

Beam Systematics and Ground Signal Characterization in the BICEP3 and BICEP Array CMB Experiments

A dissertation submitted to the
Graduate College of the University of Cincinnati
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the Department of Physics
of the College of Arts and Sciences

by

Christos Giannakopoulos

March 2026

supervised by

Dr. Colin Bischoff

Committee Member: Dr. Joachim Brod, Physics

Committee Member: Dr. Matthew Bayliss, Physics

Committee Member: Dr. Alexandre Sousa, Physics

Abstract

Abstract

Detection of excess B-mode polarization (parameterized by the tensor-to-scalar ratio— r) in the cosmic microwave background (CMB) at degree angular scales would provide direct evidence for primordial gravitational waves (PGWs) sourced by cosmic inflation. The BICEP/Keck (BK) series of experiments, all located at the Amundsen-Scott South Pole Station, has produced the most stringent constraints of $r < 0.036$ at 95% confidence. Achieving this sensitivity requires exquisite control of instrumental systematic errors, some of which originate in the optical response of the detectors.

The measured CMB polarization signal is reconstructed from pair-differencing of two co-located, orthogonally polarized detectors. Any mismatch in the beam response of these detectors can lead to leakage of the much brighter temperature signal into the polarization maps ($T \rightarrow P$ leakage), resulting in a biased estimate of r . Additionally, any non-uniform cross polar response will lead to leakage of E-mode polarization into B-mode polarization ($P \rightarrow P$ leakage) that could also bias the measurement of r .

In this thesis, we present the development and results of a comprehensive calibration program for the BICEP3 and BICEP Array CMB polarimeters, spanning beam characterization, sidelobe measurements, and azimuth-synchronous signal characterization. We describe and improve the now-scalable far-field beam mapping (FFBM) analysis pipeline producing high-fidelity beam maps of our difference and residual beams. These maps are utilized in our custom suite of beam simulations to directly predict the amount of $T \rightarrow P$ leakage in BK

data. We also develop a curved sky convolution framework to predict $P \rightarrow P$ leakage directly, which can also be used to model any anisotropic beam contamination without any approximations. We find the $P \rightarrow P$ leakage to be $\mathcal{O}(10^{-5})$, confirming this systematic to be subdominant for current and near-future sensitivity levels. During this process, we quantify inductive crosstalk contamination in BICEP3 that led us to extend the existing deprojection framework to remove this systematic from our data, which allows us to pass null tests that were previously failing due to this effect.

Using a linearly polarized broadband noise source (BSNS), we fully reconstruct the polarized beam response of BICEP3 detectors and indirectly measure their polarization efficiency of $\sim 99.1\% \pm 0.9\%$. We characterize the near and far sidelobes of BICEP3 and BICEP Array, validate the instrument baffling performance, and quantify the contribution of ghost beams that carry $\sim 1.1\%$ of the main beam power. These ghost beams, combined with the AM atmospheric model, explain the previously elusive detector gain systematics related to boresight rotation.

Finally, we present the first 360° direct measurements of ground emission for both BICEP3 and BICEP Array experiments. The resulting ground maps are calibrated to physical temperature units and characterize both the spatial distribution and temporal fluctuations of the ground signal. We extract a lower bound on the ground suppression factor, which exceeds 10^8 for all BICEP Array receivers.

© March 2026 — Christos Giannakopoulos

ALL RIGHTS RESERVED.

Acknowledgments

Contrary to the popular image of the physicist as a “lone genius” working in isolation, physics is not done in a vacuum. It is a collaborative effort between a community of people with a variety of different backgrounds, coming together to further humanity’s understanding of the Universe. This community extends far beyond scientists and engineers. It includes people who offer mentorship, physical, mental, emotional support, and all necessary resources to make our work as physicists possible. The intentional use of the word ‘we’ throughout this thesis reflects that no part of my work was done in isolation, and is meant to acknowledge the contributions of the people in my life that have made this journey possible.

First I want to thank my advisor Colin Bischoff. You are an incredible physicist, who has always treated me with respect, kindness, and supported me during the difficult times of the Ph.D. journey. Thank you for teaching me scientific rigor, and what it means to be an experimental physicist. Anyone should consider themselves lucky to have you as an advisor or collaborator, as your deep knowledge of our work combined with your unique perspective make you an invaluable asset to any project.

I am privileged to have worked with the principal investigators of the BICEP collaboration: John Kovac, Clem Pryke, Chao-Lin Kuo, Jamie Bock, and Kirit Karkare, whose feedback and guidance have shaped the quality of my work and have taught me thorough and careful scientific analysis.

Being part of the smallest group of the collaboration, it would have been easy to feel isolated. Thanks to the Beam Team comprised of Clara Vergès (the supreme leader), Brodi Elwood, and (most recently) Marcela Izquierdo, I have always felt a sense of belonging. Thank you

for the laughs and comic relief during our lengthy beam runs, and for always building each other up. Special thanks to Clara for being an incredible mentor and always being available for impromptu meetings to discuss various scientific analyses and to Brodi for always being willing to entertain my ideas and questions.

A heartfelt thank you to the most detail-oriented and helpful person I have had the privilege of working with, Jamie Cheshire, who has been an invaluable source of knowledge throughout my time in the collaboration (and—at minimum—pity-laughed at my jokes). I am particularly grateful to Silvia Zhang for being a great friend and support, as well as teaching me the ins and outs of working in a lab and tolerating my clumsiness. Thank you to the colleagues that I have the privilege to call friends, specifically James Cornelison, Baibhav Singari, Yuka Nakato, Sofia Fatigoni, and Thuong Hoang for always making time to chat about science and life.

Here in Cincinnati, I cannot stress enough how lucky I have been to be surrounded by a community of people who have kept me centered and shaped the person I am today. To my closest friends, Jeffrey Chin, Savannah Koch, Christian Coronel, Alex Maries, Tony Crawford, George Sias, Kathy Sias, Gabe Nowak, Dickson Boahen, Mary Catherine, and Aiden Roberts Williams, I am deeply grateful for your friendship and support. A special thank you to Dave Payne for his continued encouragement and support of my mental health throughout the Ph.D. journey and for helping me navigate the tumultuous world events of the past 7 years.

Most importantly, I want to thank my lovely wife, Katie. I feel so lucky to have you as my life partner. Your work ethic, kindness, and selflessness are an inspiration, combined with your wisdom, and compassion, you have made my dreams a reality. You have stood with me through the most difficult time of my life, and it is for these and many more reasons that this thesis is dedicated to you. Katie, thank you for offering the stability to tame the ups and downs of life throughout the entire decade we have been together. Regardless of what life has in store for us next, navigating it with you is the greatest gift I could ask for.

FOR MY WIFE, KATIE. MY REAL LIFE CHEAT CODE.

