1% of the main beam amplitude [184, 197]. Ghosts are observed even in cryogenic runs without the half-wave plate installed. These are expected to be due to other lower-amplitude internal reflections,

³In photography, this effect is commonly referred to as a lens flare.

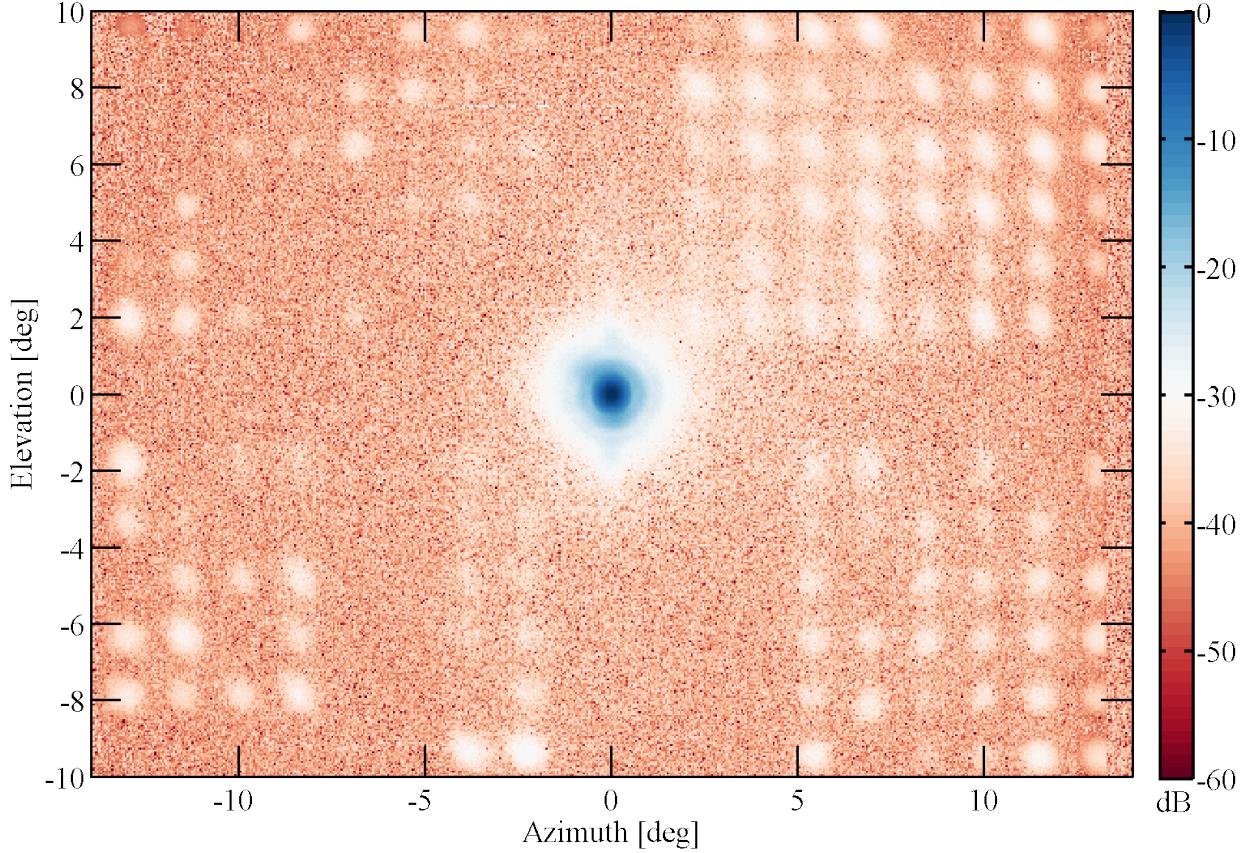


Figure 4.12: A stacked beam map from an old incarnation of X3, showing ghosts of different detectors compared to their main beam which has been centered on this map. The amplitude of the main beam is representative of the band averaged beam while the amplitude of the ghosts are suppressed as a result of the beam stacking. Note that the ghost morphology is not constant and the ghosts seem to get more extended with radial distance from the FPU. The HWP AR coat has been improved since these beam maps were acquired.

though we have been unable to identify a single optical element that is responsible for these reflections.

Based on simple geometrical arguments we expect that the ghost location will depend on the location of the detector on the focal plane. A hypothetical pixel at the center of the focal plane will have its ghost reflected directly back into the center of the beam while the ghost of a corner pixel will be reflected into the diagonal opposite of the focal plane [184]. The coordinates of the ghost center, (θ_g, ϕ_g) , relative to the main beam center can be described

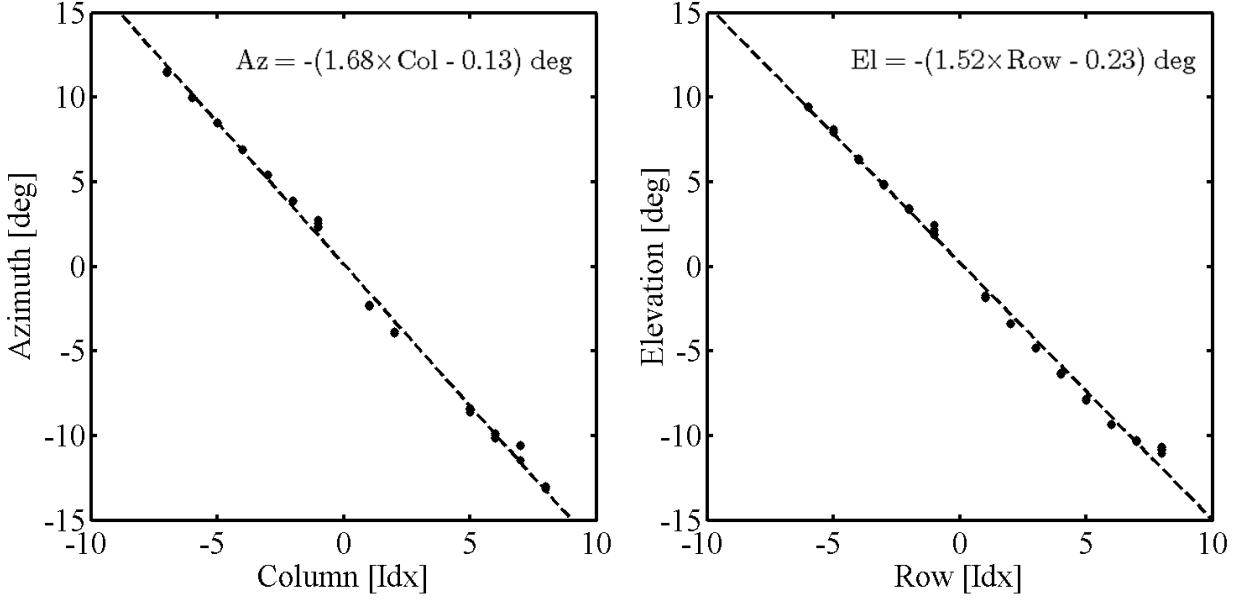


Figure 4.13: The location of X3 ghosts, as shown in Figure 4.12, in both elevation and azimuth as a function of location on the focal plane along with the best fit line, providing constraints on h_θ and h_ϕ .

according to

$$\theta_g = h_\theta \cos(\psi)c - h_\phi \sin(\psi)r, \quad (4.3)$$

$$\phi_g = h_\phi \cos(\psi)r + h_\theta \sin(\psi)c, \quad (4.4)$$

where h_θ and h_ϕ represent the linear dependence of ghost location on focal plane coordinates, ψ is the rotation angle of the telescope focal plane around the boresight of each telescope (roll), and r and c represent the physical location of the pixels on the focal plane. One can think of r and c as the physical row and column index.

Figures 4.11 and 4.12 show a stacked beam map where the distribution of ghosts around the main beam is clearly visible. Some ghosts are missing due to a combination of yield and data processing effects. Since this is a stacked map, any feature that is not common to all beams will be suppressed in amplitude. From a sub-set of the individual detector beam maps, we can determine the location of the ghosts and compare with the location of the detector

on the focal plane. Figure 4.13 shows the ghost location in both elevation and azimuth as a function of detector location on the focal plane, see r and c in Equations 4.3 and 4.4. The location of ghosts was obtained by extraction of local maximum after filtering the beam map with an 42 arcmin wide pillbox kernel. The result indicates a very repeatable pattern. According to these results, a shift by a single pixel along either a row or a column corresponds to approximately 1.6 deg shift in the ghost location on the sky. The further the pixel is away from the center of the focal plane, the larger the separation between the ghost and the main beam.

Beam measurements, performed both at Caltech and in Palestine, have allowed us to characterize the ghost amplitude, shape, and displacement as a function of position on the focal plane. With the data in hand, we observe some correlation between both ghost shape and amplitude with location on the focal plane. Although ideally, more data would help quantify the ghost characteristics, this analysis suggests we can perform first order correction to the ghosting effect in data processing, thereby reducing a potential systematic effect. Since ghost amplitude, shape, and location can be affected by delamination, a seemingly stochastic process, ghost characterization should be performed during a cryogenic run that coincides with flight.

Recent analysis of ghost shapes and amplitudes suggests that the ghosts peak at approximately -20 dB compared to the main beam. Simulations probing how B -mode systematics are induced through ghosts are discussed in MacTavish et al. [267], albeit for a different flight than currently proposed. That work suggests that a 10% ghost will not create spurious polarization exceeding a primordial B -mode signal generated by $r = 0.01$. More recent, unpublished simulations, assuming an Antarctic flight, give similar results. Those simulations assume no correction for ghost amplitude during data analysis. As first order corrections will be implemented in data processing, we can feel confident that SPIDER’s ghosts will not hinder cosmological analysis.

Chapter 5

Conclusions

In this thesis, I described the design and characterization of the SPIDER experiment which will launch from Antarctica in late 2014. The experiment is designed to advance our understanding of the B -mode power spectrum.

This experiment, poised to have the greatest instantaneous sensitivity of any CMB polarimeter so far, is currently shipping to Antarctica in preparation for a launch this austral summer [203]. I also described calibration work required for high fidelity cosmological analysis of the *Planck* HFI maps and report on flux measurements of the outer planets. These high fidelity flux measurements will allow intercalibration of CMB experiments.

By marking faint gravitational wave echoes, these CMB observatories, and others like them, are perhaps unveiling meteoric events in the embryonic universe. In combination with data gathered by other cosmological observations, such measurements are furthering our understanding of the cosmos.

In these brief conclusions, I present a qualitative discussion regarding our ability to constrain early universe models. I will also speculate about future developments. Although I try to offer an impartial view, my opinion on some matters may still have shone through.

The idea of cosmic inflation is incredibly powerful and compelling. It offers numerous predictions, many of which preceded observations, and withstood detailed scrutiny. Inflation not only solves the horizon and flatness problem, but also predicts: A nearly scale invariant

temperature power spectrum; current best limits are $n_S = 0.9603 \pm 0.0073$ [49]. A flat universe; best estimates combining CMB, Type Ia SNe, and BAO observations suggest $\Omega = (0.9995 \pm 0.0065) \Omega_{\text{Crit}}$ [49]. Minimal non-Gaussianity; so far, the best limit is $f_{\text{NL}}^{\text{local}} = 2.7 \pm 5.8$ [268]. Cosmological adiabaticity; this is well established [97]. Super horizon modes, clear from the first release of the *WMAP* analysis and from the 2003 flight of BOOMERanG [269, 270]. The idea posits a scalar field, and indirect observations consistent with predictions about the Higgs scalar field, reported in 2012 [59, 60], provide evidence for the existence of scalar fields. If this were a football match, inflation has so far dribbled, in a “Messi-like” manner, through a network of staunch defenders.¹

Inflation is really an umbrella term for a number of models which utilize a common set of ideas. Not long ago, there was a real dearth of information that differentiated constituents of the standard inflationary paradigm. Now, limits on non-Gaussianity, spectral tilt, and the *B*-mode amplitude – most recently from [49, 88] – have begun to discriminate between previously viable inflationary theories. As non-Gaussianity and *B*-mode amplitudes are further constrained, more models will slowly become inconsistent with observations.

The expression “Remarkable Simplicity*” could have been a concise title of a news article describing the *Planck* HFI 2013 results.² The data showed no significant evidence for deviations from Gaussianity in the anisotropies of the CMB. Proponents of the standard inflationary paradigm generally view this as a remarkable victory for the theory. In fact, some of the simplest models of inflation appear consistent with observations [271]. Although inflation has triumphed over multiple trials, it has been pointed out – fairly so – that this does not prove the theory.

A critic could fairly argue that discussions of the inflationary paradigm have been simplistic, with professors and grad students alike – myself included – using catch phrases that involve “holy grails” and “smoking guns.” The laws of physics are never proven – they can only

¹My apologies for the cryptic analogy; the 2014 World Cup is ongoing.

²The asterisk referring to the $\sim 3\sigma$ deviation from ΛCDM at large angular scales. Real news article titles included: “*Planck* reveals ‘almost perfect’ universe”, “First *Planck* results: the Universe is still weird and interesting”, and my favorite “Universe as an Infant: Fatter Than Expected and Kind of Lumpy.”

aspire to be falsified – and the detection of primordial B -modes would not necessarily mean that inflation took place [96]. Others suggest that the inflationary paradigm is incomplete or even defective [272]. For example, critics might contend that these theories represent unimaginable fine-tuning [273, 274], although counterviews exist [275, 276, 277]. Similar critiques of inflation often describe the concept of eternal inflation, a consequence of quantum fluctuations in the inflaton potential, as highly unfeasible and unpredictable [278]. Nevertheless, it is the case that the weight of circumstantial evidence is consistent with the expectations of the simplest models of in the paradigm. Whether this points to a fundamental theory or a convenient phenomenological framework is a matter of debate. This discussion has many facets, some bordering on philosophy; it will likely continue indefinitely. We can, however, feel optimistic that the group of currently funded CMB observatories will significantly constrain the model space. A large subset of the following experiments will publish results in the near future: ACTPol [80], ABS [72, 109], EBEX [186], *Planck* [64], POLARBEAR [92], SPIDER [203], and SPTpol [279].

Constraining primordial B -modes requires electromagnetic, spatial, and angular constraints. Together, *Planck* HFI and SPIDER stand poised to fulfill these three requirements. The complete sky coverage allows us to constrain both the degree angular scales, where primordial B -modes are expected to peak, as well as the reionization bump on the largest angular scales. Similarly, SPIDER’s 8% sky coverage will allow us to verify isotropy of any B -mode detection. Finally, *Planck*’s wide spectral coverage, and the frequency windows afforded to balloon-borne experiments, will enable effective characterization of Galactic foregrounds.

The inflationary paradigm currently offers the best model of the early universe. In combination with novel instrumentation technologies [109, 280], future developments in the field hinge on our ability to maintain growth in production rate of photon-limited detectors [110]. This offers a plethora of challenges centered on automation and efficiency. SPIDER is the next step in that direction.

Appendix A

Thermal Modeling

This appendix describes the theory and some of the assumptions that were used in creating a thermal model for the SPIDER flight cryostat.

A.1 Conduction

Thermal conduction between two points in a linear isotropic medium with a given cross-sectional area A and length L can be calculated if the thermal conduction, k , is known as a function of temperature:

$$Q = \frac{A}{L} \int_{T_1}^{T_2} k(T) dT. \quad (\text{A.1})$$

The model evaluates this integral for various materials connecting different stages, including stainless steel, phosphor bronze, aluminum alloys, and G-10.

The heat input to the main tank goes into boiling off cryogenics such that a steady gas flow out of the MT will be established at equilibrium. This gas is forced to go through the heat exchangers at VCS1 and VCS2, providing negative feedback. The cooling power supplied

by the heat exchangers on VCS1 can be written as

$$Q = \dot{m}\eta \int_{T_{\text{MT}}}^{T_{\text{VCS1}}} C_p(T) dT, \quad (\text{A.2})$$

where \dot{m} is the mass flow through the heat exchangers, η is the heat exchanger's efficiency, C_p is the specific heat of the gas, and T_{MT} and T_{VCS1} are the mean temperatures of the MT and VCS1 respectively. The cooling power to VCS2 can be written in a similar manner.

The flight cryostat has MLI installed in such a way that reduces compression, which would otherwise increase conduction through the insulation, and enables proper evacuation of interstitial gas [153, 154, 281]. However, thermal conduction is inevitable when adjacent layers are in sporadic contact. A considerable fraction of the thermal budget is attributed to MLI conduction. Due to complex geometries, uncertainties about compressive loads, etc. it becomes difficult to estimate the MLI thermal conductance as a function of temperature. Instead, somewhat empirical estimates are adopted. In the absence of a leak, the pressure inside the VV remains roughly constant after equilibrium has been reached. Gas particles that have not condensed will then conduct heat between stages at a steady rate. If the pressure is sufficiently low the heat conduction through the MLI due to rarefied gas is negligible. Thermal conduction is then due to electron and phonon propagation. As the rarefied gas pressure increases, there comes a point when gas conductance becomes non-negligible. This can be due to outgassing and poor evacuation of MLI layers. MLI conduction is then roughly linearly proportional to gas number densities. At this point, gas is still in the free-molecular regime, although it is starting to contribute significantly to heat transfer [282, 283]. This behavior can also be seen in cryogenic systems that have microscopic leaks. Small leaks can be negated by the installation of activated charcoal or zeolite adsorbers [164, 162].

The thermal model has free parameters which describe the effective thermal conductance, excluding radiative transfer, through the MLI. Thermal characterization tests have helped improve this estimate. It is assumed that the conductivity through MLI, denoted k_{MLI} ,

is in the range of 0.5–2.0 $\mu\text{W}/\text{cm}/\text{K}$, which is comparable to observed values from other experiments [284].

A.2 Radiation

Estimating radiative coupling between various stages is difficult due to the complex structure of the SPIDER flight cryostat. No attempts have been made to calculate view factors for the intricate cylindrical geometries inside the flight cryostat. Instead, different stages are assumed to couple radiatively in the same way as parallel planes. This approximation becomes more accurate when the physical separation between stages is small. The estimated effective area of each stage then becomes a free parameter incorporating properties such as view factors.

The Hagen–Rubens relation, derived from Maxwell’s equations, states that the normal spectral emissivity, ϵ , of a conductor is [285, 286]

$$\epsilon = \sqrt{16\pi\varepsilon_0\rho\nu}, \quad (\text{A.3})$$

where ρ is the resistivity of the conductor, ε_0 is the vacuum permittivity, and ν is the frequency of radiation emitted. It can also be shown that emissivity is angle-dependent, and in the case of conductors, that the emissivity is highest at large angles normal to the plane. In order to simplify calculations, this model uses the total hemispherical emissivity (from here on simply referred to as emissivity), which is the emissivity averaged over frequency and angle. Temperature dependence of emissivity is generally assumed to follow a power law. The most simple approximation assumes a linear temperature profile through the MLI [153].