Contents

1	Introduction	1
1.1	Λ CDM Cosmology	1
1.1.1	General Relativity and the Einstein Field Equations	2
1.1.2	The Friedmann–Lemaître–Robertson–Walker Metric	3
1.1.3	The Friedmann Equations	4
1.1.4	Density Parameters and the Critical Density	7
1.1.5	Cosmological Parameters	7
1.2	Inflation	8
1.2.1	The Inflationary Mechanism	11
1.2.2	Slow-Roll Inflation	13
1.2.3	Primordial Perturbations from Inflation	15
1.3	The Cosmic Microwave Background	17
1.3.1	Recombination and the Last-Scattering Surface	20
1.3.2	Temperature Anisotropies	21
1.3.3	CMB Polarization	24
1.3.4	Sources of E-Mode and B-Mode Polarization	27
1.3.5	Gravitational Lensing of the CMB	30

1.3.6	Current Observational Constraints on r	32
1.4	Foregrounds	33
1.4.1	Thermal Dust Emission	34
1.4.2	Synchrotron Radiation	35
1.4.3	Multi-Frequency Foreground Separation	35
1.5	Thesis Outline	36
2	BICEP/Keck Experiments	38
2.1	Program History	38
2.2	Instrument Overview	40
2.2.1	BICEP3	40
2.2.2	BICEP Array	43
2.3	Detectors and Readout	46
2.4	Baffling	48
2.5	Data Taking and Observation Strategy	49
2.5.1	Starpointing	52
2.6	Pointing Model	53
2.7	From Timestreams to Power Spectra	58
2.7.1	Data reduction	59
2.7.2	Mapmaking	64
2.7.3	Deprojection	67
2.7.4	Absolute calibration	72
2.7.5	Power spectra	72
2.7.6	Null Tests	77

2.8	Multi-Component Likelihood Analysis	79
3	Calibration Measurements Overview and Analysis	83
3.1	Beam Mapping	84
3.2	Far Field Beam Mapping Setup	87
3.3	Sources	90
3.3.1	Thermal chopper	90
3.3.2	Broad Spectrum Polarized Source	91
3.4	Coordinate Systems	92
3.5	Analysis Pipeline	105
3.5.1	Timestream Demodulation	106
3.5.2	Beam Parameterization and Pair-Centroid Centering	110
3.5.3	Binning and Masking	112
3.5.4	Pipeline Performance	113
3.5.5	Automatic Cuts	115
4	Beam Characterization Results	118
4.1	Unpolarized Far Field Beam Mapping	119
4.1.1	Composite Maps and Beam Window Functions	120
4.1.2	Beam Parameter Results	123
4.1.3	Residual Beams and Temperature-to-Polarization Leakage Simulations	127
4.1.4	Crosstalk Beams	132
4.2	Polarized Beam Mapping	138
4.2.1	Per-detector Mapmaking for Polarized Beam Mapping	140

4.3	Far Sidelobe Mapping	149
5	BICEP3 Systematics Studies With Linearly Polarized Source	160
5.1	Polarization-to-Polarization Leakage Simulations	161
5.2	Gain Systematics Due to Boresight Rotation	170
5.2.1	Ghost Beams	172
5.2.2	Atmospheric Modeling	178
5.2.3	Elnod Gain Model	180
6	Ground Mapping	185
6.1	Measurement Overview	187
6.2	Analysis Pipeline, Ground Maps, and Ground Movies	188
6.3	Ground Suppression Factor	198
6.4	Temporal Fluctuation	202
6.5	Future Work	204
7	Conclusion	209
	Bibliography	211
	Appendices	223
A	BA30/40 Detector Time Constant Measurements	223
B	BA150 Corrugation features	233
C	Noise modeling in Ground Spectra	238

List of Figures

1.1	Inflation to the Rescue	12
1.2	CMB sky temperature anisotropy	18
1.3	CMB temperature anisotropy power spectrum	19
1.4	E-mode and B-mode polarization patterns	26
1.5	Theoretical CMB power spectra	31
1.6	Constraint on r BK18	33
2.1	BICEP/Keck Program History	39
2.2	BICEP3 and BA Bandpasses and Atmospheric Transmission	41
2.3	BICEP3 Optical Path	42
2.4	BICEP3 Cryostat Cutout	43
2.5	BICEP3 Tile Orientations	44
2.6	BICEP Array Cryostat Cutout and Optical Rays	45
2.7	BA30/40 Focal Plane	46
2.8	BA vs B3 Wafers	47
2.9	BICEP3 Mount and Baffling	50
2.10	BICEP/Keck Observing Field	51
2.11	BICEP Array CMB Phases	52

2.12	BICEP Array Receiver Orientations on the Drum	54
2.13	CMB Pointing Reconstruction	57
2.14	BICEP3 Raw Timestreams	60
2.15	Observed Elnod Gains	62
2.16	BICEP3 Azimuth-fixed Template	63
2.17	2024 BA150 Single Year Coadded Map	68
2.18	Beam Mismatch Modes	70
2.19	Fourier Domain EE Power	75
2.20	Total BB Suppression Factor BK18	81
2.21	MCMC Results BK18	82
3.1	Far-Field Beam Mapping Terminology	86
3.2	South Pole Facilities	88
3.3	B3 Mirror Setup	89
3.4	Far-field Beam Mapping Sources	92
3.5	Mirror Reflection Setup	95
3.6	Detector-Centered Coordinate	98
3.7	Focal Plane Orientation	99
3.8	Source Coordinates	100
3.9	Mirror-Centered Coordinates	102
3.10	Pencil Beam in Mirror Centered Coordinates	105
3.11	Timestream Demodulation	108
3.12	Chop Reference Repair	109
3.13	Mirror Projection onto Focal Plane	113

3.14	Thermal Chopper FFBM Map	116
4.1	BA150 Sample Composite Map	120
4.2	B3/BA array-averaged composite maps	122
4.3	BA150 FPU Total Ellipticity	125
4.4	BA Total Differential Ellipticity	126
4.5	BICEP3 and BA150 Residual Beams	128
4.6	BK18 BB Leakage Auto and Cross with Real	132
4.7	BICEP3 Crosstalk Beam Distance	133
4.8	BICEP3 Crosstalk Residual Beams	134
4.9	BICEP3 MCE0 Crosstalk Residual Beams	135
4.10	BICEP3 Total Crosstalk Amplitude	137
4.11	BICEP3 Polarized Component Maps	139
4.12	BICEP3 Polarized Beam Coadds	145
4.13	BICEP3 RPS Polarized Beam Coadds	146
4.14	BICEP3 BSNS and RPS Radial Profiles	148
4.15	BA30/40 FSL Mapping Setup	150
4.16	BICEP3 Array Coadded FSL Maps – Forebaffle On and Off	154
4.17	BA40 Array Coadded FSL Maps – Forebaffle On and Off	155
4.18	BICEP3 Composite Beam Profiles	157
4.19	BA40 Composite Beam Profiles	158
5.1	Max and Min of B_U Maps	163
5.2	Simulated Q and U Maps	164
5.3	Point Cloud Locations on Rectangular Beam	166