The radiative heat load per unit area between two infinite blackbody planes, $\epsilon = 1$, at temperatures T_C and T_H is

$$Q = \sigma(T_H^4 - T_C^4), \quad (\text{A.4})$$

where σ is the Stefan–Boltzmann constant. The net radiative heat transfer per unit area

between two parallel infinite plains at temperature T_H and T_C with emissivities ϵ_H and ϵ_C can be shown to be

$$Q = \sigma \left(\frac{\epsilon_H \epsilon_C}{\epsilon_H + \epsilon_C - \epsilon_H \epsilon_C} \right) (T_H^4 - T_C^4). \quad (\text{A.5})$$

This is done by summing up infinite contributions due to reflection and absorption of radiation, assuming no transmission. More complex geometries require the calculation of view factors [152]. For $N + 1$ parallel infinite surfaces, here representing the layers in an MLI blanket, the net radiative heat transfer per unit area can be shown to be

$$Q = \sigma \left(\sum_{i=1}^N \frac{1}{\epsilon_{i,i+1}} \right)^{-1} (T_{N+1}^4 - T_1^4), \quad (\text{A.6})$$

where

$$\epsilon_{i,i+1} \equiv \frac{\epsilon_i \epsilon_{i+1}}{\epsilon_i + \epsilon_{i+1} - \epsilon_i \epsilon_{i+1}} \quad (\text{A.7})$$

is the effective emissivity between layers i and $i + 1$, having emissivities ϵ_i and ϵ_{i+1} , and T_1 and T_{N+1} are the temperatures of the first and the last layers respectively. Equation A.6 is generally inversely proportional to N . Emissivity is temperature-dependent so the temperature of all layers do affect the radiative load. In estimating radiative coupling between the VV, VCS2, and VCS1, Equation (A.6) is employed with $N = 52$, and $N = 16$ respectively. No MLI is present between VCS1 and the main tank and so $N = 1$ in that particular case.

Appendix B

Fourier Transform Spectroscopy

In this appendix we discuss mathematical tools that have been useful for the design of Martin Pupplet Interferometers for the SPIDER experiment.

B.1 Wire Spacing of Polarizing Beamsplitters

The SPIDER FTS's use wire grids to effectively split radiation between the two arms of the interferometer. The transmission of wire grids as well as the optimal wire radius and spacing as a function of wavelength is studied in [287]. To first order, it is found that the optimal¹ wire radius, a , and spacing, d , is

$$a \simeq \left[\frac{\lambda^5}{(1 - [\sin \alpha \sin \beta]^2)^4 \pi^7 \sigma Z_0} \right]^{1/6}, \quad (\text{B.1})$$

$$d \simeq 2\pi a \quad (\text{B.2})$$

where λ is the wavelength of the incident radiation, α is the rotation angle of the wire grid with respect to the principles axis of the beam, in our case chosen to be $\alpha = \pi/4$, β is the incident angle of the beam with respect to the plane of the wire grid, also chosen to be $\beta = \pi/4$, σ is the conductance of the wire and $Z_0 = 119.9\pi \Omega$ is the impedance of free space.

¹Here optimal refers to the minimal cross polarization

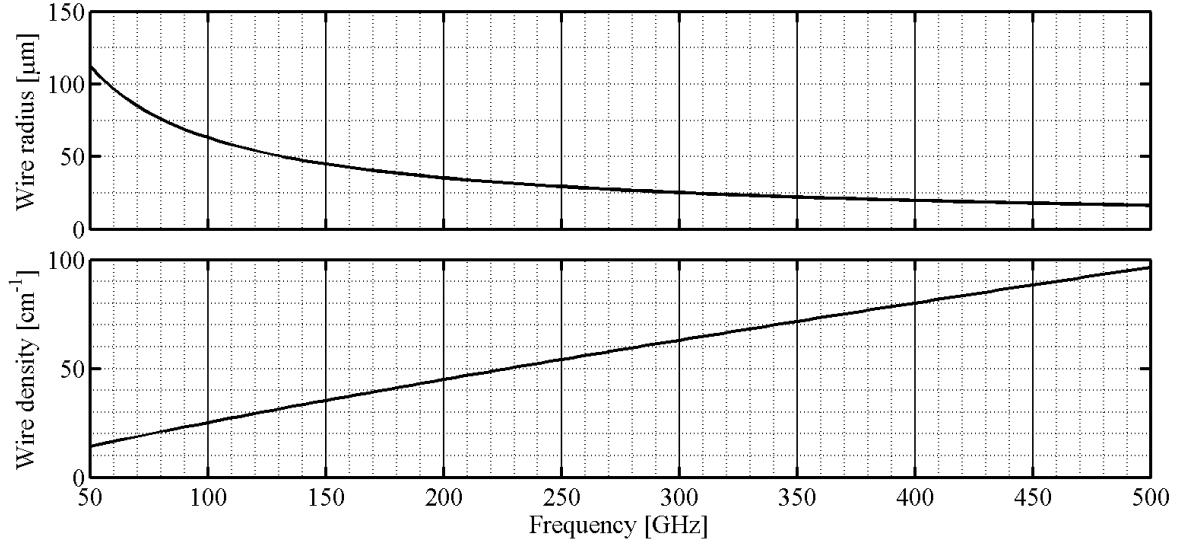


Figure B.1: The wire grid dimensions that give the lowest cross polar leakage as a function of frequency, see Equations B.1 and B.2.

The derivation is made in the approximation that $d \ll \lambda$. Figure B.1 shows the predictions of the above equation for the frequency range of interest for a CMB experiment. The complex reflection coefficients to lowest order become

$$R_{\parallel} \simeq \frac{-1}{1 + j[(2d \sin \alpha)/\lambda] \ln[d/(2\pi a)]}, \quad (\text{B.3})$$

$$R_{\perp} \simeq -j \sin \alpha \frac{\pi^2 a^2}{\lambda d}. \quad (\text{B.4})$$

All of the SPIDER FTS wire grids use gold plated tungsten wire with radius $a = 12.7 \mu\text{m}$ and wire spacing of 200 wires per inch, corresponding to $d = 39.4 \mu\text{m}$. Figure B.2 shows the absolute value of the complex reflection coefficients for three different wire grids. Solid lines represent $(1 - R_{\parallel})$, the wire grid transmission of light with an electric field parallel to the wires, while dashed lines represent R_{\perp} , the reflectance of light that is perpendicular to the wires in the grid.

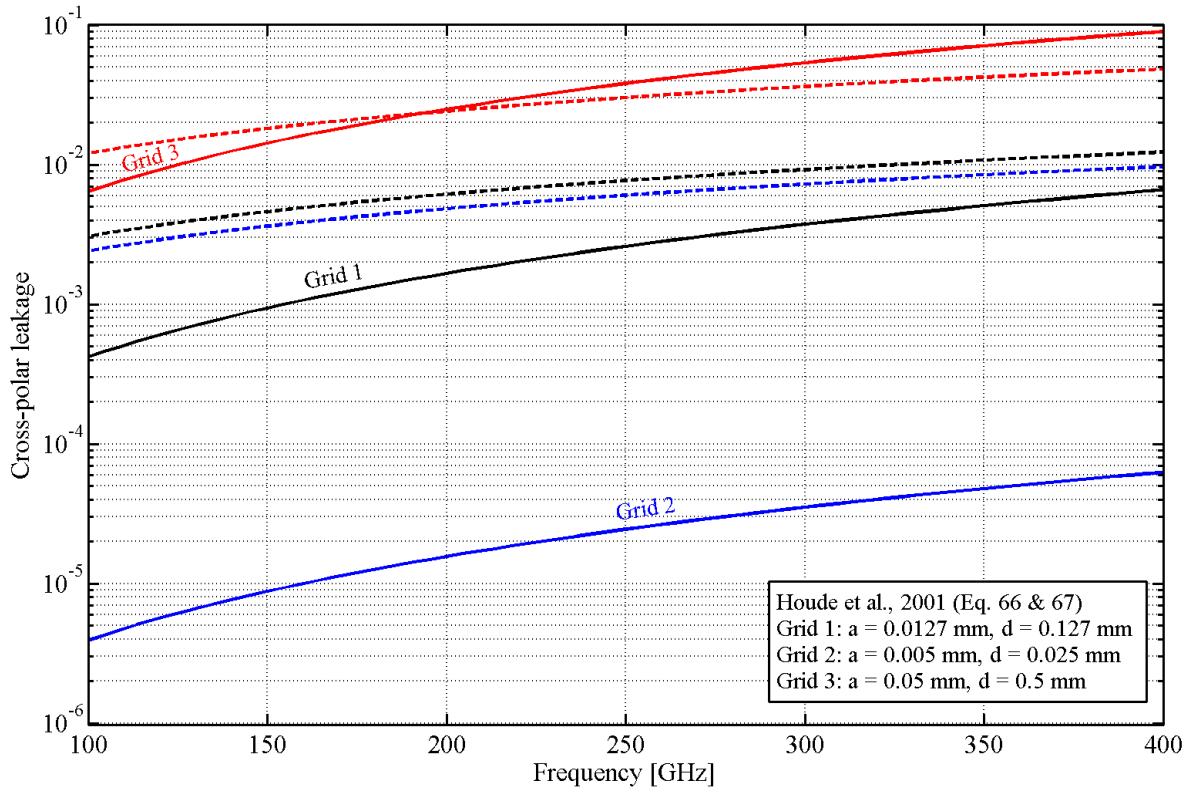


Figure B.2: The cross polar leakage for three different wire grids as calculated by Equations B.3 and B.4. Solid lines represent $(1 - R_{||})$, the wire grid transmission of light with an electric field parallel to the wires, while dashed lines represent R_{\perp} , the reflectance of light that is perpendicular to the wires in the grid.

B.2 Propagation of Electric Fields in an FTS

Let's define a coordinate system in the plane that is perpendicular to the propagation of light at any time. In some cases, a bundle of light will propagate towards a wire grid, in that case Figure B.3 represents the view from the wavefront along the direction of propagation. The blue lines represent the wires in a wire grid. In this system, \hat{x} and \hat{y} define a right hand rule coordinate system with \hat{y} pointing up from the plane of the FTS, or perpendicular to the plane of incidence, whereas \hat{a} and \hat{b} are parallel and perpendicular to the wires in the wire grid respectively, such that

$$\hat{a} = \frac{1}{\sqrt{2}}(\hat{x} + \hat{y}), \quad \hat{b} = \frac{1}{\sqrt{2}}(\hat{y} - \hat{x}).$$

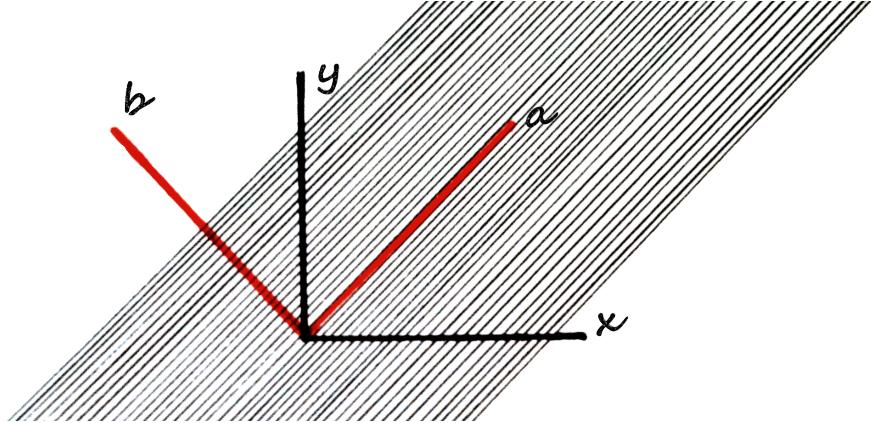


Figure B.3: Two coordinate systems formed with two sets of orthonormal basis vectors. The blue lines represent the wires in the wire grid. This is the image that a mirror sees looking towards the beamsplitter and the mobile rooftop mirror.

In the ideal scenario, when light interacts with the wire grid, any electric field oscillation which is aligned with the wires will be reflected, while the perpendicular part of the electric field is transmitted through the wires. The electric field of an arbitrarily polarized light bundle being emitted from a monochromatic source with wavelength λ inside the FTS can be written as

$$\mathbf{E} = X_0 \cos(\omega t) \hat{\mathbf{x}} + Y_0 \cos(\omega t) \hat{\mathbf{y}}, \quad (\text{B.5})$$

where for unpolarized light, $X_0 = Y_0$. Note that $\omega = 2\pi\nu = 2\pi c/\lambda$. The light is collimated by the parabolic mirror and directed towards the wire grid. The electric field of the light which is reflected of the wire grid is $\mathbf{E}_r = (\mathbf{E} \cdot \hat{\mathbf{a}})(\hat{\mathbf{x}} + \hat{\mathbf{y}})$, or

$$\mathbf{E}_r = \frac{1}{\sqrt{2}}(X_0 + Y_0)(\hat{\mathbf{x}} + \hat{\mathbf{y}}) \cos(\omega t), \quad (\text{B.6})$$

whereas the transmitted light $\mathbf{E}_t = (\mathbf{E} \cdot \hat{\mathbf{b}})(\hat{\mathbf{y}} - \hat{\mathbf{x}})$,

$$\mathbf{E}_t = \frac{1}{\sqrt{2}}(Y_0 - X_0)(\hat{\mathbf{y}} - \hat{\mathbf{x}}) \cos(\omega t). \quad (\text{B.7})$$

The bundles of light, are reflected twice of the rooftop mirrors as they traverse the two arms of the FTS. After two reflections the electric field in the plane of incidence, $\hat{\mathbf{x}}$, has flipped directions, resulting in

$$\mathbf{E}'_r = \frac{1}{\sqrt{2}}(Y_0 + X_0)(\hat{\mathbf{y}} - \hat{\mathbf{x}}) \cos(\omega t + \Delta), \quad (\text{B.8})$$

$$\mathbf{E}'_t = \frac{1}{\sqrt{2}}(Y_0 - X_0)(\hat{\mathbf{y}} + \hat{\mathbf{x}}) \cos(\omega t), \quad (\text{B.9})$$

where we have also added a relative phase shift to $\hat{\mathbf{E}}'_r$ accounting for any differences in the optical path lengths between the two arms. Before the two light bundles can combine on the other side of the wire grid they have had to interact with the wires in the grid a second time. The bundle that was transmitted now has to be reflected and vice versa. Calculating the inner products like before we find $\mathbf{E}''_r = (\mathbf{E} \cdot \hat{\mathbf{b}})(\hat{\mathbf{y}} - \hat{\mathbf{x}})$,

$$\mathbf{E}''_r = \frac{1}{2}(Y_0 + X_0)(\hat{\mathbf{y}} - \hat{\mathbf{x}}) \cos(\omega t + \Delta), \quad (\text{B.10})$$

and $\mathbf{E}''_t = (\mathbf{E} \cdot \hat{\mathbf{a}})(\hat{\mathbf{x}} + \hat{\mathbf{y}})$,

$$\mathbf{E}''_t = \frac{1}{2}(Y_0 - X_0)(\hat{\mathbf{y}} + \hat{\mathbf{x}}) \cos(\omega t). \quad (\text{B.11})$$

The two bundles of light combine on the other side of the wire grid. The total electric field becomes the sum of the two $\mathbf{E}_f = \mathbf{E}''_r + \mathbf{E}''_t$,

$$\begin{aligned} \mathbf{E}_f = \hat{\mathbf{x}}[(Y_0 - X_0) \cos(\omega t) - (Y_0 + X_0) \cos(\omega t + \Delta)]/2 \\ + \hat{\mathbf{y}}[(Y_0 + X_0) \cos(\omega t + \Delta) + (Y_0 - X_0) \cos(\omega t)]/2. \end{aligned} \quad (\text{B.12})$$

In the limit of a completely unpolarized source, $X_0 = Y_0$, we arrive at

$$\mathbf{E}_{f,\text{unpol}} = \frac{1}{2}(-\hat{\mathbf{x}}(Y_0 + X_0) \cos(\omega t + \Delta) + \hat{\mathbf{y}}(Y_0 + X_0) \cos(\omega t + \Delta)). \quad (\text{B.13})$$

The signal that is measured in a detector is proportional to the squared magnitude of the inner product of the detector polarization vector, $\hat{\mathbf{e}}$, with the electric field. In other words,

$$I \propto |\mathbf{E}_f \cdot \hat{\mathbf{e}}|^2. \quad (\text{B.14})$$

Writing $\hat{\mathbf{e}} = x_0 \hat{\mathbf{x}} + y_0 \hat{\mathbf{y}}$ we find that

$$\mathbf{E}_{f,\text{unpol}} \cdot \hat{\mathbf{e}} = \frac{1}{2} (-x_0(Y_0 + X_0) \cos(\omega t + \Delta) + y_0(Y_0 + X_0) \cos(\omega t + \Delta)). \quad (\text{B.15})$$

Calculating the squared modulus and taking the time average we find that there will be no signal modulation. Note that the time average of terms such as $\cos^2(\omega t + \Delta)$ is simply 1/2. Any time-dependent phase shift $\Delta = 2kx_m(t)$ is negated by the fast signal modulation provided by the term $\omega t = 2\pi\nu t$ where $\nu \approx 100$ GHz. Here $x_m(t)$ could encapsulate the fact that one of the mirrors can be translated by a wire grid. Assuming that the source is 100% polarized, which is true for example if $X_0 = E_0$ and $Y_0 = 0$, we instead get

$$\begin{aligned} 2\mathbf{E}_{f,\text{pol}} &= -\hat{\mathbf{x}}E_0[\cos(\omega t) + \cos(\omega t + \Delta)] + \hat{\mathbf{y}}E_0[\cos(\omega t + \Delta) + \cos(\omega t)] \\ &= -\hat{\mathbf{x}}E_0 \cos\left(\omega t + \frac{\Delta}{2}\right) \cos\left(\frac{\Delta}{2}\right) + \hat{\mathbf{y}}E_0 \cos\left(\omega t + \frac{\Delta}{2}\right) \cos\left(\frac{\Delta}{2}\right), \end{aligned}$$

where we have used the following trigonometric identities

$$2 \cos \alpha \cos \beta = \cos(\alpha + \beta) + \cos(\alpha - \beta), \quad (\text{B.16})$$

$$\cos^2(\alpha) = \frac{1 + \cos(2\alpha)}{2}. \quad (\text{B.17})$$

Aligning our detector with the $\hat{\mathbf{x}}$ direction and measuring the squared and time-averaged modulus gives

$$|\mathbf{E}_f \cdot \hat{\mathbf{x}}|^2 = \frac{E_0^2}{8} \cos^2\left(\frac{\Delta}{2}\right), \quad (\text{B.18})$$

as the signal power as measured by our detector. The phase shift, Δ , is time-dependent if the mobile mirror is being moved by the linear stage. Let's assume that the mirror is moving at a steady rate $x(t) = 2v_0t$ such that $\Delta = 4\pi v_0 t / \lambda$, where the additional factor of two accounts for the fact that the light path is twice the translation of the mirror. This will cause signal modulation with the frequency $\omega_{\text{mirror}} = 4\pi v_0$.

If instead we align our detector such that $\hat{\mathbf{e}} = (\hat{\mathbf{x}} + \hat{\mathbf{y}})/\sqrt{2}$ we get

$$\begin{aligned}\mathbf{E}_{\text{f,pol}} \cdot \hat{\mathbf{e}} &= -E_0/\sqrt{2} \cos\left(\omega t + \frac{\Delta}{2}\right) \cos\left(\frac{\Delta}{2}\right) + E_0/\sqrt{2} \cos\left(\omega t + \frac{\Delta}{2}\right) \cos\left(\frac{\Delta}{2}\right) \\ &= 0.\end{aligned}$$

From this discussion it is clear that the signal measured by our detector will depend on its polarization orientation w.r.t. the output of the FTS. Maximum signal amplitude is observed if the polarization of the detectors is parallel to either $\hat{\mathbf{x}}$ or $\hat{\mathbf{y}}$. If on the other hand the detector is at a 45 degree angle w.r.t. the $\hat{\mathbf{x}}\text{-}\hat{\mathbf{y}}$ axis no scan modulation will be observed in the interferogram. For any orientation other than the two extreme examples listed above, the final electric field that is measured by the detector has an interesting time-dependence. It will fluctuate between being linearly polarized and circularly polarized and then linearly polarized again. All the time in between it will be elliptically polarized.