List of Figures

5.4	Sample HEALPix Point Cloud	168
5.5	$P \rightarrow P$ Leakage Spectra	169
5.6	BICEP3 Observed Elnod Gain Dependence on Dk	170
5.7	Ghost Beam Sketch	173
5.8	Buddy Beams seen with BSNS and Thermal Chopper	174
5.9	Buddy Beam Amplitude	176
5.10	Buddy Region Integral	177
5.11	BICEP3 Bandpass & Atmospheric Brightness	179
5.12	Detector Pointings on Curved Sky	181
5.13	Observed vs Modeled Elnod Gains	183
5.14	Observed vs Modeled Elnod Gains Correlation	184
6.1	South Pole Ground Panorama	186
6.2	South Pole Atmospheric Model	190
6.3	BA150 Atmospheric Fit	191
6.4	BICEP3 Ground Maps for V and H Detectors	193
6.5	Coadded Ground Maps 1	194
6.6	Coadded Ground Maps 2	195
6.7	BA150 Time Varying Ground Signal - Real Space	197
6.8	BA150 Azimuth-Fixed + Ground Signals Power Spectra	201
6.9	Ground Suppression Factor	202
6.10	BA150 Ground Signal Power Spectra at 20 Minute Timescale	203
6.11	Ground Signal Temporal Fluctuation	204
A.1	Time Constant Measurement Setup	227

A.2	Time Constant Fits	229
A.3	Time Constant Bias Dependence	230
A.4	BA30/40 Time Constants from BSNS Summary	230
A.5	Time Constant Correlation	231
B.1	BA150 Corrugation Schematic	234
B.2	Tile Features in BA150 FFBM	235
B.3	BA150 Tile 5 Beam Features Amplitude and Location	236

List of Tables

2.1	Stage 3 Receiver history	40
3.1	Far-field distances	85
3.2	Far-field beam mapping schedule parameters	90
4.1	FFBM Measurement Setup Parameters	119
4.2	Beam Parameter Summary	123
4.3	FSL Observation setup parameters	151
4.4	FSL Forebaffle loading measurements	156
6.1	Ground mapping setup	187

Chapter 1

Introduction

1.1 Λ CDM Cosmology

Cosmology, a Greek-derived word made up of two components, *cosmos* (meaning world) and *logos* (meaning study, or word), is the study of the origin of our Universe and its evolution. The Universe as we understand it today is vast, expanding, and mostly empty space. Over decades of observations, cosmologists have developed the Big Bang Theory as the current cosmological framework to describe the origin of the Universe. This theory suggests that the Universe began in a hot and very dense state as a primordial plasma with frequent particle interactions between atomic nuclei and photons. As the primordial plasma cooled, stable elements like Hydrogen and Helium began to form. At that moment, interactions between particles became much less frequent, and photons were able to stream freely through the Universe. These photons can be observed today at mm wavelengths and form the Cosmic Microwave Background (CMB) radiation. The CMB offers a unique window of the Early Universe and is the basis for cosmological observations. It has a uniform average temperature of 2.725 K [1], and fluctuations in CMB temperature at the level of $\sim 100 \mu\text{K}$ [2] that probe the density fluctuations in the early Universe.

Over the past century, theoretical advances and precision astronomical observations have allowed for the development of a ‘standard cosmological model’ with its parameters constrained

to remarkable precision. This model summarizes our current understanding of the Universe, and is known as Λ CDM.

Under the Λ CDM model, the Universe originated approximately 13.8 billion years ago in an extremely hot, dense state and has been expanding and cooling ever since. Its energy budget is dominated by a cosmological constant Λ (associated with dark energy driving the accelerated expansion of the Universe) and cold dark matter (CDM) needed to explain the observed large-scale structure, with ordinary baryonic matter constituting only a small fraction of the total [3].

The Λ CDM model assumes Einstein's general theory of relativity as the correct theory of gravity at cosmological scales, and the cosmological principle, which asserts that the Universe is spatially homogeneous and isotropic on sufficiently large scales ($\gtrsim 100$ Mpc). The resulting prediction of an expanding Universe is supported by many observations, including the redshift of distant galaxies and measurements of the CMB radiation. In this section, we provide a brief overview of the theoretical framework of the Λ CDM model, highlighting its key assumptions, and experimental observations that it leaves unaddressed. Then, following the prescription by Baumann [4], we describe cosmic inflation, a strongly motivated extension of the Λ CDM model that elegantly resolves these unaddressed topics. Although no smoking gun evidence for inflation has yet been found, the theory has become the leading paradigm for understanding the earliest moments of cosmic history and the origin of the primordial perturbations that seeded all structure in the Universe.

1.1.1 General Relativity and the Einstein Field Equations

The theoretical backbone of modern cosmology is Einstein's general theory of relativity. It relates the geometry of spacetime to the distribution of matter and energy within it. Using natural units: $c = \hbar = k_B = 1$, and the convention $(-, +, +, +)$ for the metric, the relationship is expressed through the Einstein field equations:

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu}, \quad (1.1)$$

where $R_{\mu\nu}$ is the Ricci curvature tensor, $R = g^{\mu\nu} R_{\mu\nu}$ is the Ricci scalar, $g_{\mu\nu}$ is the spacetime metric tensor describing the geometry of spacetime that depends on the spatial coordinates, Λ is the cosmological constant, G is Newton's gravitational constant, and $T_{\mu\nu}$ is the stress-energy tensor that describes the energy and momentum content of the Universe. The left-hand side of Equation 1.1 captures the curvature and geometry of spacetime, while the right-hand side captures the dynamics of matter and radiation. There are a total of 10 independent equations in the most general case, first put forth by Einstein in 1915 [5].

1.1.2 The Friedmann–Lemaître–Robertson–Walker Metric

The direct consequence of the cosmological principle is that the spacetime is maximally symmetric and homogeneous. Such 3-spaces have constant curvature, that is either positive, negative, or zero. The resulting general form of the spacetime metric during the early Universe is attributed to the work of Friedmann, Lemaître, Robertson, and independently by Walker (FLRW) metric. In spherical coordinates (r, θ, φ) , the line element takes the form:

$$ds^2 = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - k r^2} + r^2 (d\theta^2 + \sin^2 \theta d\varphi^2) \right], \quad (1.2)$$

where $a(t)$ is the dimensionless scale factor that characterizes the relative size of hypersurfaces at different times (set to 1 at the present time, t_0), k describes the spatial curvature ($\in \{-1, 0, +1\}$), and t is the cosmic time. The scale factor is related to the redshift z of an object observed today as $a(t) = 1/(1 + z)$. Curvature parameter $k = +1$ corresponds to a closed (spherical) hypersurface, $k = 0$ to a spatially flat hypersurface, and $k = -1$ to an open (hyperbolic) hypersurface.¹ r is not an observable distance, but rather a ‘comoving coordinate’. A more useful distance measure is the ‘proper distance’, defined as $D(t) = a(t) r$. From the expression of the velocity of a comoving object, $v = \dot{D} = H D + v_{\text{pec}}$, we get the

¹This can be more easily seen by defining a dimensionless radial coordinate $\chi = \frac{dr}{\sqrt{1 - k r^2}}$.