B.3 Collimating mirror

A number of papers derive the surface equations for off-axis parabolic mirrors [288, 289, 290, 291].² A few parabolic off-axis collimating mirrors with focal lengths $f = 5.5$ inch were machined on a three axis CNC mill. One can derive the surface equation for such a mirror by starting with the equation for a parabolic surface

$$Z = \frac{1}{4f}(X^2 + Y^2), \quad (\text{B.19})$$

²I am not convinced that all of them give the same result!

and performing a coordinate transformation followed by rotation. The resulting system of equations is

$$X = x_0 + x \cos(\theta) - z \sin(\theta), \quad (\text{B.20})$$

$$Y = y, \quad (\text{B.21})$$

$$Z = x_0^2/4f + x \sin(\theta) + z \cos(\theta). \quad (\text{B.22})$$

Where x, y, z are the coordinates after a transformation into a system where z is normal to the machined plate and x, y define the plane with origin centered on the plate. Here, x_0 corresponds to the translation that needs to be made in the direction of X . For that reason, we must have $x_0 = 2f \tan(\theta)$. Solving this system of equations for z with the included constraint on x_0 and assuming $\theta = \pi/4$ leads to the following equation for the surface of the mirror

$$z = 4\sqrt{2}f + x \pm \sqrt{8f(4f + \sqrt{2}x) - y^2}. \quad (\text{B.23})$$

Note that the effective focal length of this system, x_0 , is equal to $2f$, twice the focal length of the original parabola.

The collimating mirrors that were constructed for the SPIDER FTS's have a diameter of $D = 6$ inch, but the wire grid mirror constricts the maximum diameter of a circular beam that could propagate through the system to $d = 4.6$ inch, resulting in an effective f-number of f/2.6. The mirror was designed so that the power inside a 5 deg FWHM Gaussian beam located at the effective focus would propagate through the system.

B.4 G-code for collimating mirror

The following code segment was used to cut the parabolic mirror described in Section 2.9.1 and shown in Figure 2.29. This script might be useful for people wanting to cut a mirror on the CNC in the Princeton Physics machine shop.

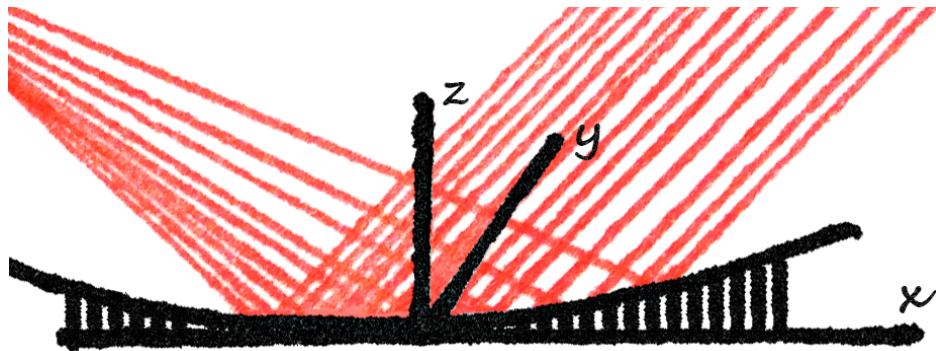


Figure B.4: The shape of the collimating mirror as calculated from Equation B.23. Note that the mirror surface is only left-right symmetric, meaning that there is only one proper mounting orientation for the mirror inside the FTS.

```

1 G0G90X0Y0T3M6
2 ; Speed
3 S2000M3
4 ; Using a 1/4 inch ball mill for cutting, all units are inches
5 J1 = 0.010; Distance between mill passes
6 R1 = 4.3426; The cutting radius
7 A1 = 5.5; The focal length
8 B1 = 0.4852; The furthest down the mirror will go
9
10 Y1 = (-1.0)*R1
11 ; Move to home position above mirror
12 G0X0.Y1Z0.25
13 M8
14
15 ; Increment Y up and calculate X on circle given Y
16 L2
17 Y1 = Y1+J1
18 X2 = (-1.0)*SQR(R1*R1-Y1*Y1)
19 X1 = J1
20
21 ; Increment X down just below calculated value
22 L3

```

```

23 X1 = X1-J1
24 ?X1<GT>X2
25 =L3!
26
27 ; Move to correct X and Y position
28 X2 = X2*(-1)
29 G0X1Y1Z.25
30
31 ; Increment X up and check if has reached its max for this pass
32 L4
33 X1 = X1+J1
34 ?X1<GT>X2; If X at max, go back, increment Y and begin next pass
35 =L2!
36
37 Z1 = -B1+4*SQR(2)*A1+Y1-SQR(8*A1*(4*A1+SQR(2)*Y1)-2*X1*X1)
38
39 G1X1Y1Z1F10.0; Feed mill to new position
40 ?Y1<LT>R1; If Y is inside circle, continue machining
41 =L4!; Otherwise, end program
42
43 G0X0.Y0.Z0.25; Back to home position at end of program

```

Appendix C

Gauss-Hermite Math

A natural function basis for Gaussian beam parametrization is found in the Gauss-Hermite functions. This appendix describes some of the math related to these functions. This information is relevant for discussion in Section 3.4.

C.1 Preliminaries

The “physicist’s” Hermite polynomials, suitable to describe the eigenstates of the quantum harmonic oscillator are defined as

$$H_n(x) = e^{x^2/2} \left(x - \frac{d}{dx} \right)^n e^{-x^2/2}. \quad (\text{C.1})$$

The first few polynomials are

$$\begin{aligned} H_0(x) &= 1, \\ H_1(x) &= 2x, \\ H_2(x) &= 4x^2 - 2, \\ H_3(x) &= 8x^3 - 12x, \\ H_4(x) &= 16x^4 - 48x^2 + 12. \end{aligned} \quad (\text{C.2})$$

Note that they alternate between being even or odd functions [292]. The explicit expression for these polynomials is

$$H_n(x) = n! \sum_{m=0}^{\lfloor n/2 \rfloor} \frac{(-1)^m}{m!(n-2m)!} (2x)^{n-2m}. \quad (\text{C.3})$$

where $\lfloor \cdot \rfloor$ represents the floor function [293].

General forms of Gaussian integrals have well known analytical expressions. In particular

$$\int_{-\infty}^{\infty} x^n e^{-x^2/2\sigma^2} = \begin{cases} \sqrt{2\pi}(n-1)!!\sigma^{n+1} & \text{if } n \text{ is even} \\ 0 & \text{if } n \text{ is odd} \end{cases} \quad (\text{C.4})$$

Equations C.3 and C.4 suggest that

$$\begin{aligned} \int_{-\infty}^{\infty} H_n(x/\sigma) e^{-x^2/2\sigma^2} dx &= n! \sum_{m=0}^{\lfloor n/2 \rfloor} \frac{(-1)^m 2^{n-2m}}{m!(n-2m)!} \int_{-\infty}^{\infty} (x/\sigma)^{n-2m} e^{-x^2/2\sigma^2} dx \\ &= \sqrt{2\pi}\sigma n! \sum_{m=0}^{\lfloor n/2 \rfloor} \frac{(-1)^m 2^{n-2m}}{m!(n-2m)!} \int_{-\infty}^{\infty} u^{n-2m} e^{-u^2/2} du \\ &= \sqrt{2\pi}\sigma n! \sum_{m=0}^{\lfloor n/2 \rfloor} \frac{(-1)^m 2^{n-2m}}{m!(n-2m)!} (n-2m-1)!! \\ &= \sqrt{2\pi}\sigma n! \sum_{m=0}^{\lfloor n/2 \rfloor} \frac{(-1)^m 2^{n-2m}}{m!(n-2m)!!}, \end{aligned} \quad (\text{C.5})$$

if n is even, but zero if n is odd.

C.2 Gauss-Hermite Decomposition

As mentioned in Section 3.4, the two dimensional Gauss-Hermite functions are defined as

$$\psi_{n_1, n_2}(x, y) = \frac{H_{n_1}(x/\sigma_x) H_{n_2}(y/\sigma_y)}{\sqrt{2^{n_1+n_2} n_1! n_2!}} e^{-\frac{1}{2}(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2})}. \quad (\text{C.6})$$

The Gauss-Hermite functions represent an orthogonal function basis \mathbb{R}^2 .

We can use these functions to describe beams that can be approximated by an elliptical Gaussian. Although it is possible to fit all parameters at the same time, we choose to fit the parameters of an elliptical Gaussian to the data before calculating the Gauss-Hermite decomposition. This way the Gauss-Hermite function bases is fixed. The parameters of the elliptical Gaussian fit are σ_x and σ_y , which describe the width of the beam, the tilt angle ϕ which defines the tilt of the elliptical Gaussian with respect to the nominal x and y axis in the map, and the beam center with coordinates (c_x, c_y) corresponding to a translation from the center pixel. The Gauss-Hermite decomposition is then performed with the rotated coordinates where

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} \begin{pmatrix} x - c_x \\ y - c_y \end{pmatrix}. \quad (\text{C.7})$$

The Gauss-Hermite coefficients can be decomposed from a raw map by calculating the overlap integral

$$\hat{s}_n = \int d\hat{\mathbf{x}} \psi_n(\hat{\mathbf{x}}') M(\hat{\mathbf{x}}) N_{\text{obs}}(\hat{\mathbf{x}}), \quad (\text{C.8})$$

where $M(\hat{\mathbf{x}})$ is the raw map and $N_{\text{obs}}(\hat{\mathbf{x}})$ is a map of the number of observations, or time stream samples, going into each pixel. This form of observation weighing reduces the effect of pixels with poor sampling. Poor sampling can be caused by a number of effects, including scan strategy and the flagging of time-ordered-data neighboring cosmic ray events. Since the maps are discretized the integrals are estimated as sums

$$\hat{s}_n = \sum_i dA \psi_n(\hat{\mathbf{x}}'_i) M(\hat{\mathbf{x}}_i) N_{\text{obs}}(\hat{\mathbf{x}}_i), \quad (\text{C.9})$$

where i represent the pixel indexing and dA is the area of each pixel. The planetary maps are finite and the Gauss-Hermite functions form a complete basis only on the infinite 2-D plane. This tends to bias the decomposed coefficients. Debiasing is done through calculation

of the overlap matrix

$$I_{mn} = \sum dA \psi_n(\hat{\mathbf{x}}'_i) \psi_m(\hat{\mathbf{x}}_i) N_{\text{obs}}(\hat{\mathbf{x}}_i), \quad (\text{C.10})$$

which is then used to calculate

$$s_n = I_{mn}^{-1} \hat{s}_m, \quad (\text{C.11})$$

where s_n are the debiased Gauss-Hermite coefficients. The Gauss-Hermite coefficients can also be decomposed in the time-domain and we have found that this approach is more precise as it preserves all positional information. The time-domain approach is described in Huffenberger et al. [234]. In essence, this approach replaces the map response, $M(\hat{\mathbf{x}}_i)$, with the timestream, removing N_{obs} from the above equations, and summing over all samples. I chose to present the map domain approach for completeness.

The *Planck* scan strategy is such that the peak in a planetary transit can easily be missed. By fitting a Gaussian function to the raw map the true value of the peak signal can be recovered to some extent. The subsequent Gauss-Hermite decomposition helps to capture the non-Gaussian properties of the beam. The beam can then be reconstructed at any resolution using the parameters of the elliptical Gaussian and the coefficients of the Gauss-Hermite functions. The estimated response function is

$$\Psi(\hat{\mathbf{x}}) = \sum_n s_n \psi_n(\hat{\mathbf{x}}'), \quad (\text{C.12})$$

from which maps can be reconstructed and numerically integrated to produce the total beam solid angle, Ω , according to

$$\Omega = \sum_i \Psi(\hat{\mathbf{x}}_i) dA / \Psi_{\text{max}}, \quad (\text{C.13})$$

where Ψ_{max} is the maximum response in the reconstructed map, dA is the pixel area, and the summation is performed over all pixels. The discussion above suggests that a semi-analytical expression for the total beam solid angle should exist. Using Equations C.3 and C.4 and the decomposed Gauss-Hermite coefficients the explicit expression for the beam solid angle, Ω , is

found to be

$$\begin{aligned}
\Omega &= \sum_{n_1, n_2=0}^{N_{\max}-1} s_{n_1, n_2} \iint \frac{H_{n_1}(x/\sigma_x) H_{n_2}(y/\sigma_y)}{\sqrt{2^{n_1+n_2} n_1! n_2!}} e^{-\frac{1}{2}(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2})} dx dy / \Psi_{\max} \\
&= \frac{1}{\Psi_{\max}} \sum_{n_1, n_2}^{N_{\max}-1} s_{n_1, n_2} \times \int_{-\infty}^{\infty} \frac{H_{n_1}(x/\sigma_x)}{\sqrt{2^{n_1} n_1!}} e^{-\frac{x^2}{2\sigma_x^2}} dx \times \int_{-\infty}^{\infty} \frac{H_{n_2}(y/\sigma_y)}{\sqrt{2^{n_2} n_2!}} e^{-\frac{y^2}{2\sigma_y^2}} dy \\
&= \frac{2\pi\sigma_x\sigma_y}{\Psi_{\max}} \sum_{n_1, n_2}^{\text{Even}} s_{n_1, n_2} \sqrt{\frac{n_1! n_2!}{2^{n_1+n_2}}} \times \sum_{m=0}^{n_1/2} \frac{(-1)^m 2^{n_1-2m}}{m!(n_1-2m)!!} \sum_{k=0}^{n_2/2} \frac{(-1)^k 2^{n_2-2k}}{k!(n_2-2k)!!}.
\end{aligned} \tag{C.14}$$

The summation only takes place for even values of n_1 and n_2 , since the odd valued functions integrate to zero. This means that only functions with n_1 and n_2 as even numbers contribute to the solid angle. This is an important property of the Gauss-Hermite polynomials. It should also be noted that the Gauss-Hermite coefficients are decomposed such that they describe the amplitude of Gauss-Hermite functions on a coordinate system that has been translated, scaled, and finally rotated with respect to some nominal coordinate system. However, neither the translation, nor the rotation, will affect the total solid angle of the Gauss-Hermite beam, while the scaling is accounted for in the above derivation.

Unfortunately we do not have an expression for Ψ_{\max} and so it must be extracted from the reconstructed map. Note that $\Psi_{\max} = \Psi(c_x, c_y)$ does not necessarily hold although it is often a very good approximation. Equation C.14 can be implemented with a few lines of code and is less computationally demanding than the numerical approach. Fast methods, like the semi-analytical expression described above, could prove useful for experiments with tens of thousands of detectors.

C.3 Window Functions

The signal on the sky as measured by any CMB experiment is the convolution of the instrument response function with the true sky signal. We are interested in the power spectrum of

the true sky signal and need to be able to correct for the finite angular resolution of the instrument. This involves calculating the window function. We follow the derivation in Dodelson [97]

$$W_\ell = \int d\Omega \int d\Omega' \Psi(\hat{\mathbf{x}}) \Psi(\hat{\mathbf{x}}') P_\ell(\cos(\hat{\mathbf{x}} \cdot \hat{\mathbf{x}}')), \quad (\text{C.15})$$

where Ψ is the beam response¹, P_ℓ is the Legendre polynomial of order ℓ , and the following relation between spherical harmonics and Legendre polynomials has been used

$$\sum_m Y_{\ell m}(\hat{\mathbf{x}}) Y_{\ell m}^*(\hat{\mathbf{x}}') = (2\ell + 1) P_\ell(\hat{\mathbf{x}} \cdot \hat{\mathbf{x}}'). \quad (\text{C.16})$$

For what follows we will use the integral representation of the Bessel function of the first kind [294]

$$J_n(\ell) = \frac{1}{2\pi} \int_{-\pi}^{\pi} d\phi \exp(i[n\ell - \ell \sin(\phi)]). \quad (\text{C.17})$$

We can simplify this expression for the case where $n = 0$ to find

$$J_0(\ell) = \frac{1}{2\pi} \int_0^{2\pi} d\phi \exp(-i\ell \cos(\phi)), \quad (\text{C.18})$$

where we have used the phase relation and periodicity of trigonometric functions. In the limit of large ℓ and small $|\hat{\mathbf{x}} - \hat{\mathbf{x}}'|$ (the small angle approximation) Equation 8.722.2 in [294] claims that one can write

$$\begin{aligned} \lim_{\ell \rightarrow \infty} P_\ell(\cos(|\hat{\mathbf{x}} - \hat{\mathbf{x}}'|)) &= J_0([\ell + 0.5]|\hat{\mathbf{x}} - \hat{\mathbf{x}}'|) \\ &= \frac{1}{2\pi} \int_0^{2\pi} d\phi e^{-i[\ell + 0.5]|\hat{\mathbf{x}} - \hat{\mathbf{x}}'| \cos(\phi)} \\ &= \frac{1}{2\pi} \int_0^{2\pi} d\phi e^{-i\ell \cdot (\hat{\mathbf{x}} - \hat{\mathbf{x}}')}. \end{aligned} \quad (\text{C.19})$$

To arrive at this last expression we had to promote $\ell + 0.5$ to a two dimensional vector ℓ and assume that ϕ is the angle between ℓ and $(\hat{\mathbf{x}} - \hat{\mathbf{x}}')$. One can decompose ℓ in Cartesian

¹Note that this is identical to Equation 3.3. We've changed symbols however, using $\hat{\mathbf{x}}$ instead of $\hat{\mathbf{n}}$.

coordinates, ℓ_x and ℓ_y , following

$$\ell + 0.5 = \sqrt{\ell_x^2 + \ell_y^2},$$

$$\phi = \arctan(\ell_y/\ell_x). \quad (\text{C.20})$$

For future reference, we define $\ell' = \ell + 0.5$. The inverse transformation is then

$$\ell_x = \ell' \cos(\phi),$$

$$\ell_y = \ell' \sin(\phi). \quad (\text{C.21})$$

Wielding these mathematical acrobatics, we note that Equation C.15 can be written as

$$\begin{aligned} W_\ell &= \int d\Omega \int d\Omega' \Psi(\hat{\mathbf{x}}) \Psi(\hat{\mathbf{x}}') P_\ell(\cos(|\hat{\mathbf{x}} - \hat{\mathbf{x}}'|)), \\ &= \frac{1}{2\pi} \int_0^{2\pi} d\phi \int d\mathbf{x} \Psi(\hat{\mathbf{x}}) e^{-i\ell \cdot \hat{\mathbf{x}}} \int d\mathbf{x}' \Psi(\hat{\mathbf{x}}') e^{i\ell \cdot \hat{\mathbf{x}}'}, \\ &= \frac{1}{2\pi} \int_0^{2\pi} d\phi \mathcal{F}(\Psi) \mathcal{F}^*(\Psi). \end{aligned} \quad (\text{C.22})$$

This shows that the Fourier transform of the beam response function is related to the beam window function in the flat space approximation.

We now direct our attention to the Gauss-Hermite functions. The generating function for the Hermite polynomials is [292]

$$\exp(2xt - t^2) = \sum_{n=0}^{\infty} H_n(x) \frac{t^n}{n!}. \quad (\text{C.23})$$

Using this expression, the Fourier transform of the Gauss-Hermite functions can be easily derived. One finds that

$$\mathcal{F}(H_n(x) e^{-x^2/2}) = (-i)^n e^{-k^2/2} H_n(k). \quad (\text{C.24})$$

This means that the Fourier transform of a Gauss-Hermite function is a Gauss-Hermite function with an added phase. In other words, the Gauss-Hermite functions are eigenfunctions of the Fourier transform operator. This also suggests that the Fourier transform of the instrument response is

$$\begin{aligned}
\mathcal{F}(\Psi) &= \mathcal{F} \left(\sum_{n_1, n_2} s_{n_1, n_2} \frac{H_{n_1}(x/\sigma_x) H_{n_2}(y/\sigma_y)}{\sqrt{2^{n_1+n_2} n_1! n_2!}} e^{-\frac{1}{2}(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2})} \right) (\ell), \\
&= \sum_{n_1, n_2} s_{n_1, n_2} \frac{\mathcal{F}(H_{n_1}(x/\sigma_x) e^{-x^2/2\sigma_x^2})}{\sqrt{2^{n_1} n_1!}} \times \frac{\mathcal{F}(H_{n_2}(y/\sigma_y) e^{-y^2/2\sigma_y^2})}{\sqrt{2^{n_2} n_2!}}, \\
&= \sigma_x \sigma_y \sum_{n_1, n_2} s_{n_1, n_2} \frac{(-i)^{n_1} e^{-l_x^2 \sigma_x^2/2} H_{n_1}(l_x \sigma_x)}{\sqrt{2^{n_1} n_1!}} \times \frac{(-i)^{n_2} e^{-l_y^2 \sigma_y^2/2} H_{n_2}(l_y \sigma_y)}{\sqrt{2^{n_2} n_2!}},
\end{aligned} \tag{C.25}$$

where $\ell = (l_x, l_y)$ are used as the Fourier transform parameters and we have used the relation

$$\mathcal{F}(f(x/a))(l) = aF(la), \tag{C.26}$$

with $\mathcal{F}(f(x)) \equiv F(x)$. The Gauss-Hermite functions are defined on the coordinate system $x' - y'$ which has been translated, rotated, and scaled with respect to some other Cartesian coordinate system, $x - y$. The translation of points in the $x - y$ coordinate system results in a phase shift in Fourier-space which can be easily accounted for. The Fourier transform is angle preserving, meaning that any rotation of vectors in real-space will result in an identical rotation in Fourier-space. We can then choose to perform the rotation after the Fourier transform has been taken.

As Equation C.22 suggests, the calculation of the window function involves finding the angular average through an integral. For a given pixelization of the sky, and knowledge about the *Planck* scanning strategy, one should be able to calculate both the location within a pixel and the angle of incidence of the scanning beam with respect to some Cartesian coordinate system. The calculation of a pixel-based window function that takes into account

all observations that fall inside that pixel should then be straightforward. We find

$$\begin{aligned}
W_\ell(n) &= \frac{1}{2\pi} \int_0^{2\pi} d\phi \mathcal{F}(\Psi) \mathcal{F}^*(\Psi), \\
&= \frac{(\sigma_x \sigma_y)^2}{2\pi} \int_0^{2\pi} d\phi \sum_{n_1, n_2} s_{n_1, n_2}^2 \\
&\times \left| \frac{e^{-i(\ell' \cos(\phi) c_x + \frac{n_1 \pi}{2})} H_{n_1}(\ell' \cos(\phi) \sigma_x)}{e^{\ell'^2 \cos^2(\phi) \sigma_x^2/2} \sqrt{2^{n_1} n_1!}} \times \frac{e^{-i(\ell' \sin(\phi) c_y + \frac{n_2 \pi}{2})} H_{n_2}(\ell' \sin(\phi) \sigma_y)}{e^{\ell'^2 \sin^2(\phi) \sigma_y^2/2} \sqrt{2^{n_2} n_2!}} \right|^2, \quad (\text{C.27})
\end{aligned}$$

where c_x and c_y represent the translation of the beam with respect to some origin located within a distance from the pixel where the flat sky approximation is still valid. This semi-analytical expression for the scanning beam window function can be used in concert with information about the instrument scanning to calculate an effective beam window function. To find the window function for multiple observations one should perform an additional summation and normalization over samples before taking the squared modulus and calculating the numerical integral.

C.4 Ruze Envelope

Imperfections of reflector surfaces will scatter light and reduce the forward gain of an optical system [243]. The following functional form describes the Ruze envelope that is caused by scattering off a reflector system with a uniform distribution of surface deformations

$$G(\theta, \phi) = G_0(\theta, \phi) e^{-\overline{\rho^2}} + \left(\frac{2\pi c}{\lambda} \right)^2 e^{-\overline{\rho^2}} \sum_{n=1}^{\infty} \frac{\overline{\rho^2}^n}{n \times n!} e^{-(\pi c \sin(\theta)/\lambda)^2/n} \quad (\text{C.28})$$

where c is the correlation length of the surface deformation, λ is the wavelength, $G_0(\theta, \phi)$ is the ideal beam shape and $\overline{\rho^2}$ represents the variance of the phase errors. We want to calculate the window function corresponding to the second term in the above equation. We choose to make the small angle approximation, $\sin \theta \approx \theta$ and calculate the Fourier transform on the

flat sky. Finding

$$\begin{aligned}
\mathcal{F}(\Psi_{\text{Ruze}}) &= \mathcal{F} \left(\left(\frac{2\pi c}{\lambda} \right)^2 e^{-\overline{\rho^2}} \sum_{n=1}^{\infty} \frac{\overline{\rho^2}^n}{n \times n!} e^{-(\pi c \sin(\theta)/\lambda)^2/n} \right), \\
&\approx \left(\frac{2\pi c}{\lambda} \right)^2 e^{-\overline{\rho^2}} \sum_{n=1}^{\infty} \frac{\overline{\rho^2}^n}{n \times n!} \left(\mathcal{F}(e^{-\frac{\pi^2 c^2}{\lambda^2} \frac{x^2}{n}}) \mathcal{F}(e^{-\frac{\pi^2 c^2}{\lambda^2} \frac{y^2}{n}}) \right), \\
&= \left(\frac{2\pi c}{\lambda} \right)^2 e^{-\overline{\rho^2}} \sum_{n=1}^{\infty} \frac{\overline{\rho^2}^n}{n \times n!} \frac{\lambda \sqrt{n}}{\pi c} e^{-\frac{n\lambda^2}{\pi^2 c^2} l_x^2} e^{-\frac{n\lambda^2}{\pi^2 c^2} l_y^2}, \\
&= \frac{4\pi c}{\lambda} e^{-\overline{\rho^2}} \sum_{n=1}^{\infty} \frac{\overline{\rho^2}^n}{\sqrt{n} \times n!} e^{-\frac{n\lambda^2}{\pi^2 c^2} l^2}.
\end{aligned} \tag{C.29}$$

The azimuthal symmetry should not come as a surprise. The Ruze contribution to the window function is then the squared modulus of this value. A translation of the Ruze envelope with respect to some Cartesian coordinate system will only result in a phase shift that can be easily accounted for. The window function that combines the main beam and the Ruze envelope is easily derived from Equations C.27 and C.29. It is the angular average of the sum of the squared moduli of all terms. Since the Ruze envelope is azimuthally symmetric by construction we expect all $b_{\ell m}$ to be zero where $m \neq 0$. The complete spherical harmonic decomposition therefore follows naturally from the above expression.

Appendix D

Photometric Calibration Using Planet Flux Estimates

Here we discuss general concepts and some details of the analysis pipeline used to calculate the flux densities presented in Section 3.7. This analysis has helped inform the absolute photometric calibration of the instrument.

D.1 Describing Flux

Luminosity represents an intrinsic property of a light source, whereas flux depends on the separation between the observer and the source, and is therefore an extrinsic property. While luminosity is measured in units of Watts, flux is measured in units of Watts/m². Another extrinsic quantity is called spectral flux density, this one accounts for variation in perceived power depending on the spectral bandpass of the receiver used to perform the measurement. It implies that an integral is required in order to find the total flux from an object. In astronomy, spectral flux density, also known as spectral irradiance, is commonly quoted using a non-SI unit called Jansky (Jy), with $1 \text{ Jy} = 10^{-26} \text{ W/m}^2/\text{Hz}$.

Instead of reporting spectral flux density, astronomers will often use quantities that describe the spectral radiance of planets, an intrinsic quantity, and expect others to derive the

corresponding flux density appropriate for their instrument.¹ The parameters referred to as brightness temperature and Rayleigh-Jeans temperature are frequently used for this purpose. The Rayleigh-Jeans temperature is often used at low frequencies or when the effective source temperature is great. Unlike brightness temperature, the Rayleigh-Jeans temperature is linearly related to spectral flux density. On the other hand, brightness temperature is naturally related to Planck's law for blackbody radiation.

Suppose a calibrated bolometer measures a peak signal, B , when pointed directly towards a source that is much brighter than its surrounding area on the sky. Let's also assume that the calibration is such that the detector timestream is in units of Jy/sr. The flux from that object, F_p , can then be found by the following expression

$$F_p = B\Omega_p, \quad (\text{D.1})$$

where Ω_p is the solid angle extended by the source. Note how B , the spectral radiance, differs from spectral flux density in that it does not incorporate the solid angle of the emitter.

The spectral radiance of a source with brightness temperature T_b as measured by a bolometer with band center frequency ν_c is

$$B(\nu_c, T_b) = \frac{2h\nu_c^3}{c^2} \frac{1}{e^{h\nu_c/kT_b} - 1}. \quad (\text{D.2})$$

The above equation is the well known Planck formula for a blackbody with temperature T_b . In the limit of high temperature or low frequency, the expression reduces to

$$B(\nu_c, T_{\text{RJ}}) \approx \frac{2\nu_c^2 k T_{\text{RJ}}}{c^2}, \quad (\text{D.3})$$

where T_{RJ} is the Rayleigh-Jeans temperature. If a reference quotes a brightness temperature

¹When reporting luminosity-related parameters, no single standard exists for experiments operating at mm-wavelengths. Unit conversion can be cumbersome; relying on accurate estimates of the spectral bandpasses of the detectors as well as the spectrum of the sources that are used for calibration.

or a Rayleigh-Jeans temperature, the converse can be found by setting

$$\frac{2h\nu_c^3}{c^2} \frac{1}{e^{h\nu_c/kT_b} - 1} = \frac{2\nu_c^2 k T_{\text{RJ}}}{c^2} \quad (\text{D.4})$$

and solving for either T_b or T_{RJ} .

The outer planets of our solar system can be approximated as Rayleigh-Jeans sources in the spectral range of the *Planck* HFI detectors, but unfortunately none of the planets are perfect blackbodies. This means that both the Rayleigh-Jeans temperature and the brightness temperature are frequency dependent. The choice between using Rayleigh-Jeans temperature or brightness temperature is therefore equally motivated and simply a matter of taste.

D.2 Estimating Flux from Peak Signal

The following equation plays a central role in antenna theory [200]

$$\lambda^2 = A_{\text{eff}} \Omega_b. \quad (\text{D.5})$$

For any element of a lossless optical system, the product of its effective area, A , and the beam solid angle, Ω_b , is constant and equal to the wavelength of the radiation squared. The $A\Omega$ product is often referred as the optical throughput or étunde. In antenna theory, a detector's gain is defined as the ratio of power emitted in a particular direction compared to a detector that emits its supplied power evenly over all 4π steradians

$$G(\theta, \phi) = \frac{P(\theta, \phi)}{P_{\text{avg}}/4\pi}. \quad (\text{D.6})$$

Under this definition, the gain of an isotropic radiator is unity.

The time-dependent signal from a bolometer, $s(t)$, can be expressed by the following

convolution

$$s(t) = (Z * P)(t), \quad (\text{D.7})$$

where Z is the time dependent transfer function and P is the instantaneous power dissipated in the receiver. Assuming an infinitely fast time response, a detector observing a point like blackbody head on will measure an instantaneous background removed power according to

$$\begin{aligned} P &= g \iint d\Omega d\nu \tilde{F}(\nu) \Psi(\theta, \phi) A_{\text{eff}}(\nu) B(\nu, T) \\ &= g \frac{1}{\Omega_b} \iint d\Omega d\nu N \lambda^2 \tilde{F}(\nu) \Psi(\theta, \phi) B(\nu, T) \\ &= g \frac{\Omega_p}{\Omega_b} \int d\nu N \lambda^2 \tilde{F}(\nu) B(\nu, T) \end{aligned} \quad (\text{D.8})$$

where g is the absolute gain, representing the efficiency with which incident radiation is converted into energy of different form within the detector, $\Psi(\theta, \phi)$ is the instrument beam, normalized to unity at peak, $\tilde{F}(\nu)$ is the normalized detector bandpass, Ω_p is the time varying planetary solid angle, Ω_b is the scanning beam solid angle, and $B(\nu, T)$ is Planck's blackbody function. The above description assumes that the spatial response, $\Psi(\theta, \phi)$, does not vary over the bandpass and adopts the relationship between the effective area, beam solid angle, wavelength, and effective number of radiation modes N [295, 296].

$$N \lambda^2 = A_{\text{eff}} \Omega_b \quad (\text{D.9})$$

The integral in Equation D.8 is performed over all frequencies and the solid angle subtended by the source. Assuming the source is very small compared to the resolution of the detector, the integral over solid angle simply produces the solid angle of the source, Ω_p .² Most real sources are not blackbodies, meaning that the brightness temperature is frequency dependent. This is indeed the case for planets. Unfortunately, we can only sparsely sample this frequency dependence with the HFI detectors. The exact definition of brightness temperature is

² $\Psi(\theta, \phi)$ is approximately constant and equal to one over the solid angle spanned by a point source.

somewhat ill-defined in the context of detectors with a finite bandpass. Assuming an impulse shaped transfer function the signal at a given point in time becomes³

$$s = g \frac{\Omega_p}{\Omega_b} \int d\nu F(\nu) B(\nu, T), \quad (\text{D.10})$$

where $F(\nu)$ has been redefined to incorporate $N\lambda^2$ from Equation D.8 since $N\lambda^2\tilde{F}(\nu)$ is the standard output from a measurement of a detector's spectral response. The lower frequency channels are single-moded, so that $N = 1$ for all frequencies, but the 545 and 857 GHz bands are multi-moded, having sacrificed resolution for increased optical throughput. This is also measured and incorporated into $F(\nu)$ during analysis of the *Planck* HFI spectral response [214].

For CMB experiments it is customary to convert raw detector units into CMB temperature equivalent units. This is done to first order by taking the temperature derivative of the signal evaluated at T_{CMB}

$$\begin{aligned} \frac{ds}{dT} \Big|_{T_{\text{CMB}}} &= \frac{d}{dT} \left(g \frac{\Omega_b}{\Omega_b} \int d\nu F(\nu) B(\nu, T) \right) \Big|_{T_{\text{CMB}}} \\ &= g \int d\nu F(\nu) \frac{dB(\nu, T)}{dT} \Big|_{T_{\text{CMB}}}. \end{aligned} \quad (\text{D.11})$$

In this case the instrument response is integrated over the whole sphere giving the detector beam solid angle and not the planet solid angle. Higher order terms are not required as $B(\nu, T)$ is to a very good approximation a linear function of T for small values of ΔT relative to the overall signal amplitude. By finding the ratio of the expressions described in Equations D.10 and D.11 we can convert our results to CMB temperature. The corresponding CMB temperature change from a planet, ΔT_p , is therefore

$$\Delta T_p = \frac{g \frac{\Omega_p}{\Omega_b} \int d\nu F(\nu) B(\nu, T)}{g \int d\nu F(\nu) \frac{dB(\nu, T)}{dT} \Big|_{T_{\text{CMB}}}}, \quad (\text{D.12})$$

³Or a proper deconvolution of the detector time-response function. The assumption is also valid if the window size is small compared to the time it takes $P(t)$ to vary considerably.

where

$$\frac{dB}{dT}(\nu, T) = \frac{2\nu^2 k}{c^2} \left(\frac{h\nu}{kT}\right)^2 \frac{e^{\frac{h\nu}{kT}}}{(e^{\frac{h\nu}{kT}} - 1)^2} \quad (\text{D.13})$$

$$= \frac{2\nu^2 k}{c^2} \frac{x^2 e^x}{(e^x - 1)^2}, \quad (\text{D.14})$$

with $x = \frac{h\nu}{kT}$. Using this, the functional form for ΔT_p can be written out as

$$\begin{aligned} \Delta T_p &= \frac{\frac{\Omega_p}{\Omega_b} \int d\nu F(\nu) B(\nu, T)}{\int d\nu F(\nu) \frac{dB(\nu, T)}{dT} |_{T_{\text{CMB}}}} \\ &= \frac{\Omega_p}{\Omega_b} \frac{\int d\nu F(\nu) \frac{2h\nu^3}{c^2} / (e^{\frac{h\nu}{kT}} - 1)}{\int d\nu F(\nu) \frac{2\nu^2 k}{c^2} \left(\frac{h\nu}{kT_{\text{CMB}}}\right)^2 e^{\frac{h\nu}{kT_{\text{CMB}}}} / (e^{\frac{h\nu}{kT_{\text{CMB}}}} - 1)^2} \\ &= \frac{kT_{\text{CMB}}^2}{h} \frac{\Omega_p}{\Omega_b} \frac{\int d\nu F(\nu) \nu^3 / (e^{\frac{h\nu}{kT}} - 1)}{\int d\nu F(\nu) \nu^4 e^{\frac{h\nu}{kT_{\text{CMB}}}} / (e^{\frac{h\nu}{kT_{\text{CMB}}}} - 1)^2}, \end{aligned} \quad (\text{D.15})$$

where $T_{\text{CMB}} \equiv 2.725$ K. This equation can be used to calculate the expected value of peak temperature, ΔT_p , if the brightness temperature is known. Conversely, with a measurement of some peak temperature one can solve for the constant T that best fits measurements. This involves numerically solving a transcendental equation.

In the Rayleigh-Jeans limit, where $h\nu \ll kT$, the above equation can be approximated by

$$\Delta T_p = T \left(\frac{k^2 T_{\text{CMB}}^2}{h^2} \frac{\Omega_p}{\Omega_b} \right) \frac{\int d\nu F(\nu) \nu^2}{\int d\nu F(\nu) \nu^4 e^{\frac{h\nu}{kT_{\text{CMB}}}} / (e^{\frac{h\nu}{kT_{\text{CMB}}}} - 1)^2}. \quad (\text{D.16})$$

In that case the brightness temperature is linearly dependent on parameters such as Ω , Ω_p , and ΔT_p . This approximation is more accurate for large brightness temperature and low frequencies.

Due to the finite bandpass of the HFI detectors these absolute photometric calibrations are linked to the spectral shape of the calibrators. Photometric analysis of sources that have a different spectral shape could therefore be inaccurate. We combat this by applying

color corrections, see Section D.3. The color correction is a simple scaling factor that can be applied to the LHS of Equation D.15 before solving for T .