Hubble parameter:

$$H(t) = \frac{\dot{a}(t)}{a(t)}, \quad (1.3)$$

which quantifies the instantaneous rate of expansion of the Universe.² Overdots denote derivatives with respect to cosmic time t . The Hubble parameter's value at the present epoch, $H_0 \equiv H(t_0)$, is known as the Hubble constant. A more useful time coordinate η is also defined as:

$$\eta = \int \frac{dt}{a(t)}, \quad (1.4)$$

which has the convenient property that null geodesics (the paths of photons) appear as straight lines on a spacetime diagram plotted in conformal time versus comoving distance. It can be interpreted as a 'clock' that slows down as the Universe expands and makes it straightforward to assess whether two regions of spacetime could have been in causal contact,³ which will become important when we discuss the horizon problem in Section 1.2.

1.1.3 The Friedmann Equations

With the spacetime metric defined under the cosmological principle assumption, we turn to the right-hand side of 1.1. The most general form of the stress-energy tensor consistent with the symmetries of the FLRW metric (Equation 1.2) is that of a perfect fluid. Hence, the matter content of the Universe is modeled as such with energy density ρ and pressure p yielding a stress-energy tensor of the form:

² v_{pec} is the 'peculiar velocity' of an object, which is the velocity relative to the Hubble flow. Usually, due to gravitational forces on an object.

³Their comoving separation must be less than η .

$$T^\mu{}_\nu = \begin{pmatrix} -\rho & 0 & 0 & 0 \\ 0 & p & 0 & 0 \\ 0 & 0 & p & 0 \\ 0 & 0 & 0 & p \end{pmatrix}, \quad (1.5)$$

Inserting the FLRW metric into the Einstein field equations yields two independent equations (out of the 10 total in the most general case) governing the dynamics of the scale factor. The first Friedmann equation, obtained from the time-time component of equation 1.1 that is given by:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{k c^2}{a^2} + \frac{\Lambda c^2}{3}, \quad (1.6)$$

where the left-hand side of the equation represents the Hubble parameter squared, and is related to the energy density ρ , the curvature k , and the cosmological constant Λ . The second Friedmann equation, follows from the spatial components of the Einstein field equations which are all identical due to the isotropy assumption, and is given by:

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}(\rho + 3p) + \frac{\Lambda}{3}. \quad (1.7)$$

This equation reveals which conditions lead to accelerated ($\ddot{a} > 0$) or decelerated ($\ddot{a} < 0$) expansion. When matter (corresponding to density ρ) or radiation (corresponding to pressure p) dominate, the expansion decelerates, whereas when a cosmological constant (associated here with negative pressure) dominates, the expansion accelerates. To understand which component dominates at a given time, we need to understand how the energy density of each component evolves with the expansion of the Universe. To that end, we turn to the dynamics of the energy density in the expanding Universe.

In Minkowski spacetime, energy and momentum are conserved. Therefore, the energy density must satisfy the continuity equation (density dilution equals the divergence of the energy flux). Generalizing to curved spacetime, this implies that $\nabla_\mu T^\mu{}_\nu = 0$ where ∇_μ denotes the

covariant derivative. Expanding it for the perfect fluid under the FLRW metric yields the cosmological continuity equation:

$$\dot{\rho} + 3H(\rho + p) = 0. \quad (1.8)$$

which is a statement of local energy conservation.⁴ It is common to specify an equation of state that relates the pressure to the energy density for each component $p = w\rho$, where w is a dimensionless parameter whose value depends on the species. Equation 1.8 can then be solved to find the scaling of the energy density of each component with the expansion of the Universe. In terms of the equation-of-state parameter w , the scaling factors for the energy density of each component are summarized as follows:

- **Non-relativistic matter** ($w = 0$): Both baryonic matter (protons, neutrons, electrons) and CDM. Energy density scales as $\rho_m \propto a^{-3}$, decreasing in proportion to the expanding volume.
- **Radiation** ($w = 1/3$): Relativistic species such as photons and neutrinos. Using the continuity equation, the energy density of radiation scales as $\rho_r \propto a^{-4}$. Intuitively, the extra factor of a^{-1} beyond the scaling of matter (a^{-3}) is attributed to the photon wavelengths being stretched as the Universe expands.
- **Dark energy** ($w = -1$): Associated with the cosmological constant Λ . Its energy density (ρ_Λ) remains constant as the Universe expands, meaning $\rho_\Lambda \propto a^0$. Necessary to explain the observed accelerated expansion of the Universe [6].

Each component scales at different rates resulting in the dominant form of energy varying over cosmic history. In the early Universe, when a was very small, radiation dominated. As the Universe expanded, matter overtook radiation at a redshift of $z \sim 3400$, marking the transition to the matter-dominated era. As the scale factor approaches present day, the constant dark energy density has come to dominate over the ever-diluting matter.

⁴Symmetry under time translations is broken for the case of an expanding Universe, hence the term local.

1.1.4 Density Parameters and the Critical Density

It is convenient to introduce the critical density, defined as the total energy density in a spatially flat Universe in the absence of a cosmological constant (obtained by setting $k = 0$ and $\Lambda = 0$ in the first Friedmann equation):

$$\rho_{\text{crit}} \equiv \frac{3 H^2}{8\pi G}. \quad (1.9)$$

Then, we define dimensionless density parameters for each component as:

$$\Omega_i \equiv \frac{\rho_i}{\rho_{\text{crit}}}. \quad (1.10)$$

where the subscript i denotes the species (e.g. m for matter, r for radiation, k for curvature, Λ for dark energy). The first Friedmann equation (Equation 1.6) is then rewritten as:

$$\frac{H^2}{H_0^2} = \Omega_{r,0} a^{-4} + \Omega_{m,0} a^{-3} + \Omega_{k,0} a^{-2} + \Omega_{\Lambda,0}, \quad (1.11)$$

where the subscript 0 indicates present-day values. The curvature density parameter is defined as: $\Omega_{k,0} \equiv -k/(a_0 H_0)^2$. This form is convenient for comparing the relative contributions of each component since the sum of all density parameters is unity by definition ($\sum_i \Omega_i = 1$).