D.2.1 Estimating Photometric Parameters

The above section shows how we can derive the brightness temperature from estimates of the beam solid angle, Ω_b , the planet solid angle, Ω_p , the detector spectral response, $F(\nu)$, and the peak signal from the planet, ΔT_p .

The planet solid angles, Ω_p , are estimated from the JPL ephemerides software [297] and corrected for planet oblateness using methods described in [246]. The signal timestream is used to estimate the time at which each channel observes the planet head on. This corresponds to the time at which the peak signal is observed. This time is then input to the time-dependent function $\Omega_p(t)$. For this analysis, we assume 0.1% fractional error in the planet solid angle estimate.

The scanning beam solid angle, Ω_p , is based on the official DX11 hybrid B-Spline beam with error estimates from the corresponding simulation pipeline [229]. The band-average fractional scanning beam errors presented in the 2013 papers corresponded to 0.5% at 100 GHz and 0.1% at all other HFI frequencies [207]. The updated DX11 hybrid B-Spline beam will quote a smaller error. In this analysis we incorporate a conservative 0.3% scanning beam error at all frequencies.

The detector spectral response, $F(\nu)$, is based on the official spectra, discussed in great detail in [214]. We incorporate the color correction error described therein. The absolute calibration used for the 2014 papers is presented in [226]. In this analysis we assume the absolute calibration error corresponds to 0.1% at 100–217 GHz, 1% at 353 GHz, and 5% at 545 and 857 GHz.

Perhaps the most poorly described parameter in the above discussion is ΔT_p . Under perfect conditions, infinite sampling of the sky, and with optimal deconvolution of the detector time response, the scanning of a point source will generate the underlying scanning beam solid

angle. Assuming the beam is Gaussian in shape, the maximum of that Gaussian, as registered in the timestream, corresponds to the parameter ΔT_p as described above. Unfortunately, non idealities such as noise, astrophysical background, and incomplete coverage make the estimate of the peak signal from a planet observation non-trivial.

For the results presented in Section 3.7, the peak signals for all of the planet observations are based on a Gauss-Hermite fit and informed by the time-domain simulations. The peak signal estimate is based on a Gauss-Hermite decomposition (using $N_{\max} = 9$) of the signal timestream. The corresponding Gauss-Hermite beam is then reconstructed from which the peak signal is extracted.

Monte-Carlo simulations have been run for all of 22 planet observations. A known input beam is used to generate a signal timestream. For every realization, the peak amplitude is varied by a known amount which is sampled from a Gaussian distribution with unity mean and 2% standard deviation. The input amplitude is consistent with the models discussed in Section 3.7. The derived peak signal is then compared to the input to estimate the reconstruction bias and error. The error and bias depend on the observation, especially on the signal to noise ratio and the angular coverage in the region of the planet.

We assume all estimated errors are representative of Gaussian random variates. We use these estimates to define a Monte-Carlo brightness temperature analysis. Each realization draws from these Gaussian distributions. This allows us to estimate the distribution of errors in measurements of brightness temperature. We assume the errors are uncorrelated between detectors. The errors from this analysis are shown in Figures 3.19–3.23.

D.2.2 Defining Brightness Temperature – Subtle Differences

Some difficulties arise when the brightness temperature as defined in this document gets compared to model predictions which generate brightness temperatures as a function of frequency. This is due to the finite bandpass of the detectors. The above definitions creates a

parameter that is constant over the spectral bandpass of the HFI detectors and allows us to calculate the power incident on our detectors, see Equation D.10.

In reality, the power that is dissipated in the detectors per unit frequency is not constant over the detector bandpass. On the other hand, the total power that is dissipated in the detectors is determined by the spectral bandpass, the nature of the source, and the orientation of our receiver. As bolometers only measure the total power, this allows for some flexibility when it comes defining parameters that describe the properties of our source such as brightness temperature. The brightness temperature that our detectors measure does not need to be defined according to Equation D.10, instead we could write

$$s = g \frac{\Omega_p}{\Omega_b} \int d\nu F(\nu) B(\nu, T(\nu)) \quad (\text{D.17})$$

which implies that the brightness temperature is a function of frequency. We could also write

$$s = g \frac{\Omega_p}{\Omega_b} B(\nu_c, T(\nu_c)) \int d\nu F(\nu). \quad (\text{D.18})$$

Equation D.17 implies that the brightness temperature is frequency dependent. This is a natural way to define brightness temperature since the spectral energy density of most sources cannot be described by a single blackbody temperature. Unfortunately, the HFI detectors cannot probe this frequency dependence. Instead, we could choose to probe the brightness temperature for the band center according to Equation D.18. This has the benefit that the output can be compared directly to model predictions at the band center frequency. The only drawback is that the representation fails to capture the fact that the spectral energy density changes with frequency.

For consistency, we use Equation D.18 when estimating brightness temperature and not Equation D.10.⁴ Therefore, all brightness temperature estimates are reported at the formal *Planck* HFI band centers, corresponding to 100, 143, 217, 353, 545, and 857 GHz. For the

⁴This also means that the derivation following Equation D.10 need to be slightly modified.

models that we have considered, we find that this definition of brightness temperature can differ from the signal predicted by Equation D.18 by a non-negligible amount.

D.3 Color Corrections

Most of the *Planck* HFI channels are calibrated to the CMB orbital dipole [226], which has a well known spectral shape. In the case of the sub-mm channels, 545 and 857 GHz, the absolute calibration was originally based on FIRAS maps [298]. The FIRAS calibration assumes that the spectral radiance of the signal follows

$$I^{FIRAS}(\nu) \propto \nu^{-1}. \quad (\text{D.19})$$

However, as discussed in [224], the FIRAS calibrations was found to be inconsistent with predictions of models on planet flux densities. Because of this, the sub-mm channels were recalibrated so that the derived flux would agree with predictions from the ESA planet flux models. Despite the recalibration, the sub-mm channels are calibrated for sources with the spectral shape described by Equation D.19. Needless to say, not all sources on the sky have this spectral shape. This is remedied by implementing a color correction which is dependent on the assumed spectral shape of the source.

Similarly, the first order color correction to the CMB channels (100–353 GHz) is brought about by the difference in spectral shape of the 2.73 K blackbody temperature derivative compared to other astrophysical sources. This includes the planets, which to a good approximation have a spectral index $\alpha = 2$ over the *Planck* HFI frequency range.

In the case of planet brightness temperature, the color corrections for the two scenarios

Table D.1: The band average color corrections appropriate for analysis of planet flux densities within the framework described above. These numbers are derived using Equations D.20 and the band average spectral response. Further detail, including error estimates, found in [214].

Band	\mathcal{C}^{RJ}	$\mathcal{C}^{\text{Mars}}$	$\mathcal{C}^{\text{Jupiter}}$	$\mathcal{C}^{\text{Saturn}}$	$\mathcal{C}^{\text{Uranus}}$	$\mathcal{C}^{\text{Neptune}}$
100 GHz	0.961	0.961	0.962	0.963	0.969	0.974
143 GHz	1.006	1.007	1.008	1.008	1.012	1.013
217 GHz	0.928	0.936	0.938	0.953	0.949	0.962
353 GHz	0.931	0.934	0.939	0.941	0.949	0.937
545 GHz	0.935	0.936	0.929	0.818	0.951	0.964
857 GHz	0.980	0.983	0.999	0.999	0.992	0.998

are:

$$\begin{aligned} \mathcal{C}_{\text{CMB}} &= \frac{\int d\nu F(\nu) B(T, \nu) / B(T_c, \nu_c)}{\int d\nu F(\nu) \frac{dB(T_{\text{CMB}}, \nu)}{dT} / \frac{dB(T_{\text{CMB}}, \nu_c)}{dT}}, \\ \mathcal{C}_{\text{FIRAS}} &= \frac{\int d\nu F(\nu) B(T, \nu) / B(T_c, \nu_c)}{\int d\nu F(\nu) I^{\text{FIRAS}}(\nu) / I^{\text{FIRAS}}(\nu_c)}, \end{aligned} \quad (\text{D.20})$$

where B is the Planck blackbody function, with the frequency dependent planet brightness temperature T as input, and I^{FIRAS} represents the power law assumed in the sub-millimeter calibration to the FIRAS data. Here ν_c and T_c represent the band center frequency, the first frequency moment weighted by the spectral response, and the corresponding brightness temperature. The temperature derivative of the blackbody function is described in Equation D.14.

The color correction is a unitless number that corrects for the difference in spectral shape between the calibration source and the observed source. These color corrections can be applied to Equation D.15 before solving for the brightness temperature. These and similar procedures are described in [214].

Under the Rayleigh-Jeans assumption the intensity of the planets, I_{planet} can be written as

$$I_{\text{planet}} \approx \frac{2kT}{c^2} \nu^2. \quad (\text{D.21})$$

The first order color correction can then be obtained by assuming that the planets are Rayleigh-Jeans sources. Equations D.20 then take the approximate form

$$\begin{aligned}\mathcal{C}_{\text{CMB}} &= \frac{\int d\nu F(\nu) \nu^2 / \nu_c^2}{\int d\nu F(\nu) \frac{dB(T_{\text{CMB}}, \nu)}{dT} / \frac{dB(T_{\text{CMB}}, \nu_c)}{dT}}, \\ \mathcal{C}_{\text{FIRAS}} &= \frac{\int d\nu F(\nu) \nu^2 / \nu_c^2}{\int d\nu F(\nu) \nu_c / \nu}.\end{aligned}\quad (\text{D.22})$$

Table D.1 highlights the color corrections appropriate for analysis of the planet flux. Clearly, the Rayleigh-Jeans approximation is not valid for all of the planets. The columns represent the band-average color corrections for each of the five outer planets based on estimates of planet spectra from ESA models.

D.4 Measured Planet Fluxes

The following tables show the band-averaged results of the *Planck* HFI planet flux measurements. Interpretations of these results are described in Section 3.7.

Table D.2: Band averaged properties of Mars and Jupiter observations.

Unix time	Ω_{Planet}	Peak Value	I	T_B
	[arcsec 2]	[K_{CMB}]	[Jy]	[K]
Mars 1				
100 GHz	1256289161	43.9 ± 0.1	0.0297 ± 0.0005	61.9 ± 0.6
143 GHz	1256753797	47.0 ± 0.2	0.074 ± 0.001	140 ± 1
217 GHz	1256441015	44.9 ± 0.2	0.298 ± 0.007	303 ± 4
353 GHz	1256545814	45.61 ± 0.01	1.48 ± 0.03	832 ± 7
545 GHz	1256625171	46.15 ± 0.01	980 ± 90	1950 ± 20
857 GHz	1256625892	46.15 ± 0.01	2440 ± 50	4700 ± 100
Mars 2				
100 GHz	1271303236	52.5 ± 0.1	0.0332 ± 0.0007	69 ± 1
143 GHz	1270959570	56.0 ± 0.2	0.082 ± 0.001	155 ± 1
217 GHz	1271187753	53.6 ± 0.2	0.326 ± 0.006	332 ± 7
353 GHz	1271109495	54.47 ± 0.01	1.62 ± 0.03	910 ± 8
545 GHz	1271050444	55.08 ± 0.02	1000 ± 100	2110 ± 30
857 GHz	1271049935	55.09 ± 0.01	2800 ± 100	5420 ± 90
Mars 3				
100 GHz	1324421229	52.7 ± 0.1	0.0354 ± 0.0007	73 ± 1
143 GHz	1324748380	56.1 ± 0.2	0.090 ± 0.001	169 ± 1
217 GHz	1324532241	53.8 ± 0.2	0.352 ± 0.007	358 ± 7
353 GHz	1324607842	54.67 ± 0.01	1.77 ± 0.04	995 ± 9
545 GHz	1324664652	55.27 ± 0.01	1100 ± 100	2300 ± 30
857 GHz	1324664723	55.27 ± 0.01	3120 ± 80	6000 ± 100
Jupiter 1				
100 GHz	1256804823	1278 ± 1	0.759 ± 0.007	1580 ± 10
143 GHz	1256534511	1304 ± 2	1.79 ± 0.04	3360 ± 30
217 GHz	1256713680	1287 ± 2	7.0 ± 0.2	7100 ± 100
353 GHz	1256652527	1293.2 ± 0.1	33 ± 1	18600 ± 700
545 GHz	1256606080	1297.7 ± 0.1	17000 ± 1000	34600 ± 700
857 GHz	1256605881	1297.7 ± 0.1	44000 ± 2000	87200 ± 700

Table D.3: Band averaged properties of the Jupiter observations.

Unix time	Ω_{Planet}	Peak Value	I	T_{B}
	[arcsec 2]	[K_{CMB}]	[Jy]	[K]
Jupiter 2				
100 GHz	1278225694	1292 ± 1	0.772 ± 0.007	1600 ± 10
143 GHz	1278579274	1327 ± 2	1.83 ± 0.04	3430 ± 40
217 GHz	1278345800	1304 ± 2	7.1 ± 0.1	7200 ± 100
353 GHz	1278425937	1312.5 ± 0.1	33 ± 1	18900 ± 700
545 GHz	1278485682	1318.4 ± 0.1	17900 ± 800	35000 ± 1000
857 GHz	1278486082	1318.5 ± 0.1	48000 ± 7000	90000 ± 10000
Jupiter 3				
100 GHz	1292133769	1251 ± 1	0.742 ± 0.006	1540 ± 10
143 GHz	1291744072	1289 ± 3	1.77 ± 0.04	3320 ± 30
217 GHz	1292003329	1264 ± 3	6.9 ± 0.1	7000 ± 100
353 GHz	1291915176	1272.6 ± 0.1	32 ± 1	18400 ± 600
545 GHz	1291848252	1279.2 ± 0.2	17000 ± 1000	34500 ± 700
857 GHz	1291847952	1279.3 ± 0.1	49000 ± 6000	90000 ± 10000
Jupiter 4				
100 GHz	1312417985	1232.91 ± 0.02	0.741 ± 0.006	1534 ± 9
143 GHz	1312662015	1255 ± 1	1.73 ± 0.04	3260 ± 40
217 GHz	1312506435	1241 ± 1	6.9 ± 0.2	7000 ± 100
353 GHz	1312558365	1245.86 ± 0.09	32 ± 1	18100 ± 600
545 GHz	1312597508	1249.5 ± 0.1	16400 ± 600	34300 ± 100
857 GHz	1312597688	1249.5 ± 0.1	44000 ± 3000	86000 ± 3000
Jupiter 5				
100 GHz	N/A	N/A	N/A	N/A
143 GHz	1326219900	1297 ± 2	1.78 ± 0.04	3360 ± 30
217 GHz	1326403694	1279.0 ± 0.1	6.9 ± 0.1	7100 ± 100
353 GHz	1326353402	1284.1 ± 0.1	33 ± 1	18600 ± 600
545 GHz	1326301018	1289.4 ± 0.1	17000 ± 1000	34000 ± 1000
857 GHz	1326300638	1289.4 ± 0.1	49000 ± 7000	90000 ± 10000

Table D.4: Band averaged properties of the Saturn observations.

Unix time	Ω_{Planet}	Peak Value	I	T_B
	[arcsec 2]	[K_{CMB}]	[Jy]	[K]
Saturn 1				
100 GHz 1262597043	228.1 ± 0.1	0.112 ± 0.001	234 ± 1	144.8 ± 0.8
143 GHz 1262915364	231.1 ± 0.2	0.260 ± 0.006	491 ± 3	147 ± 1
217 GHz 1262701476	229.1 ± 0.2	1.01 ± 0.02	1050 ± 10	140 ± 1
353 GHz 1262771326	229.79 ± 0.01	4.7 ± 0.1	2660 ± 20	137 ± 1
545 GHz 1262823882	230.29 ± 0.01	2400 ± 100	4200 ± 100	98 ± 2
857 GHz 1262823612	230.29 ± 0.01	5900 ± 70	11500 ± 200	113 ± 2
Saturn 2				
100 GHz 1276766784	219.8 ± 0.1	0.110 ± 0.001	230 ± 1	147 ± 1
143 GHz 1276382297	223.3 ± 0.2	0.257 ± 0.006	483 ± 3	149 ± 1
217 GHz 1276638930	221.0 ± 0.2	0.99 ± 0.03	1020 ± 10	141 ± 1
353 GHz 1276553040	221.82 ± 0.01	4.5 ± 0.1	2560 ± 30	136 ± 1
545 GHz 1276485989	222.44 ± 0.02	2400 ± 200	4170 ± 60	100 ± 1
857 GHz 1276485689	222.44 ± 0.01	5300 ± 200	10360 ± 70	107.0 ± 0.6
Saturn 3				
100 GHz 1295357862	224.8 ± 0.1	0.1097 ± 0.0009	228 ± 1	143 ± 1
143 GHz 1295630470	227.3 ± 0.1	0.257 ± 0.006	483 ± 4	147 ± 1
217 GHz 1295449103	225.6 ± 0.2	1.01 ± 0.02	1047 ± 7	141.6 ± 0.9
353 GHz 1295511028	226.217 ± 0.009	4.7 ± 0.1	2670 ± 20	139 ± 1
545 GHz 1295557771	226.64 ± 0.01	2410 ± 90	4420 ± 20	103.5 ± 0.4
857 GHz 1295558411	226.65 ± 0.01	6000 ± 200	12000 ± 600	119 ± 5
Saturn 4				
100 GHz 1309789178	212.4 ± 0.1	0.1053 ± 0.0009	219 ± 1	145.2 ± 0.9
143 GHz 1309435409	215.5 ± 0.2	0.247 ± 0.006	465 ± 4	149 ± 1
217 GHz 1309673017	213.4 ± 0.2	0.94 ± 0.02	980 ± 10	140 ± 2
353 GHz 1309593546	214.17 ± 0.01	4.4 ± 0.1	2530 ± 30	139 ± 1
545 GHz 1309532233	214.70 ± 0.01	2400 ± 200	4130 ± 90	102 ± 2
857 GHz 1309531334	214.71 ± 0.01	5700 ± 100	11400 ± 400	119 ± 3

Table D.5: Band averaged properties of the Uranus observations.