1.1.5 Cosmological Parameters

Precision measurements of the CMB, combined with observations of baryon acoustic oscillations (BAO), (see section 1.3.2), and Type Ia supernovae, have placed tight constraints on the parameters of the Λ CDM model. The most recent results from the *Planck* satellite [3] yield the following best-fit values (68% confidence):

$$\begin{aligned}
H_0 &= 67.4 \pm 0.5 \text{ km s}^{-1} \text{ Mpc}^{-1}, \\
\Omega_b h^2 &= 0.0224 \pm 0.0001, \\
\Omega_c h^2 &= 0.120 \pm 0.001, \\
\Omega_m &= 0.315 \pm 0.007, \\
\Omega_\Lambda &= 0.685 \pm 0.007,
\end{aligned} \tag{1.12}$$

where $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$. These directly correspond to the current content of the Universe: approximately 68% dark energy, 27% dark matter, and 5% baryonic matter, with a negligible contribution from radiation. The spatial curvature is constrained to be $\Omega_k = 0.001 \pm 0.002$ fully consistent with a flat Universe (see section 1.2).

We note that there is currently a $\sim 5\sigma$ discrepancy between the value of H_0 inferred from the CMB [3] and that measured from the local distance ladder using Cepheid-calibrated supernovae, which gives $H_0 = 73.0 \pm 1.0 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [7]. This is referred to as the ‘Hubble tension’ and remains an active area of research in modern cosmology.

The Λ CDM framework is an extraordinarily successful description of the observable Universe across a vast range of scales and epochs (as proven by the tight constraints on its parameters, despite the Hubble tension). Next we turn to some questions, already hinted at so far, that the Λ CDM model does not address. These questions motivate the development of an extension to this model known as the theory of cosmic inflation, which we discuss in the following section.

1.2 Inflation

High precision CMB observations indicate that the Universe is extraordinarily uniform at large scales. How is this possible given that most of the Universe appears to be causally disconnected? In addition, how is the curvature of the Universe today so close to zero,

given that any deviation from flatness is amplified by gravity over time? It could be argued that these are just manifestations of a quantum theory of gravity we do not currently understand. However, this hypothetical theory faces the problem of explaining how the small inhomogeneities observed in the Early Universe are highly correlated across large scales.⁵ These questions are named the *horizon problem*, the *flatness problem*, and the *superhorizon correlations* problem, respectively.

The theory of cosmic inflation, first proposed by A. Guth [8], resolves these issues by postulating a brief epoch of accelerated, quasi-exponential expansion in the very early Universe, prior to the radiation-dominated era. In addition, inflation puts forth a mechanism that produces the primordial density fluctuations that will evolve into the large-scale structure of the Universe we see today, fully explaining the superhorizon correlations. In this section, we review cosmic inflation and explain how it naturally solves the problems of the standard cosmological model while making testable predictions for precision cosmological observations.

The Horizon Problem

In comoving coordinates, null (light) geodesics are straight lines, making the choice of conformal time η (defined in equation 1.4) particularly convenient. If time ‘started’ at zero, then the greatest comoving distance that a photon could have traveled since the Big Bang at a given conformal time η is:

$$\Delta\eta = \int_0^t \frac{dt}{a(t)} = \int_0^a \frac{1}{a^2 H(a)} da = \int_0^{\ln a} (aH)^{-1} d\ln a, \quad (1.13)$$

This is also known as the particle horizon. The quantity $(aH)^{-1}$ is the comoving Hubble radius and hence $a = 0$ corresponds to the Big Bang singularity. At the time of recombination ($z \sim 1100$), our Universe (to a good approximation) is matter dominated, and the comoving Hubble radius monotonically grows with time. In the case of the CMB when viewed from the Earth today, this means that the particle horizon at recombination was much smaller

⁵‘Large scales’ is defined as distances larger than light could have traveled since the Big Bang.

than the comoving distance to the surface of last scattering (see section 1.3.1) i.e. $\eta_{\text{rec}} \ll \eta_0$. This implies that most parts of the CMB were never in causal contact with each other (or equivalently, their past light cones do not intersect). How is it possible that their respective temperatures are uniform to more than one part in 10^4 ? Using the current measured value of the Hubble constant, and assuming a Universe dominated by only radiation and matter at the time of recombination, we can calculate that the angle subtended by the particle horizon at recombination to be:

$$\theta_{\text{hor}} = \frac{2\eta_{\text{rec}}}{\eta_0 - \eta_{\text{rec}}} \approx 2^\circ. \quad (1.14)$$

Not only is the CMB uniform across the entire sky, but also the observed temperature fluctuations are correlated across angles much larger than θ_{hor} (i.e. superhorizon correlations).

The Flatness Problem

The curvature density parameter, Ω_k , can be expressed in terms of the Friedmann equation as:

$$\Omega_k = \frac{-k}{(aH)^2}. \quad (1.15)$$

Because Ω_k scales as $(aH)^{-2}$ and the comoving Hubble radius grows during both radiation- and matter-dominated eras, any departure from perfect flatness ($\Omega_k = 0$) is dynamically amplified over time (i.e. curvature wins over time). The fact that the Universe is observed to be spatially flat today to better than $\sim 0.2\%$ accuracy means that at times near $t = 0$ (e.g. $t \sim 10^{-43}$ s), the curvature must have satisfied $|\Omega_k| \lesssim 10^{-62}$. With no explanation of what physical mechanism could have set Ω_k so close to zero at early times, this constitutes the flatness problem. It is worth noting that the flatness and horizon problems are closely related, meaning that any solution to one is likely to solve the other as well.

The Origin of Primordial Perturbations

Although the standard Λ CDM model successfully describes how small initial density perturbations grow under gravitational instability to form the large-scale structure that we see today, it does not provide a mechanism for generating these perturbations in the first place. Inflation naturally provides such a mechanism by introducing quantum vacuum fluctuations in the *inflaton* field (see section 1.2.2), stretched to macroscopic scales by the exponential expansion. These seed all cosmic structure [9, 10].

1.2.1 The Inflationary Mechanism

With what sounds like a ‘hot fix’ to the horizon and flatness problems, inflation addresses them by postulating a period during which the comoving Hubble radius $(aH)^{-1}$ rapidly *shrinks*.⁶ This decrease during inflation allows quantum fluctuations to be stretched to superhorizon scales, only to then re-enter the horizon at later times (during radiation and matter dominated eras). Effectively, the majority of the contribution to the integral in Equation 1.13 would be from this early period. This is achieved by taking the derivative of the comoving Hubble radius with respect to time and setting it to be negative:

$$\frac{d}{dt}(aH)^{-1} = \frac{d}{dt}(\dot{a}^{-1}) = -\frac{\ddot{a}}{(\dot{a})^2}. \quad (1.16)$$

For $(d/dt)(aH)^{-1} < 0$, then $\ddot{a} > 0$ which implies accelerated expansion. Another way of thinking about it is that now that the comoving Hubble radius is dominated by *early times*, the particle horizon (Equation 1.13) can be much larger. Although sometimes the terms: particle horizon and Hubble radius are used interchangeably (especially in the context of matter and radiation dominated eras), it is important to note the two are very different during inflation.