Unix time	Ω_{Planet} [arcsec 2]	Peak Value [K_{CMB}]	I [Jy]	T_{B} [K]
Uranus 1				
100 GHz 1260424458	9.536 ± 0.002	0.0039 ± 0.0001	8.3 ± 0.2	123 ± 3
143 GHz 1260073106	9.603 ± 0.005	0.0077 ± 0.0002	14.6 ± 0.5	106 ± 4
217 GHz 1260306288	9.558 ± 0.005	0.028 ± 0.001	29.5 ± 0.6	95 ± 2
353 GHz 1260226976	9.5738 ± 0.0003	0.114 ± 0.002	65 ± 1	83 ± 1
545 GHz 1260165303	9.5856 ± 0.0004	59 ± 4	119 ± 4	70 ± 1
857 GHz 1260163194	9.5860 ± 0.0002	122 ± 5	236 ± 1	64.9 ± 0.3
Uranus 2				
100 GHz 1277961137	9.623 ± 0.003	0.0036 ± 0.0001	7.6 ± 0.2	113 ± 4
143 GHz 1278298015	9.686 ± 0.005	0.0077 ± 0.0001	14.7 ± 0.2	106 ± 1
217 GHz 1278072634	9.644 ± 0.005	0.0283 ± 0.0009	29.2 ± 0.7	94 ± 2
353 GHz 1278148270	9.6588 ± 0.0002	0.112 ± 0.006	63 ± 4	81 ± 4
545 GHz 1278205148	9.6694 ± 0.0001	62 ± 5	124.7 ± 0.5	72.4 ± 0.2
857 GHz 1278206729	9.6697 ± 0.0001	123 ± 1	243 ± 6	65 ± 1
Uranus 3				
100 GHz 1292414071	9.529 ± 0.003	0.0039 ± 0.0001	8.2 ± 0.2	122 ± 3
143 GHz 1292062524	9.596 ± 0.005	0.0079 ± 0.0002	15.0 ± 0.2	109 ± 1
217 GHz 1292299959	9.551 ± 0.005	0.028 ± 0.002	28 ± 1	94 ± 5
353 GHz 1292219674	9.5667 ± 0.0004	0.112 ± 0.007	64 ± 4	82 ± 5
545 GHz 1292158338	9.5784 ± 0.0004	59 ± 5	120 ± 1	70.9 ± 0.9
857 GHz 1292158138	9.5785 ± 0.0003	118 ± 6	232 ± 7	64 ± 1
Uranus 4				
100 GHz 1309934421	9.652 ± 0.002	0.00396 ± 0.00009	8.3 ± 0.1	121 ± 2
143 GHz 1310260667	9.713 ± 0.004	0.0081 ± 0.0002	15.3 ± 0.1	110 ± 1
217 GHz 1310042675	9.673 ± 0.004	0.029 ± 0.001	30 ± 1	97 ± 3
353 GHz 1310116967	9.6870 ± 0.0002	0.112 ± 0.005	64 ± 1	81 ± 2
545 GHz 1310174518	9.6978 ± 0.0003	61 ± 7	123 ± 3	71 ± 1
857 GHz 1310173569	9.6976 ± 0.0002	122 ± 6	236 ± 3	64.3 ± 0.7

Table D.6: Band averaged properties of the Neptune observations.

Unix time	Ω_{Planet} [arcsec 2]	Peak Value [K_{CMB}]	I [Jy]	T_{B} [K]
Neptune 1				
100 GHz 1257362717	4.0655 ± 0.0006	0.00167 ± 0.00006	3.5 ± 0.1	122 ± 5
143 GHz 1257087407	4.080 ± 0.001	0.0034 ± 0.0001	6.5 ± 0.2	112 ± 4
217 GHz 1257260311	4.0710 ± 0.0009	0.0118 ± 0.0005	12.4 ± 0.8	94 ± 6
353 GHz 1257209543	4.07376 ± 0.00008	0.047 ± 0.001	26.6 ± 0.9	80 ± 2
545 GHz 1257158872	4.0764 ± 0.0001	25 ± 2	52 ± 2	72 ± 2
857 GHz 1257158412	4.0765 ± 0.0001	50 ± 1	99 ± 3	64 ± 1
Neptune 2				
100 GHz 1274153625	4.0077 ± 0.0008	0.0016 ± 0.0001	3.5 ± 0.2	125 ± 7
143 GHz 1274495091	4.025 ± 0.001	0.0033 ± 0.0001	6.3 ± 0.1	109 ± 2
217 GHz 1274263215	4.013 ± 0.001	0.0116 ± 0.0006	12.1 ± 0.5	93 ± 3
353 GHz 1274342572	4.0177 ± 0.0001	0.045 ± 0.002	25.5 ± 0.8	78 ± 2
545 GHz 1274402518	4.0208 ± 0.0001	23 ± 2	48 ± 1	68 ± 1
857 GHz 1274402277	N/A	N/A	N/A	N/A
Neptune 3				
100 GHz 1289100641	4.0689 ± 0.0007	0.0015 ± 0.0001	3.3 ± 0.2	115 ± 9
143 GHz 1288822732	4.083 ± 0.001	0.0031 ± 0.0001	5.9 ± 0.2	102 ± 4
217 GHz 1289011048	4.073 ± 0.001	0.0115 ± 0.0006	12.0 ± 0.4	92 ± 3
353 GHz 1288949782	4.0770 ± 0.0002	0.049 ± 0.005	28 ± 3	84 ± 9
545 GHz 1288896433	4.07989 ± 0.00008	24 ± 2	50.1 ± 0.1	69.5 ± 0.2
857 GHz 1288896853	4.07986 ± 0.00006	52 ± 3	104 ± 5	66 ± 2
Neptune 4				
100 GHz 1305932952	4.0114 ± 0.0007	0.00149 ± 0.00008	3.1 ± 0.1	110 ± 6
143 GHz 1306289285	4.030 ± 0.001	0.0031 ± 0.0001	5.9 ± 0.2	103 ± 4
217 GHz 1306051735	4.017 ± 0.001	0.0117 ± 0.0009	12.3 ± 0.7	95 ± 5
353 GHz 1306123564	4.0214 ± 0.0002	0.046 ± 0.002	25.8 ± 0.9	79 ± 2
545 GHz 1306188841	4.0249 ± 0.0002	23 ± 2	48.3 ± 0.3	68.2 ± 0.4
857 GHz 1306190022	4.0249 ± 0.0001	54 ± 3	107 ± 4	68 ± 2

Bibliography

- [1] G. van Rossum. Python - web site. <https://www.python.org/>, August 2014.
- [2] MATLAB. *version 7.10.0 (R2010a)*. The MathWorks Inc., Natick, Massachusetts, 2010.
- [3] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing In Science & Engineering*, 9(3):90–95, 2007.
- [4] K. M. Górski, E. Hivon, A. J. Banday, et al. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *Astrophysical Journal*, 622:759–771, April 2005.
- [5] G. Chon, A. Challinor, S. Prunet, E. Hivon, and I. Szapudi. Fast estimation of polarization power spectra using correlation functions. *Monthly Notices of the RAS*, 350:914–926, May 2004.
- [6] A. Lewis, A. Challinor, and A. Lasenby. Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models. *Astrophysical Journal*, 538:473–476, August 2000.
- [7] H. S. Leavitt and E. C. Pickering. Periods of 25 Variable Stars in the Small Magellanic Cloud. *Harvard College Observatory Circular*, 173:1–3, March 1912.
- [8] E. Hubble. A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15:168–173, March 1929.
- [9] Anon. Finds Spiral Nebulae Are Stellar Systems. *New York Times*, November 1924.
- [10] E. Opik. An Estimate of the Distance of the Andromeda Nebula. *Astrophysical Journal*, 55:406–410, June 1922.
- [11] G. Lemaître. Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. *Annales de la Societe Scietifique de Bruxelles*, 47:49–59, 1927.
- [12] G. Lemaître. Expansion of the Universe, A Homogeneous Universe of Constant Mass and Increasing Radius Accounting for the Radial Velocity of Extra-galactic Nebulae. *Monthly Notices of the RAS*, 91:483–490, March 1931.

- [13] G. Lemaître. The Beginning of the World from the Point of View of Quantum Theory. *Nature*, 127:706, May 1931.
- [14] J. Farrell. *The Day Without Yesterday*. Thunder's Mouth Press, New York, NY, 2005.
- [15] R. A. Alpher, H. Bethe, and G. Gamow. The Origin of Chemical Elements. *Phys. Rev.*, 73:803–804, Apr 1948.
- [16] R. A. Alpher and R. Herman. Evolution of the Universe. *Nature*, 162:774–775, November 1948.
- [17] E. M. Burbidge, G. R. Burbidge, W. A. Fowler, and F. Hoyle. Synthesis of the Elements in Stars. *Rev. Mod. Phys.*, 29:547–650, Oct 1957.
- [18] P. G. Roll and D. T. Wilkinson. Cosmic Background Radiation at 3.2 cm-Support for Cosmic Black-Body Radiation. *Phys. Rev. Lett.*, 16:405–407, Mar 1966.
- [19] R. H. Dicke and P. J. Peebles. Gravitation and Space Science. *Space Science Reviews*, 4:419–460, June 1965.
- [20] P. J. E. Peebles, L. A. Page, and R. B. Partridge. *Finding the Big Bang*. Cambridge University Press, Cambridge, UK, 2009.
- [21] A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophysical Journal*, 142:419–421, July 1965.
- [22] R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson. Cosmic Black-Body Radiation. *Astrophysical Journal*, 142:414–419, July 1965.
- [23] A. McKellar. Evidence for the Molecular Origin of Some Hitherto Unidentified Interstellar Lines. *Publications of the ASP*, 52:187, June 1940.
- [24] G. Herzberg. *Molecular Spectra and Molecular Structure. Vol.1: Spectra of Diatomic Molecules*. 1950.
- [25] E. A. Ohm. Receiving System. *Bell System Technical Journal*, 40(4):1065–1094, 1961.
- [26] W. C. Jones. *A Measurement of the Temperature and Polarization Anisotropies in the Cosmic Microwave Background Radiation*. PhD thesis, California Institute of Technology, 2006.
- [27] D. J. Fixsen. The Temperature of the Cosmic Microwave Background. *Astrophysical Journal*, 707:916–920, December 2009.
- [28] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.
- [29] F. Zwicky. On the Masses of Nebulae and of Clusters of Nebulae. *Astrophysical Journal*, 86:217, October 1937.

- [30] V. C. Rubin, N. Thonnard, and W. K. Ford, Jr. Extended Rotation Curves of High-luminosity Spiral Galaxies. IV - Systematic Dynamical Properties, SA Through SC. *Astrophysical Journal, Letters*, 225:L107–L111, November 1978.
- [31] V. C. Rubin, W. K. J. Ford, and N. Thonnard. Rotational Properties of 21 SC Galaxies With a Large Range of Luminosities and Radii, from NGC 4605 ($R = 4\text{kpc}$) to UGC 2885 ($R = 122\text{ kpc}$). *Astrophysical Journal*, 238:471–487, June 1980.
- [32] J. P. Ostriker, P. J. E. Peebles, and A. Yahil. The Size and Mass of Galaxies, and the Mass of the Universe. *Astrophysical Journal, Letters*, 193:L1–L4, October 1974.
- [33] M. Markevitch, A. H. Gonzalez, D. Clowe, et al. Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657-56. *Astrophysical Journal*, 606:819–824, May 2004.
- [34] A. H. Guth. *The Inflationary Universe*. Basic Books, New York, NY, 1997.
- [35] S. Singh. *Big Bang: The Origin of the Universe*. Harper Perennial, New York, NY, 2005.
- [36] S. Weinberg. *The First Three Minutes*. Basic Books, New York, NY, 1993.
- [37] M. Colless, G. Dalton, S. Maddox, et al. The 2dF Galaxy Redshift Survey: spectra and redshifts. *Monthly Notices of the RAS*, 328:1039–1063, December 2001.
- [38] V. Springel, S. D. M. White, A. Jenkins, et al. Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature*, 435:629–636, June 2005.
- [39] Y. Rasera, P.-S. Corasaniti, J.-M. Alimi, et al. Cosmic-variance limited Baryon Acoustic Oscillations from the DEUS-FUR Λ CDM simulation. *ArXiv e-prints*, November 2013.
- [40] M. Colless. The 2df galaxy redshift survey - web site. <http://www2.aoa.gov.au/2dFGRS/>, February 2014.
- [41] P. Lubin, T. Villela, G. Epstein, and G. Smoot. A Map of the Cosmic Background Radiation at 3 Millimeters. *Astrophysical Journal, Letters*, 298:L1–L5, November 1985.
- [42] S. Hayakawa, T. Matsumoto, H. Matsuo, H. Murakami, and S. Sato. Cosmological implication of a new measurement of the submillimeter background radiation. *Publications of the ASJ*, 39:941–948, 1987.
- [43] J. C. Mather, E. S. Cheng, R. E. Eplee, Jr., et al. A preliminary measurement of the cosmic microwave background spectrum by the Cosmic Background Explorer (COBE) satellite. *Astrophysical Journal, Letters*, 354:L37–L40, May 1990.
- [44] D. J. Fixsen, E. S. Cheng, D. A. Cottingham, et al. Cosmic microwave background dipole spectrum measured by the COBE FIRAS instrument. *Astrophysical Journal*, 420:445–449, January 1994.

- [45] J. C. Mather, E. S. Cheng, D. A. Cottingham, et al. Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument. *Astrophysical Journal*, 420:439–444, January 1994.
- [46] A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354:769–822, 1916.
- [47] A. Friedmann. Über die Krümmung des Raumes. *Zeitschrift fur Physik*, 10:377–386, 1922.
- [48] A. E. Lange, P. A. Ade, J. J. Bock, et al. Cosmological parameters from the first results of Boomerang. *Physical Review D*, 63(4):042001, February 2001.
- [49] Planck Collaboration, P. A. R. Ade, N. Aghanim, et al. Planck 2013 Results. XVI. Cosmological parameters. *ArXiv e-prints*, March 2013.
- [50] L. M. Krauss and M. S. Turner. The cosmological constant is back. *General Relativity and Gravitation*, 27:1137–1144, November 1995.
- [51] A. G. Riess, A. V. Filippenko, P. Challis, et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The Astronomical Journal*, 116(3):1009, 1998.
- [52] S. Perlmutter, G. Aldering, G. Goldhaber, et al. Measurements of Ω and Λ from 42 High-Redshift Supernovae. *The Astrophysical Journal*, 517(2):565, 1999.
- [53] S. Cole, W. J. Percival, J. A. Peacock, et al. The 2dF Galaxy Redshift Survey: power-spectrum analysis of the final data set and cosmological implications. *Monthly Notices of the RAS*, 362:505–534, September 2005.
- [54] D. J. Eisenstein, I. Zehavi, D. W. Hogg, et al. Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *Astrophysical Journal*, 633:560–574, November 2005.
- [55] A. H. Guth. Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems. *Phys. Rev. D*, 23(2):347–356, Jan 1981.
- [56] A. Albrecht and P. J. Steinhardt. Cosmology for Grand Unified Theories with Radiatively Induced Symmetry Breaking. *Physical Review Letters*, 48:1220–1223, April 1982.
- [57] A. D. Linde. A New Inflationary Universe Scenario: A Possible Solution of the Horizon, Flatness, Homogeneity, Isotropy and Primordial Monopole Problems. *Physics Letters B*, 108:389–393, February 1982.
- [58] A. A. Starobinsky. A New Type of Isotropic cosmological Models Without Singularity. *Physics Letters B*, 91(1):99 – 102, 1980.

[59] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1):1 – 29, 2012.

[60] CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC . *Physics Letters B*, 716(1):30 – 61, 2012.

[61] A. R. Liddle and D. H. Lyth. *Cosmological Inflation and Large-Scale Structure*. June 2000.

[62] W. Hu and M. White. Acoustic Signatures in the Cosmic Microwave Background. *Astrophysical Journal*, 471:30, November 1996.

[63] Y. B. Zel'dovich, V. G. Kurt, and R. A. Syunyaev. Recombination of Hydrogen in the Hot Model of the Universe. *Soviet Journal of Experimental and Theoretical Physics*, 28:146, January 1969.

[64] Planck collaboration. Planck 2013 Results. XV. CMB power spectra and likelihood. *ArXiv e-prints*, March 2013.

[65] Planck Collaboration. Planck 2013 Results. XXII. Constraints on inflation. *ArXiv e-prints*, March 2013.

[66] Planck Collaboration. Planck 2013 Results. XVII. Gravitational lensing by large-scale structure. *ArXiv e-prints*, March 2013.

[67] Planck Collaboration. Planck 2013 Results. XXVII. Doppler boosting of the CMB: Eppur si muove. *ArXiv e-prints*, March 2013.

[68] LAMBDA. Legacy archive for microwave background analysis - web site. <http://lambda.gsfc.nasa.gov/>, May 2014.

[69] W. Hu and M. White. CMB anisotropies: Total angular momentum method. *Physical Review D*, 56:596–615, July 1997.

[70] H. C. Chiang, P. A. R. Ade, D. Barkats, et al. Measurement of Cosmic Microwave Background Polarization Power Spectra from Two Years of BICEP Data. *The Astrophysical Journal*, 711(2):1123, 2010.

[71] QUIET Collaboration, D. Araujo, C. Bischoff, et al. Second Season QUIET Observations: Measurements of the Cosmic Microwave Background Polarization Power Spectrum at 95 GHz. *Astrophysical Journal*, 760:145, December 2012.

[72] T. Essinger-Hileman, J. W. Appel, J. A. Beal, et al. The Atacama B-Mode Search: CMB Polarimetry with Transition-Edge-Sensor Bolometers. In B. Young, B. Cabrera, and A. Miller, editor, *American Institute of Physics Conference Series*, volume 1185 of *American Institute of Physics Conference Series*, pages 494–497, December 2009.

- [73] H. T. Nguyen, J. Kovac, P. A. R. Ade, et al. BICEP2/SPUD: searching for inflation with degree scale polarimetry from the South Pole. volume 7020, page 70201F. SPIE, 2008.
- [74] M. Hedman. *The Princeton IQU Experiment and Constraints on the Polarization of the Cosmic Microwave Background at 90 GHz*. PhD thesis, Princeton University, 2002.
- [75] M. Zaldarriaga and U. Seljak. All-sky analysis of polarization in the microwave background. *Physical Review D*, 55:1830–1840, February 1997.
- [76] S. Weinberg. *Cosmology*. Oxford University Press, New York, NY, 2008.
- [77] R. K. Sachs and A. M. Wolfe. Perturbations of a Cosmological Model and Angular Variations of the Microwave Background. *Astrophysical Journal*, 147:73, January 1967.
- [78] J. Silk. Cosmic Black-Body Radiation and Galaxy Formation. *Astrophysical Journal*, 151:459, February 1968.
- [79] T. M. Crawford, K. K. Schaffer, S. Bhattacharya, et al. A Measurement of the Secondary-CMB and Millimeter-wave-foreground Bispectrum using 800 deg² of South Pole Telescope Data. *Astrophysical Journal*, 784:143, April 2014.
- [80] S. Naess, M. Hasselfield, J. McMahon, et al. The Atacama Cosmology Telescope: CMB Polarization at $200 < \ell < 9000$. *ArXiv e-prints*, May 2014.
- [81] M. L. Brown, P. Ade, J. Bock, et al. Improved Measurements of the Temperature and Polarization of the Cosmic Microwave Background from QUaD. *Astrophysical Journal*, 705:978–999, November 2009.
- [82] G. Hinshaw, D. Larson, E. Komatsu, et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *Astrophysical Journal, Supplement*, 208:19, October 2013.
- [83] C. L. Bennett, D. Larson, J. L. Weiland, et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. *Astrophysical Journal, Supplement*, 208:20, October 2013.
- [84] C. Bischoff, L. Hyatt, J. J. McMahon, et al. New Measurements of Fine-Scale CMB Polarization Power Spectra from CAPMAP at Both 40 and 90 GHz. *Astrophysical Journal*, 684:771–789, September 2008.
- [85] J. L. Sievers, C. Achermann, J. R. Bond, et al. Implications of the Cosmic Background Imager Polarization Data. *Astrophysical Journal*, 660:976–987, May 2007.
- [86] E. M. Leitch, J. M. Kovac, N. W. Halverson, et al. Degree Angular Scale Interferometer 3 Year Cosmic Microwave Background Polarization Results. *Astrophysical Journal*, 624:10–20, May 2005.
- [87] BICEP1 Collaboration, D. Barkats, R. Aikin, et al. Degree-Scale CMB Polarization Measurements from Three Years of BICEP1 Data. *ArXiv e-prints*, October 2013.