Referring to the second Friedmann equation (Equation 1.7), the condition on the derivative

⁶How to fix things not being in causal contact? Postulate that at some early times they were!

of the comoving Hubble radius is met whenever the pressure satisfies $p < -\rho/3$, i.e., the dominant energy component must have an equation-of-state parameter $w < -1/3$. With a shrinking comoving Hubble radius ⁷, causally disconnected regions today were in causal contact prior to the onset of inflation (i.e. the entire observable Universe was within the particle horizon at the start of inflation). By assuming that the Hubble parameter H is approximately constant during inflation, $da/a \approx H dt$ which implies that the scale factor grows exponentially as $a(t) \propto e^{Ht}$, and the comoving Hubble radius shrinks as $(aH)^{-1} \propto e^{-Ht}$. The exponential expansion pushes causally connected regions beyond causal contact, but they were already equilibrated while within the particle horizon. Hence, horizon problem is solved. Similarly, any pre-existing spatial curvature is exponentially diluted during inflation. If the scale factor grows by a factor of e^N during inflation, then $\Omega_k \propto (aH)^{-2}$ is suppressed by e^{-2N} . For $N \gtrsim 60$ e-folds of expansion, this drives Ω_k to negligibly small values, naturally explaining the observed flatness without a need for extraordinary initial conditions. Figure 1.1 illustrates how inflation solves the horizon and flatness problems.




Figure 1.1. Illustration of how inflation solves the horizon and flatness problems. The comoving Hubble radius shrinks during inflation, allowing regions that are causally disconnected today to have been in causal contact prior to inflation. The gray line represents a fluctuation mode that starts within the Hubble radius prior to inflation, exits the Hubble radius during inflation, and re-enters during the later epochs (radiation, matter, dark energy) seen by us today as *superhorizon correlations*. For example, a superhorizon mode that enters the Hubble radius 60 e-folds after the Big Bang would correspond to a mode that exited the Hubble radius 60 e-folds before the end of inflation. Plot from [4].

⁷The comoving Hubble radius is more convenient to work here as it makes no assumptions about times earlier than t_i . As opposed to the particle horizon that depends on the entire history of the Universe prior to inflation.

1.2.2 Slow-Roll Inflation

For inflation to happen in a way that is consistent with the Universe we see today, the parameters of the model need to vary *slowly*. Prior to any assumptions on the dynamics of the inflationary expansion, we can directly quantify what we mean by ‘slowly’ by looking at the derivative of the Hubble radius with respect to time (Equation 1.16). It can be rewritten as:

$$\frac{d}{dt}(aH)^{-1} = -\frac{1}{a} \left(1 - \frac{\dot{H}}{H^2} \right). \quad (1.17)$$

Looking at the term in parentheses, we see that for a shrinking comoving Hubble radius, we need $\dot{H}/H^2 < 1$. This is the first slow-roll condition:

$$\epsilon \equiv -\frac{\dot{H}}{H^2} < 1. \quad (1.18)$$

It says that inflation occurs for as long as the fractional change in the Hubble parameter is small compared to the expansion rate itself.

The second slow-roll condition is placed by the fact that inflation must last long enough to produce the flatness and causal decoupling we see today (usually $\sim 40 - 60$ e-folds), and it must do so by keeping $\epsilon < 1$. This is satisfied by the second slow-roll condition:

$$\kappa \equiv \frac{\dot{\epsilon}}{H \epsilon} < 1. \quad (1.19)$$

The dynamics of inflation under the slow-roll conditions can be modeled by a single scalar field ϕ (the ‘inflaton’) evolving under the influence of a potential $V(\phi)$ ⁸. Minimizing the action for this scalar field leads to the Klein-Gordon equation (equation 1.22) in an expanding background, which leads to the following density and pressure for the inflaton field:

⁸Pick a potential and make your own new theory of inflation!

$$\rho_\phi = \frac{1}{2}\dot{\phi}^2 + V(\phi), \quad (1.20)$$

$$p_\phi = \frac{1}{2}\dot{\phi}^2 - V(\phi). \quad (1.21)$$

For $\dot{\phi}^2 \ll V(\phi)$, the energy density is dominated by the potential, yielding $p_\phi \approx -\rho_\phi$, which corresponds to $w \approx -1$ and drives exponential expansion.

The equation of motion for ϕ in an expanding FLRW background is obtained by substituting Equations 1.20 and 1.21 into the continuity equation (Equation 1.8):

$$\ddot{\phi} + 3H\dot{\phi} + V'(\phi) = 0, \quad (1.22)$$

where $V'(\phi) \equiv dV/d\phi$. The $3H\dot{\phi}$ term (commonly referred to as Hubble friction) slows the field's roll down the potential, acting like a friction term. In the slow-roll regime ($\ddot{\phi} \approx 0$ and $\dot{\phi}^2 \ll V$), the equation of motion and the Friedmann equation simplify to:

$$3H\dot{\phi} \approx -V'(\phi), \quad (1.23)$$

$$H^2 \approx \frac{V(\phi)}{3M_{\text{Pl}}^2}, \quad (1.24)$$

These are coupled equations that govern the dynamics of the inflaton field and the expansion of the Universe during inflation under slow-roll conditions. $M_{\text{Pl}} = (8\pi G)^{-1/2}$ is the reduced Planck mass.

The slow-roll conditions, now in terms of the inflaton potential are:

$$\epsilon_V \equiv \frac{M_{\text{Pl}}^2}{2} \left(\frac{V'}{V} \right)^2, \quad (1.25)$$

$$\kappa_V \equiv M_{\text{Pl}}^2 \frac{V''}{V}, \quad (1.26)$$

where primes denote derivatives with respect to ϕ . Inflation occurs as long as both $\epsilon_V < 1$ and $|\kappa_V| < 1$. Inflation ends when the rolling of the parameters increases and the inflaton reaches the minimum of the potential, all while its energy is converted into standard model particles during a process known as reheating [11].

1.2.3 Primordial Perturbations from Inflation

So far, inflation has successfully solved the horizon, flatness problems, and superhorizon correlations problem. However, the key prediction of inflation is that it provides a mechanism for generating the primordial density perturbations that seed the large-scale structure of the Universe we see today. Inflation posits that they arise from quantum vacuum fluctuations in the inflaton and spacetime metric. During inflation, these fluctuations are stretched to wavelengths exceeding the Hubble radius, until they re-enter the horizon at a later epoch. By introducing small perturbations to the FLRW metric tensor and quantizing the inflaton field as a collection of quantum harmonic oscillators, we can calculate observables such as the power spectra of scalar and tensor perturbations⁹. The tensor perturbations correspond to primordial gravitational waves (PGWs) and directly probe the energy scale of inflation. If detected at sufficient confidence, they would provide ‘smoking-gun’ evidence of inflation.

Scalar Perturbations

Per the generalization of the Helmholtz decomposition theorem to tensors, perturbations to the spacetime metric can be decomposed into scalar, vector, and tensor components (SVT

⁹Vector perturbations are also included here. However, due to the phenomenology of inflation they get completely washed out by the accelerated expansion.

decomposition) as $\mathbf{V} = \nabla\Phi + \nabla \times \mathbf{A}$ where Φ is a scalar and \mathbf{A} is divergence-free. These are commonly characterized by the gauge-invariant curvature perturbation \mathcal{R} . In the slow-roll approximation, the power spectrum of scalar perturbations is [4]:

$$P_{\mathcal{R}}(k) = \frac{1}{8\pi^2} \frac{H^2}{M_{\text{Pl}}^2 \epsilon_V} \Big|_{k=aH} \equiv 2\pi^2 A_s k^{-3} \left(\frac{k}{k_*}\right)^{n_s-1}, \quad (1.27)$$

where the quantities are evaluated at horizon crossing ($k = aH$). A_s is the scalar amplitude, n_s is the scalar spectral index, and k_* is a pivot scale (we adopt a scale of $k_* = 0.05 \text{ Mpc}^{-1}$ throughout all our publications including this dissertation).

A perfectly scale-invariant spectrum corresponds to $n_s = 1$. Slow-roll inflation generically predicts a slight ‘red tilt’, $n_s = 1 - 6\epsilon_V + 2\eta_V$, meaning that there is marginally more power on large scales than on small scales. The measured value, $n_s = 0.965 \pm 0.004$ [3], represents a detection of this tilt and is key evidence of the inflationary Universe.

Tensor Perturbations

In addition to scalar perturbations, inflation also sources tensor perturbations (PGWs) corresponding to fluctuations in the transverse, traceless components of the spacetime metric. Their power spectrum in the slow-roll approximation is [4]:

$$P_t(k) = \frac{2}{\pi^2} \frac{H^2}{M_{\text{Pl}}^2} \Big|_{k=aH} \equiv 2\pi^2 A_t k^{-3} \left(\frac{k}{k_*}\right)^{n_t}, \quad (1.28)$$

where A_t is the tensor amplitude and n_t is the tensor spectral index. A scale-invariant tensor spectrum corresponds to $n_t = 0$, and slow-roll inflation predicts $n_t = -2\epsilon_V$, a small negative tilt.

The Tensor-to-Scalar Ratio

The key observable that links the tensor and scalar perturbation spectra is the tensor-to-scalar ratio:

$$r \equiv \frac{P_t(k_*)}{P_{\mathcal{R}}(k_*)} = \frac{A_t}{A_s} = 16 \epsilon_V. \quad (1.29)$$

A non-zero detection of r would be linked to the first slow-roll parameter ϵ and would directly probe the energy scale of inflation as can be seen in equation 1.30:

$$V^{1/4} \approx 1.06 \times 10^{16} \text{ GeV} \left(\frac{r}{0.01} \right)^{1/4}. \quad (1.30)$$

The energy scale of inflation is a key unknown in cosmology, and is expected to be far above the current energy reach of particle colliders, making it a unique window into high-energy physics and the early Universe.

As we will see in the next section, tensor perturbations imprint a unique signature on the polarization of the CMB in the form of B-mode polarization. This is a direct observational avenue for measuring or constraining r , which is particularly powerful given that a big component of any theory is the ability to make *testable* predictions.

1.3 The Cosmic Microwave Background

As mentioned in the previous sections, the CMB has been a big contributor to establishing the Λ CDM model of cosmology and in constraining its parameters. It has given us detailed understanding of the composition, geometry, and dynamics of the Universe, as well as evidence of the large-scale anisotropy correlations of the primordial quantum density fluctuations. What makes the CMB such a powerful probe of the early Universe is that it captures the early-Universe conditions at their infancy, prior to the large-scale structure evolution that occurred in later epochs. The physics of the later epochs is complicated and governed by the Einstein-Boltzmann equations, which are highly non-linear¹⁰. On the contrary, the CMB and its structure was frozen at the time of recombination, a period when the Universe had undergone little evolution and can therefore be described by linear perturbation theory.

¹⁰Not to mention the high source of uncertainty in some of their parameters.

Figure 1.2 shows the temperature anisotropy map of the CMB as measured by the *Planck* satellite [12]. Simulations of the CMB temperature anisotropy field can be generated using the best-fit Λ CDM parameters and the Code for anisotropies in the Microwave Background CAMB [13]. It captures the intensity (temperature for us experimental cosmologists) variations of the CMB photons across the entire sky.




Figure 1.2. Measured temperature anisotropy map of the CMB as measured by the *Planck* satellite [12].

These variations reflect density inhomogeneities in the primordial hot plasma at the time when photons were able to free stream through space (see section 1.3.1) Figure 1.3 shows a power spectrum of the CMB temperature anisotropies as measured by the *Planck* satellite [12] together with data from the South Pole Telescope (SPT) [14] denoted as dots, and the theoretical prediction of the Λ CDM model (gray line) made using the Cosmic Linear Anisotropy Solving System (CLASS) [15].

The different regions of the power spectrum correspond to different physical processes that the primordial plasma underwent before recombination. The low- ℓ region (large angular scales) correspond to superhorizon (outside the particle horizon) scales at the time of recombination. Therefore, fluctuations on these scales did not evolve prior to photon decoupling, offering a direct probe at the initial conditions of the Universe.




Figure 1.3. CMB temperature anisotropy power spectrum as measured by the *Planck* satellite [12] (black) together with data from the South Pole Telescope (SPT) [14] (gray) made using the Cosmic Linear Anisotropy Solving System (CLASS) [15]. The solid line corresponds to the Λ CDM model prediction. Plot from [4]. The different regions correspond to regimes that the primordial plasma underwent before recombination.

The intermediate region (labeled as sound waves) forms distinctive peaks known as the ‘acoustic peaks’ and correspond to oscillations of the tightly coupled photon-baryon fluid. Unsurprisingly, it is the small and intermediate scales (degree angular scales) that we are most interested in when looking for signatures of inflation. The region of small angular scales (or high ℓ) is dominated by photon damping caused by the random diffusion.

The goal of this section is to review how we model the anisotropy and polarization spectra of the CMB in search of the characteristic signature of inflation.

1.3.1 Recombination and the Last-Scattering Surface

In the earliest times (post inflation during the radiation dominated era) after the Big Bang, the Universe was a hot, dense plasma in which photons, electrons, and baryons were tightly coupled through Compton scattering and Coulomb interactions. As the Universe expanded and cooled, it eventually reached a temperature of $T \sim 3000$ K (corresponding to a redshift of $z \approx 1100$), at which point the thermal energy of the plasma dropped below the binding energy of hydrogen. Electrons and protons then combined to form neutral hydrogen atoms in a process known as recombination¹¹.