- [88] BICEP2 Collaboration, P. A. R. Ade, R. W. Aikin, et al. BICEP2 I: Detection Of B-mode Polarization at Degree Angular Scales. *ArXiv e-prints*, March 2014.
- [89] T. E. Montroy, P. A. R. Ade, J. J. Bock, et al. A Measurement of the CMB EE Spectrum from the 2003 Flight of BOOMERANG. *Astrophysical Journal*, 647:813–822, August 2006.
- [90] QUIET Collaboration, D. Araujo, C. Bischoff, et al. Second Season QUIET Observations: Measurements of the Cosmic Microwave Background Polarization Power Spectrum at 95 GHz. *Astrophysical Journal*, 760:145, December 2012.
- [91] D. Hanson, S. Hoover, A. Crites, et al. Detection of B-Mode Polarization in the Cosmic Microwave Background with Data from the South Pole Telescope. *Physical Review Letters*, 111(14):141301, October 2013.
- [92] The POLARBEAR Collaboration, P. A. R. Ade, Y. Akiba, et al. A Measurement of the Cosmic Microwave Background B-Mode Polarization Power Spectrum at Sub-Degree Scales with POLARBEAR. *ArXiv e-prints*, March 2014.
- [93] L. Knox and M. S. Turner. Detectability of Tensor Perturbations through Anisotropy of the Cosmic Background Radiation. *Physical Review Letters*, 73:3347–3350, December 1994.
- [94] L. Knox. Determination of Inflationary Observables by Cosmic Microwave Background Anisotropy Experiments. *Physical Review D*, 52:4307–4318, October 1995.
- [95] C. Cheung, A. L. Fitzpatrick, J. Kaplan, L. Senatore, and P. Creminelli. The Effective Field Theory of Inflation. *Journal of High Energy Physics*, 3:14–014, March 2008.
- [96] R. H. Brandenberger. Is the Spectrum of Gravitational Waves the "Holy Grail" of Inflation? *ArXiv e-prints*, April 2011.
- [97] S. Dodelson. *Modern cosmology*. Academic Press. Academic Press, Incorporated, 2003.
- [98] A. R. Liddle. An Introduction to Cosmological Inflation. In A. Masiero, G. Senjanovic, and A. Smirnov, editors, *High Energy Physics and Cosmology, 1998 Summer School*, page 260, 1999.
- [99] D. Baumann and H. V. Peiris. Cosmological Inflation: Theory and Observations. *Advanced Science Letters*, 2:105–120(16), 2009.
- [100] D. Baumann. Cosmology - Lecture Notes for Part III Mathematical Tripos, 2013. University of Cambridge.
- [101] R. Durrer. Cosmological Perturbation Theory. In K. Tamvakis, editor, *The Physics of the Early Universe*, volume 653 of *Lecture Notes in Physics*, Berlin Springer Verlag, page 31, 2005.
- [102] S. Hanany, M. D. Niemack, and L. Page. *CMB Telescopes and Optical Systems*, page 431. 2013.

- [103] A. D. Miller, R. Caldwell, M. J. Devlin, et al. A Measurement of the Angular Power Spectrum of the Cosmic Microwave Background from $\ell = 100$ to 400. *Astrophysical Journal, Letters*, 524:L1–L4, October 1999.
- [104] J. M. Kovac, E. M. Leitch, C. Pryke, et al. Detection of polarization in the cosmic microwave background using DASI. *Nature*, 420:772–787, December 2002.
- [105] S. Das, B. D. Sherwin, P. Aguirre, et al. Detection of the Power Spectrum of Cosmic Microwave Background Lensing by the Atacama Cosmology Telescope. *Physical Review Letters*, 107(2):021301, July 2011.
- [106] Polarbear Collaboration, P. A. R. Ade, Y. Akiba, et al. Gravitational Lensing of Cosmic Microwave Background Polarization. *ArXiv e-prints*, December 2013.
- [107] J. P. Kaufman, N. J. Miller, M. Shimon, et al. Self-Calibration of BICEP1 Three-Year Data and Constraints on Astrophysical Polarization Rotation. *ArXiv e-prints*, December 2013.
- [108] BICEP2 Collaboration, P. A. R Ade, R. W. Aikin, et al. BICEP2 II: Experiment and Three-Year Data Set. *ArXiv e-prints*, March 2014.
- [109] A. Kusaka, T. Essinger-Hileman, J. W. Appel, et al. Modulation of CMB polarization with a warm rapidly-rotating half-wave plate on the Atacama B-Mode Search (ABS) instrument. *ArXiv e-prints*, October 2013.
- [110] K. N. Abazajian, K. Arnold, J. Austermann, et al. Neutrino Physics from the Cosmic Microwave Background and Large Scale Structure. *ArXiv e-prints*, September 2013.
- [111] A. G. Riess, L. Macri, S. Casertano, et al. A 3% Solution: Determination of the Hubble Constant with the Hubble Space Telescope and Wide Field Camera 3. *The Astrophysical Journal*, 730:119, April 2011.
- [112] C. Blake, E. A. Kazin, F. Beutler, et al. The WiggleZ Dark Energy Survey: mapping the distance-redshift relation with baryon acoustic oscillations. *Monthly Notices of the RAS*, 418:1707–1724, December 2011.
- [113] B. Bassett and R. Hlozek. *Baryon Acoustic Oscillations*, page 246. 2010.
- [114] J. Beringer, J. F. Arguin, R. M. Barnett, et al. Review of particle physics. *Phys. Rev. D*, 86:010001, Jul 2012.
- [115] B. P. Crill, P. A. R. Ade, E. S. Battistelli, et al. SPIDER: A Balloon-borne Large-scale CMB Polarimeter. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7010 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, August 2008.
- [116] J. P. Filippini, P. A. R. Ade, M. Amiri, et al. SPIDER: a Balloon-Borne CMB Polarimeter for Large Angular Scales. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.

- [117] E. S. Battistelli, M. Amiri, B. Burger, et al. Functional Description of Read-out Electronics for Time-Domain Multiplexed Bolometers for Millimeter and Sub-millimeter Astronomy. *Journal of Low Temperature Physics*, 151:908–914, May 2008.
- [118] M. Dobbs, M. Halpern, K. D. Irwin, et al. Multiplexed Readout of CMB Polarimeters. *Journal of Physics: Conference Series*, 155(1):012004, 2009.
- [119] S. J. Benton, P. A. R. Ade, M. Amiri, et al. BLASTbus Electronics: General-Purpose Readout and Control for Balloon-Borne Experiments. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, June 2014.
- [120] J. A. Shariff, P. A. R. Ade, M. Amiri, et al. Pointing Control for the Spider Balloon-borne Telescope. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, June 2014.
- [121] J. Dunkley, A. Amblard, C. Baccigalupi, et al. Prospects for polarized foreground removal. In S. Dodelson, D. Baumann, A. Cooray, et al., editors, *American Institute of Physics Conference Series*, volume 1141 of *American Institute of Physics Conference Series*, pages 222–264, June 2009.
- [122] A. A. Fraisse, J.-A. C. Brown, G. Dobler, et al. Foreground Science Knowledge and Prospects. In S. Dodelson, D. Baumann, A. Cooray, et al., editors, *American Institute of Physics Conference Series*, volume 1141 of *American Institute of Physics Conference Series*, pages 265–310, June 2009.
- [123] A. A. Fraisse. *Journey to the Surface of Last Scattering*. PhD thesis, Princeton University, 2010.
- [124] A. A. Fraisse, P. A. R. Ade, M. Amiri, et al. SPIDER: Probing the Early Universe with a Suborbital Polarimeter. *J. Cosmology Astropart. Phys.*, 4:47, April 2013.
- [125] R. Flauger, J. C. Hill, and D. N. Spergel. Toward an Understanding of Foreground Emission in the BICEP2 Region. *ArXiv e-prints*, May 2014.
- [126] E. D. Kovetz and M. Kamionkowski. Cosmic Bandits: Exploration versus Exploitation in CMB B-Mode Experiments. *ArXiv e-prints*, August 2013.
- [127] A. Lange. Future Detectors Needs for CMB Observations. In *Far-IR, Sub-mm and MM Detector Technology Workshop*, 2002.
- [128] J. R. Pardo, J. Cernicharo, and E. Serabyn. Atmospheric transmission at microwaves (atm): an improved model for millimeter/submillimeter applications. *Antennas and Propagation, IEEE Transactions on*, 49(12):1683–1694, Dec 2001.
- [129] J. Delabrouille, M. Betoule, J.-B. Melin, et al. The pre-launch Planck Sky Model: a model of sky emission at submillimetre to centimetre wavelengths. *Astronomy and Astrophysics*, 553:A96, May 2013.

- [130] D. T. O'Dea, P. A. R. Ade, M. Amiri, et al. SPIDER Optimization. II. Optical, Magnetic, and Foreground Effects. *The Astrophysical Journal*, 738(1):63, 2011.
- [131] R. Hanbury Brown and R. Q. Twiss. Correlation Between Photons in Two Coherent Beams of Light. *Nature*, 177:27–29, January 1956.
- [132] J. M. Lamarre. Photon noise in photometric instruments at far-infrared and submillimeter wavelengths. *Appl. Opt.*, 25(6):870–876, Mar 1986.
- [133] M. Bowden, A. N. Taylor, K. M. Ganga, et al. Scientific optimization of a ground-based CMB polarization experiment. *Monthly Notices of the RAS*, 349:321–335, March 2004.
- [134] C. Armitage-Caplan, J. Dunkley, H. K. Eriksen, and C. Dickinson. Impact on the tensor-to-scalar ratio of incorrect Galactic foreground modelling. *Monthly Notices of the RAS*, 424:1914–1924, August 2012.
- [135] Planck Collaboration, P. A. R. Ade, N. Aghanim, et al. Planck 2013 Results. XII. Component separation. *ArXiv e-prints*, March 2013.
- [136] J. Dunkley, A. Amblard, C. Baccigalupi, et al. Prospects for polarized foreground removal. In S. Dodelson, D. Baumann, A. Cooray, et al., editors, *American Institute of Physics Conference Series*, volume 1141 of *American Institute of Physics Conference Series*, pages 222–264, June 2009.
- [137] M. J. Mortonson and U. Seljak. A joint analysis of Planck and BICEP2 B modes including dust polarization uncertainty. *ArXiv e-prints*, May 2014.
- [138] Planck Collaboration. Planck 2013 Results. VI. High Frequency Instrument data processing. *ArXiv e-prints*, March 2013.
- [139] J. Bock, A. Cooray, S. Hanany, et al. The Experimental Probe of Inflationary Cosmology (EPIC): A Mission Concept Study for NASA's Einstein Inflation Probe. *ArXiv e-prints*, May 2008.
- [140] The COrE Collaboration, C. Armitage-Caplan, M. Avillez, et al. COrE (Cosmic Origins Explorer) A White Paper. *ArXiv e-prints*, February 2011.
- [141] D. D. Gregory and W. E. Stepp. NASA's long duration balloon program: the last ten years and the next ten years. *Advances in Space Research*, 33:1608–1612, January 2004.
- [142] Columbia Scientific Balloon Facility. *Structural Requirements and Recommendations for Balloon Gondola Design*, April 2013.
- [143] J. D. Soler. *In Search of an Imprint of Magnetization in the Balloon-borne Observations of the Polarized Dust Emission from Molecular Clouds*. PhD thesis, University of Toronto, 2013.
- [144] S. J. Benton. *Mapping Submillimetre Polarization with BLASTPol*. PhD thesis, University of Toronto, 2014.

- [145] J. A. Shariff. *In preparation*. PhD thesis, University of Toronto, 2014.
- [146] N. N. Gandilo. *In preparation*. PhD thesis, University of Toronto, 2014.
- [147] S. Peterzen, S. Rabbia, A. Cardillo, I. Musso, and A. Leonardi. Development of a long duration balloon program in the polar regions - Svalbard, Norway. In B. Warmbein, editor, *European Rocket and Balloon Programmes and Related Research*, volume 530 of *ESA Special Publication*, pages 213–216, August 2003.
- [148] P. de Bernardis, S. Masi, and OLIMPO and LSPE Teams. Precision CMB measurements with long-duration stratospheric balloons: activities in the Arctic. In M. G. Burton, X. Cui, and N. F. H. Tothill, editors, *IAU Symposium*, volume 288 of *IAU Symposium*, pages 208–213, January 2013.
- [149] D. D. Gregory. Private communications, 2014.
- [150] J. E. Gudmundsson, P. A. R. Ade, M. Amiri, et al. Thermal Architecture for the SPIDER Flight Cryostat. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.
- [151] W. M. Kays and A. L. London. *Compact Heat Exchangers*. McGraw-Hill, New York, second edition, 1964.
- [152] F. P. Incropera, D. P. DeWitt, T. L. Bergman, and A. S. Lavine. *Fundamentals of Heat and Mass Transfer*. John Wiley & Sons, Hoboken, NJ, sixth edition, 2006.
- [153] S. Bapat. Performance Prediction of Multilayer Insulation. *Cryogenics*, 30(8):700–710, August 1990.
- [154] Y. Eyssa and O. Okasha. Thermodynamic Optimization of Thermal Radiation Shields for a Cryogenic Apparatus. *Cryogenics*, 18(5):305–307, May 1978.
- [155] M. C. Runyan and W. C. Jones. Thermal conductivity of thermally-isolating polymeric and composite structural support materials between 0.3 and 4 K. *Cryogenics*, 48:448–454, September 2008.
- [156] N. Selden, N. Gimelshein, S. Gimelshein, and A. Ketsdever. Analysis of accommodation coefficients of noble gases on aluminum surface with an experimental/computational method. *Physics of Fluids (1994-present)*, 21(7):–, 2009.
- [157] D. Steffensrud, J. Bahls, E. Christenson, et al. Efficiency of short heat exchangers for helium vapor cooling. *Review of Scientific Instruments*, 62(1):214–216, 1991.
- [158] H. Okamoto and D. Chen. A low-loss, ultrahigh vacuum compatible helium cryostat without liquid nitrogen shield. *Review of Scientific Instruments*, 72(2):1510–1513, 2001.
- [159] A. S. Rahlin. *In preparation*. PhD thesis, Princeton University, 2014.

- [160] J. W. Ekin. *Experimental techniques for low-temperature measurements*. Oxford University Press, New York, 2006.
- [161] T. M. Flynn. *Cryogenic Engineering*. CRC Press, New York, second edition, 2005.
- [162] L. Duband, A. Ravex, and J. Chaussy. Adsorption isotherms of helium on activated charcoal. *Cryogenics*, 27(7):397 – 400, 1987.
- [163] A. J. Kidnay and M. J. Hiza. Physical adsorption in cryogenic engineering. *Cryogenics*, 10(4):271 – 277, 1970.
- [164] M. Kumita, J. Ozaki, J. Kobayashi, et al. Characteristics of helium adsorption on microporous solids under cryogenic conditions. *Journal of Chemical Engineering of Japan*, 28(2):159–164, 1995.
- [165] K. Kaneko. Specific Intermolecular Structures of Gases Confined in Carbon Nanospace. *Carbon*, 38(2):287 – 303, 2000.
- [166] K. Murata, M. El-Merraoui, and K. Kaneko. A new Determination Method of Absolute Adsorption Isotherm of Supercritical Gases Under High Pressure with a Special Relevance to Density-functional Theory Study. *Journal of Chemical Physics*, 114:4196–4205, March 2001.
- [167] L. E. DeLong, O. G. Symko, and J. C. Wheatley. Continuously Operating He-4 Evaporation Refrigerator. *Review of Scientific Instruments*, 42(1):147–150, 1971.
- [168] J. Wrubel, G. Gabrielse, W.S. Kolthammer, et al. Pumped Helium System for Cooling Positron and Electron Traps to 1.2 K. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 640(1):232 – 240, 2011.
- [169] N. K. Das, J. Pradhan, Md. Z. A. Naser, et al. Design and Performance of a 4He-Evaporator at < 1.0 K. *Cryogenics*, (0):–, 2012.
- [170] Y. Fujiyoshi, T. Mizusaki, K. Morikawa, et al. Development of a Superfluid Helium Stage for High-Resolution Electron Microscopy. *Ultramicroscopy*, 38:241–251, 1991.
- [171] J. Talpe and V. Bekeris. Low Dead Time 1 K Cryostat. *Cryogenics*, 29(8):854 – 856, 1989.
- [172] A. Singsaas and G. Ahlers. Entropy of He II from 1.6 K to the λ Line. *Phys. Rev. B*, 29:4951–4960, May 1984.
- [173] S.-i. Koh. The Onset of Superfluidity in Capillary Flow of Liquid Helium 4. *ArXiv e-prints*, August 2008.
- [174] J. Lau, M. Benna, M. Devlin, S. Dicker, and L. Page. Experimental Tests and Modeling of the Optimal Orifice Size for a Closed Cycle 4He Sorption rRefrigerator. *Cryogenics*, 46(11):809 – 814, 2006.

[175] J. F. Allen and A. D. Misener. Flow of Liquid Helium II. *Nature*, 141:75, January 1938.

[176] L. Landau. Theory of the Superfluidity of Helium II. *Physical Review*, 60:356–358, August 1941.

[177] I. Amdur. Low Temperature Transport Properties of Gases. I. Helium. *The Journal of Chemical Physics*, 15(7), 1947.

[178] U. Schotte. He II phase separation with slits and porous plugs for space cryogenics. *Cryogenics*, 24(10):536 – 548, 1984.

[179] C. R. Lages, R. H. Torii, and D. B. DeBra. Evaporation of Superfluid Helium in a Capillary. *Cryogenics*, 35(1):31–33, 1995.

[180] K. W. Yoon, P. A. R. Ade, D. Barkats, et al. The Robinson Gravitational Wave Background Telescope (BICEP): a bolometric large angular scale CMB polarimeter. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6275 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2006.

[181] R. W. Aikin, P. A. Ade, S. Benton, et al. Optical performance of the BICEP2 Telescope at the South Pole. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.

[182] M. C. Runyan, P. A. R. Ade, M. Amiri, et al. Design and Performance of the SPIDER Instrument. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.

[183] A. Mager. Magnetic shields. *Magnetics, IEEE Transactions on*, 6(1):67–75, 1970.

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

[185] S. A. Bryan, P. A. R. Ade, M. Amiri, et al. Modeling and characterization of the SPIDER half-wave plate. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.

[186] B. Reichborn-Kjennerud, A. M. Aboobaker, P. Ade, et al. EBEX: a balloon-borne CMB polarization experiment. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.

[187] L. M. Fissel, P. A. R. Ade, F. E. Angilè, et al. The balloon-borne large-aperture submillimeter telescope for polarimetry: BLAST-Pol. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.

- [188] S. A. Bryan, T. E. Montroy, and J. E. Ruhl. Modeling dielectric half-wave plates for cosmic microwave background polarimetry using a Mueller matrix formalism. *Applied Optics*, 49:6313, November 2010.
- [189] T. Essinger-Hileman. Transfer matrix for treating stratified media including birefringent crystals. *Applied Optics*, 52:212, January 2013.
- [190] L. Moncelsi, P. A. R. Ade, F. E. Angil  , et al. Empirical modelling of the BLASTPol achromatic half-wave plate for precision submillimetre polarimetry. *Monthly Notices of the RAS*, 437:2772–2789, January 2014.
- [191] K. D. Irwin. An application of electrothermal feedback for high resolution cryogenic particle detection. *Applied Physics Letters*, 66(15):1998–2000, 1995.
- [192] A. R. Trangsrud. *The SPIDER CMB Polarimeter*. PhD thesis, California Institute of Technology, 2012.
- [193] C. L. Kuo, J. J. Bock, J. A. Bonetti, et al. Antenna-coupled TES bolometer arrays for CMB polarimetry. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7020 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, August 2008.
- [194] K. D. Irwin and G. C. Hilton. *Cryogenic particle detection*, chapter Transition-edge sensors. Springer, 2005.
- [195] J. C. Mather. Bolometer noise: nonequilibrium theory. *Appl. Opt.*, 21(6):1125–1129, Mar 1982.
- [196] J. C. Mather. Bolometers: ultimate sensitivity, optimization, and amplifier coupling. *Applied Optics*, 23:584–588, February 1984.
- [197] R. S. Tucker. *Characterization of Detectors and Instrument Systematics for the SPIDER CMB Polarimeter*. PhD thesis, California Institute of Technology, 2014.
- [198] D. H. Martin and E. Puplett. Polarised interferometric spectrometry for the millimetre and submillimetre spectrum. *Infrared Physics*, 10(2):105 – 109, 1970.
- [199] E. M. Bierman, T. Matsumura, C. D. Dowell, et al. A millimeter-wave galactic plane survey with the bicep polarimeter. *The Astrophysical Journal*, 741(2):81, 2011.
- [200] J. D. Kraus. *Antennas*. First edition, 1950.
- [201] T. E. Hileman. *Probing Inflationary Cosmology: The Atacama B-Mode Search (ABS)*. PhD thesis, Princeton University, 2011.
- [202] J. D. Soler, P. A. R. Ade, M. Amiri, et al. A Light-weight Gondola for the Balloon-borne SPIDER Experiment. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, June 2014.

- [203] A. S. Rahlin, P. A. R. Ade, M. Amiri, et al. Pre-Flight Integration and Characterization of the SPIDER Balloon-Borne Telescope. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, June 2014.
- [204] N. N. Gandilo, P. A. R. Ade, M. Amiri, et al. Attitude Determination for Balloon-borne Experiments. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, June 2014.
- [205] E. Pascale, P. A. R. Ade, J. J. Bock, et al. The Balloon-borne Large Aperture Submillimeter Telescope: BLAST. *Astrophysical Journal*, 681:400–414, July 2008.
- [206] The Planck Collaboration. The Scientific Programme of Planck. *ArXiv Astrophysics e-prints*, April 2006.
- [207] Planck Collaboration. Planck 2013 Results. VII. HFI time response and beams. *ArXiv e-prints*, March 2013.
- [208] T. Stute. The Planck Telescope Reflectors. In J. Antebi and D. Lemke, editors, *Astronomical Structures and Mechanisms Technology*, volume 5495 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 1–10, September 2004.
- [209] S. Triqueneaux, L. Sentis, P. Camus, A. Benoit, and G. Guyot. Design and performance of the dilution cooler system for the Planck mission. *Cryogenics*, 46:288–297, April 2006.
- [210] ESA. Multimedia Gallery - Planck Spacecraft Cut-away View. <http://sci.esa.int/planck/44151-planck-spacecraft-cut-away-view/>, April 2014.
- [211] R. A. Sunyaev and Y. B. Zeldovich. Small-Scale Fluctuations of Relic Radiation. *Astrophysics and Space Science*, 7:3–19, April 1970.
- [212] W. C. Jones, R. S. Bhatia, J. J. Bock, and A. E. Lange. A Polarization Sensitive Bolometric Detector for Observations of the Cosmic Microwave Background. *ArXiv Astrophysics e-prints*, September 2002.
- [213] W. A. Holmes, J. J. Bock, B. P. Crill, et al. Initial Test Results on Bolometers for the Planck High Frequency Instrument. *Applied Optics*, 47:5996–6008, November 2008.
- [214] Planck Collaboration. Planck 2013 Results. IX. HFI spectral response. *ArXiv e-prints*, March 2013.
- [215] J. A. Murphy, R. Colgan, E. Gleeson, et al. Corrugated horn design for HFI on PLANCK. In M. de Petris and M. Gervasi, editors, *Experimental Cosmology at Millimetre Wavelengths*, volume 616 of *American Institute of Physics Conference Series*, pages 282–289, May 2002.

- [216] J. A. Murphy, T. Peacocke, B. Maffei, et al. Multi-mode horn design and beam characteristics for the Planck satellite. *Journal of Instrumentation*, 5:4001, April 2010.
- [217] J. A. Murphy and R. Padman. Radiation patterns of few-moded horns and condensing lightpipes. *Infrared Physics*, 31:291–299, 1991.
- [218] J. A. Tauber, H. U. Norgaard-Nielsen, P. A. R. Ade, et al. Planck Pre-launch Status: The Optical System. *Astronomy and Astrophysics*, 520:A2, September 2010.
- [219] B. Maffei, F. Noviello, J. A. Murphy, et al. Planck Pre-launch Status: HFI Beam Expectations from the Optical Optimisation of the Focal Plane. *Astronomy and Astrophysics*, 520:A12, September 2010.
- [220] X. Dupac and J. Tauber. Scanning Strategy for Mapping the Cosmic Microwave Background Anisotropies with Planck. *Astronomy and Astrophysics*, 430:363–371, January 2005.
- [221] J. A. Tauber, N. Mandolesi, J.-L. Puget, et al. Planck pre-launch status: The Planck mission. *Astronomy and Astrophysics*, 520:A1, September 2010.
- [222] Planck Collaboration, P. A. R. Ade, N. Aghanim, et al. Planck early results. I. The Planck mission. *Astronomy and Astrophysics*, 536:A1, December 2011.
- [223] Planck Collaboration, P. A. R. Ade, N. Aghanim, et al. Planck 2013 Results. I. Overview of products and scientific results. *ArXiv e-prints*, March 2013.
- [224] Planck Collaboration. Planck 2013 Results. VIII. HFI photometric calibration and mapmaking. *ArXiv e-prints*, March 2013.
- [225] G. Hinshaw, J. L. Weiland, R. S. Hill, et al. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Data Processing, Sky Maps, and Basic Results. *Astrophysical Journal, Supplement*, 180:225–245, February 2009.
- [226] Planck Collaboration. Planck 2014 Results. I. Calibration & Temperature and Polarization Maps. In prep., October 2014.
- [227] M. White and M. Srednicki. Window functions of cosmic microwave background experiments. *Astrophysical Journal*, 443:6–10, April 1995.
- [228] T. Souradeep and B. Ratra. Window Function for Noncircular Beam Cosmic Microwave Background Anisotropy Experiment. *The Astrophysical Journal*, 560:28–40, October 2001.
- [229] Planck Collaboration. Planck 2014 Results. I. Time-Ordered Information and Beam Processing. In prep., October 2014.
- [230] D. Hanson, A. Lewis, and A. Challinor. Asymmetric beams and CMB statistical anisotropy. *Physical Review D*, 81(10):103003, May 2010.

- [231] S. Mitra, G. Rocha, K. M. Górski, et al. Fast Pixel Space Convolution for Cosmic Microwave Background Surveys with Asymmetric Beams and Complex Scan Strategies: FEBECoP. *Astrophysical Journal, Supplement*, 193:5, March 2011.
- [232] R. S. Hill, J. L. Weiland, N. Odegard, et al. Five-Year Wilkinson Microwave Anisotropy Probe Observations: Beam Maps and Window Functions. *Astrophysical Journal, Supplement*, 180:246–264, February 2009.
- [233] R. A. Monsalve. Beam characterization for the QUIET Q-Band instrument using polarized and unpolarized astronomical sources. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7741 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010.
- [234] K. M. Huffenberger, B. P. Crill, A. E. Lange, K. M. Górski, and C. R. Lawrence. Measuring Planck Beams with Planets. *Astronomy and Astrophysics*, 510:A58+, February 2010.
- [235] G. Roudier. *Contraintes sur la birefringence cosmique à partir de l'analyse des données polarisées du fond diffus cosmologique de Planck*. PhD thesis, Université Paris Diderot - Paris 7, 2012.
- [236] M. Shimon, B. Keating, N. Ponthieu, and E. Hivon. CMB polarization systematics due to beam asymmetry: Impact on inflationary science. *Physical Review D*, 77:083003, Apr 2008.
- [237] N. J. Miller, M. Shimon, and B. G. Keating. CMB beam systematics: Impact on lensing parameter estimation. *Physical Review D*, 79(6):063008, March 2009.
- [238] D. O’Dea, A. Challinor, and B. R. Johnson. Systematic errors in cosmic microwave background polarization measurements. *Monthly Notices of the RAS*, 376:1767–1783, April 2007.
- [239] P. J. E. Peebles. Statistical Analysis of Catalogs of Extragalactic Objects. I. Theory. *Astrophysical Journal*, 185:413–440, October 1973.
- [240] E. Hivon, K. M. Górski, C. B. Netterfield, et al. MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets. *Astrophysical Journal*, 567:2–17, March 2002.
- [241] E. L. Wright. Infrared Brightness Temperature of Mars, 1983-2103. *ArXiv Astrophysics e-prints*, March 2007.
- [242] E. Lellouch and H. Amri. Mars brightness model. <http://www.lesia.obspm.fr/perso/emmanuel-lellouch/mars/>, 2008.
- [243] J. Ruze. Antenna tolerance theory – a review. *Proceedings of the IEEE*, 54(4):633–640, April 1966.

[244] M. J. Klein and M. D. Hofstadter. Long-term variations in the microwave brightness temperature of the Uranus atmosphere. *Icarus*, 184:170–180, September 2006.

[245] C. Kramer, R. Moreno, and A. Greve. Long-term observations of Uranus and Neptune at 90 GHz with the IRAM 30 m telescope. (1985-2005). *Astronomy and Astrophysics*, 482:359–363, April 2008.

[246] J. L. Weiland, N. Odegard, R. S. Hill, et al. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Planets and Celestial Calibration Sources. *Astrophysical Journal, Supplement*, 192:19, February 2011.

[247] M. Hasselfield, K. Moodley, J. R. Bond, et al. The Atacama Cosmology Telescope: Beam Measurements and the Microwave Brightness Temperatures of Uranus and Saturn. *The Astrophysical Journal Supplement Series*, 209(1):17, 2013.

[248] T. Louis, G. E. Addison, M. Hasselfield, et al. The Atacama Cosmology Telescope: Cross Correlation with Planck maps. *ArXiv e-prints*, March 2014.

[249] A. B. Goldin, M. S. Kowitt, E. S. Cheng, et al. Whole-Disk Observations of Jupiter, Saturn, and Mars in Millimeter/Submillimeter Bands. *Astrophysical Journal, Letters*, 488:L161–L164, October 1997.

[250] M. C. Runyan, P. A. R. Ade, R. S. Bhatia, et al. ACBAR: The Arcminute Cosmology Bolometer Array Receiver. *Astrophysical Journal, Supplement*, 149:265–287, December 2003.

[251] R. Moreno. Neptune and Uranus planetary brightness temperature tabulation. Tech. rep., ESA Herschel Science Center, 2014.

[252] M. Showalter. Planet Viewers. <http://pds-rings.seti.org/tools/index.html>, March 2014.

[253] D. J. Rudy, D. O. Muhleman, G. L. Berge, B. M. Jakosky, and P. R. Christensen. Mars - VLA observations of the northern hemisphere and the north polar region at wavelengths of 2 and 6 CM. *Icarus*, 71:159–177, July 1987.

[254] D. O. Muhleman and G. L. Berge. Observations of Mars, Uranus, Neptune, Io, Europa, Ganymede, and Callisto at a wavelength of 2.66 MM. *Icarus*, 92:263–272, August 1991.

[255] H. M. Böttger, S. R. Lewis, P. L. Read, and F. Forget. Modelling water transport in the Martian atmosphere using the European Mars GCM. In *Bulletin of the American Astronomical Society*, volume 33 of *Bulletin of the American Astronomical Society*, pages 1068–+, November 2001.

[256] E. W. Weisstein. *Millimeter/Submillimeter Fourier Transform Spectroscopy of Jovian Planet Atmospheres*. PhD thesis, California Institute of Technology, 1996.

[257] E. W. Weisstein and E. Serabyn. Submillimeter line search in jupiter and saturn. *Icarus*, 123(1):23 – 36, 1996.

[258] J. Gibson, W. J. Welch, and I. de Pater. Accurate jovian radio flux density measurements show ammonia to be subsaturated in the upper troposphere. *Icarus*, 173:439–446, February 2005.

[259] R. Moreno. Private communications, 2014.

[260] J. N. Cuzzi, J. A. Burns, S. Charnoz, et al. An evolving view of saturn’s dynamic rings. *Science*, 327(5972):1470–1475, 2010.

[261] R. M. Canup. Origin of Saturn’s rings and inner moons by mass removal from a lost Titan-sized satellite. *Nature*, (7326):943–926, 2010.

[262] M. Hasselfield, K. Moodley, J. R. Bond, et al. The Atacama Cosmology Telescope: Beam Measurements and the Microwave Brightness Temperatures of Uranus and Saturn. *Astrophysical Journal, Supplement*, 209:17, November 2013.

[263] M.J. Griffin and G.S. Orton. The near-millimeter brightness temperature spectra of uranus and neptune. *Icarus*, 105(2):537 – 547, 1993.

[264] Y. D. Takahashi, P. A. R. Ade, D. Barkats, et al. Characterization of the BICEP Telescope for High-precision Cosmic Microwave Background Polarimetry. *Astrophysical Journal*, 711:1141–1156, March 2010.

[265] M. Born and E. Wolf. *Principles of Optics*. October 1999.

[266] E. Hecht and A. R. Ganesan. *Optics*. Pearson, 2008.

[267] C. J. MacTavish, P. A. R. Ade, E. S. Battistelli, et al. Spider Optimization: Probing the Systematics of a Large-Scale B-Mode Experiment. *Astrophysical Journal*, 689:655–665, December 2008.

[268] Planck Collaboration. Planck 2013 Results. XXIV. Constraints on primordial non-Gaussianity. *ArXiv e-prints*, March 2013.

[269] H. V. Peiris, E. Komatsu, L. Verde, et al. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Implications For Inflation. *Astrophysical Journal, Supplement*, 148:213–231, September 2003.

[270] F. Piacentini, P. A. R. Ade, J. J. Bock, et al. A Measurement of the Polarization-Temperature Angular Cross-Power Spectrum of the Cosmic Microwave Background from the 2003 Flight of BOOMERANG. *Astrophysical Journal*, 647:833–839, August 2006.

[271] K. Freese and W. H. Kinney. Natural Inflation: Consistency with Cosmic Microwave Background Observations of Planck and BICEP2. *ArXiv e-prints*, March 2014.

[272] A. Ijjas, P. J. Steinhardt, and A. Loeb. Inflationary schism after Planck2013. *ArXiv e-prints*, February 2014.

- [273] R. Penrose. Difficulties with inflationary cosmology. *Annals of the New York Academy of Sciences*, 571:249–264, December 1989.
- [274] S. M. Carroll and H. Tam. Unitary Evolution and Cosmological Fine-Tuning. *ArXiv e-prints*, July 2010.
- [275] M. Tegmark. How unitary cosmology generalizes thermodynamics and solves the inflationary entropy problem. *Phys. Rev. D*, 85:123517, Jun 2012.
- [276] A. Linde. Inflationary Cosmology after Planck 2013. *ArXiv e-prints*, February 2014.
- [277] A. H. Guth, D. I. Kaiser, and Y. Nomura. Inflationary paradigm after Planck 2013. *Physics Letters B*, 733:112–119, June 2014.
- [278] P. J. Steinhardt. Cosmology: The Inflation Debate. *Scientific American*, 304(4):36–43, April 2011.
- [279] J. E. Austermann, K. A. Aird, J. A. Beall, et al. SPTpol: an instrument for CMB polarization measurements with the South Pole Telescope. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8452 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, September 2012.
- [280] A. Kusaka, D. J. Fixsen, A. J. Kogut, et al. MuSE: a novel experiment for CMB polarization measurement using highly multimoded bolometers. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8452 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, September 2012.
- [281] S. Jacob, S. Kasthurirengan, and R. Karunanithi. Part 1: Calorimetric Studies. *Cryogenics*, 32(12):1137–1146, 1992.
- [282] R. J. Corruccini. Gaseous Heat Conduction at Low Pressures and Temperatures. *Vacuum*, 7-8:19–29, 1959.
- [283] P. J. Sun, J. Y. Wu, P. Zhang, L. Xu, and M. L. Jiang. Experimental Study of the Influences of Degraded Vacuum on Multilayer Insulation Blankets. *Cryogenics*, 49(12):719–726, December 2009.
- [284] R. G. Scurlock and B. Saull. Development of multilayer insulations with thermal conductivities below $0.1 \text{ } \mu\text{W cm}^{-1} \text{ K}^{-1}$. *Cryogenics*, 16(5):303 – 311, 1976.
- [285] J. J. Bock, M. K. Parikh, M. L. Fischer, and A. E. Lange. Emissivity measurements of reflective surfaces at near-millimeter wavelengths. *Applied Optics*, 34(22):4812–4816, August 1995.
- [286] J. B. Heaney. Efficiency of aluminized mylar insulation at cryogenic temperatures. *Proceedings of SPIE*, 3435:150–157, September 1998.
- [287] M. Houde, R. L. Akeson, J. E. Carlstrom, et al. Polarizing Grids, Their Assemblies, and Beams of Radiation. 113:622–638, May 2001.

[288] P. Arguijo and M. S. Scholl. Exact ray-trace beam for an off-axis paraboloid surface. *Applied Optics*, 42(16):3284–3289, 2003.

[289] T. J. Finn, N. Trappe, J. A. Murphy, and S. Withington. The gaussian beam mode analysis of off-axis aberrations in long wavelength optical systems. *Infrared Physics and Technology*, 51(4):351 – 359, 2008.

[290] D. Malacara. Some parameters and characteristics of an off-axis paraboloid. *Optical Engineering*, 30:1277–1280, September 1991.

[291] S. Withington, J. A. Murphy, and K. G. Isaak. Representation of mirrors in beam waveguides as inclined phase-transforming surfaces. *Infrared Physics and Technology*, 36(3):723 – 734, 1995.

[292] E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, Hoboken, NJ, ninth edition, 2006.

[293] M. M. Mizrahi. Generalized hermite polynomials. *Journal of Computational and Applied Mathematics*, 1(3):137 – 140, 1975.

[294] I. S. Gradshteyn, I. M. Ryzhik, A. Jeffrey, and D. Zwillinger. *Table of Integrals, Series, and Products*. 2007.

[295] M. C. Hudson. Calculation of the Maximum Optical Coupling Efficiency into Multimode Optical Waveguides. *Applied Optics*, 13(5):1029–1033, 1974.

[296] H. Hodara and C. Slemon. Throughput and coupling in optical fibers. *Applied Scientific Research*, 41:203–221, 1984. 10.1007/BF00382453.

[297] JPL HORIZONS. HORIZONS Web-Interface. <http://ssd.jpl.nasa.gov/?ephemerides>, April 2014.

[298] J. C. Mather, D. J. Fixsen, R. A. Shafer, C. Mosier, and D. T. Wilkinson. Calibrator Design for the COBE Far-Infrared Absolute Spectrophotometer (FIRAS). *Astrophysical Journal*, 512:511–520, February 1999.