The photon mean free path as the Universe cooled is $\lambda_{\text{mfp}} = (n_e \sigma_T)^{-1}$, where n_e is the electron number density and σ_T is the Thomson scattering cross-section (low energy limit of Compton scattering). The number density of free electrons (dropping quickly while the Universe is expanding) means that this mean free path became larger and larger until it exceeded the comoving Hubble radius. This is the definition of photons *free streaming* or *decoupling* from matter. Of course this process did not happen instantaneously. From our perspective, it gradually occurred in an expanding ‘shell’ of decoupling, with finite thickness that lasted between 50,000 – 100,000 years, the thickness of which would determine the rate of recombination. The surface of this shell from which these photons last scattered is called the last-scattering surface, and is a snapshot of the Universe at the end of recombination—the Cosmic Microwave Background.

Due to the thermal equilibrium of the primordial plasma (prior to last scattering), its spectrum can be modeled as a blackbody by the Planck formula (not working in natural units momentarily):

$$B_\nu(\nu, T) = \frac{2h\nu^3}{c^2} \frac{1}{e^{h\nu/(k_B T)} - 1}, \quad (1.31)$$

where h is the Planck constant, k_B is the Boltzmann constant, and ν is the photon frequency. The blackbody nature of the CMB was measured first by the FIRAS instrument on the *COBE*

¹¹A misnomer since nothing was combined before.

satellite [1], and the present-day temperature has been measured to be $T_0 = 2.72548 \pm 0.00057$ K by *Planck*.

The CMB was first detected by Penzias and Wilson in 1965 as an isotropic excess antenna temperature of approximately 3.5 K [16]. In a companion paper, Dicke, Peebles, Roll, and Wilkinson identified this signal as the predicted remnant radiation from the hot Big Bang. This discovery earned Penzias and Wilson the 1978 Nobel Prize in Physics.

1.3.2 Temperature Anisotropies

As discussed in section 1.2, the CMB is uniform but not perfectly homogeneous. Tiny temperature fluctuations at the level of $\Delta T/T \sim 10^{-5}$ were first detected by the *COBE* DMR experiment and have since been mapped with exquisite precision by the *WMAP* [17] and *Planck* satellite missions. These anisotropies are a direct probe of the density fluctuations in the early Universe.

Spherical Harmonic Decomposition

To directly probe the temperature anisotropies we first subtract the mean temperature T_0 from the observed temperature field $T(\hat{n})$. Then, we subtract the induced dipole anisotropy caused by the Earth's motion with respect to the CMB rest frame, which is approximately 3.36 mK. What is left is the CMB temperature anisotropies, of order $\sim 100 \mu\text{K}$.

The observed fluctuations come from a spherical surface (the last-scattering surface) and are therefore expanded in spherical harmonics $Y_{\ell m}(\theta, \varphi)$ where θ and φ correspond to a certain direction $\hat{n} = (\theta, \varphi)$ on the sky. The temperature fluctuations are then: $\Delta T(\hat{n}) = T(\hat{n}) - T_0$

$$\Delta T(\hat{n}) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^T Y_{\ell m}(\hat{n}), \quad (1.32)$$

where the coefficients $a_{\ell m}^T$ are the ‘multipole moments’, obtained by exploiting the orthogonality of the spherical harmonics as:

$$a_{\ell m}^T = \int \Delta T(\hat{n}) Y_{\ell m}^*(\hat{n}) d\Omega. \quad (1.33)$$

The multipole number ℓ characterizes the angular scale of the fluctuation, with the approximate correspondence $\theta \sim 180^\circ/\ell$. Low ℓ values correspond to large angular scales and high ℓ to small angular scales as can be seen in Figure 1.3.

The Angular Power Spectrum

Under the assumption that the primordial perturbations are drawn from a Gaussian random field (an assumption strongly supported by *Planck* although primordial non-Gaussianity searches are investigated, see for example [18]), all map information is contained in the two-point correlation function, or equivalently the angular power spectrum C_ℓ , defined by:

$$\langle a_{\ell m}^T a_{\ell' m'}^{T*} \rangle = C_\ell^{TT} \delta_{\ell\ell'} \delta_{mm'}, \quad (1.34)$$

where the angle brackets denote the theoretical expectation value. We estimate the power spectrum from the available m -modes:

$$\hat{C}_\ell^{TT} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{\ell m}^T|^2. \quad (1.35)$$

There is a finite number of m -modes at each ℓ (namely $2\ell + 1$), which imposes a fundamental lower bound on the uncertainty with which the power spectrum can be measured, known as the cosmic variance limit:

$$\frac{\Delta C_\ell}{C_\ell} = \sqrt{\frac{2}{2\ell + 1}}. \quad (1.36)$$

At low ℓ , the variance is large due to the small number of modes, available, while at high ℓ the variance decreases. Given this dependence on ℓ , we scale the power spectrum by the factor $\ell(\ell + 1)$ and plot $\mathcal{D}_\ell \equiv \ell(\ell + 1) C_\ell / (2\pi)$, resulting in a scale-invariant temperature

power spectrum at low ℓ .

Acoustic Oscillations and Parameter Constraints

The oscillatory structure seen in Figure 1.3 captures the dynamics of the Universe prior to recombination. The two forces at play are gravity, which pulls matter inward, and radiation pressure pushing outward, driving oscillations known as baryon acoustic oscillations (BAO). During last scattering, when photons are able to free stream, these oscillations get baked in the photon temperature field, with some modes caught at maximum compression and others at minimum, resulting in the pattern of hot and cold seen in Figure 1.2. These produce the characteristic peaks in the power spectrum.

The angular position of the first acoustic peak (at $\ell \approx 220$) captures the maximum distance that the photon-baryon fluid waves could have traveled at recombination. This distance is also known as the sound horizon at recombination, since it is reminiscent of the particle horizon but for sound waves in the plasma. The sound horizon can be related to the geometry of the Universe at last scattering, and provides strong evidence for spatial flatness [19]. The first peak corresponds to maximum compression of the photon-baryon fluid, therefore the second peak corresponds to minimum compression or ‘rarefaction’. Each even peak corresponds to a rarefaction and each odd peak corresponds to a compression, with the peaks getting smaller and smaller as we go to higher ℓ due to Silk damping (photon diffusion) [20] during last scattering.

The height difference between successive peaks tells us about the composition of the Universe, specifically about the ratio of baryons to photons. Since we know the photon density from the CMB temperature, the relative heights of the peaks can be used to infer the baryon density. In turn, knowing the baryon density and the total matter density, we can infer the dark matter density [21].

Theoretical predictions for the CMB spectrum are generated by solving the coupled Einstein-Boltzmann system of equations, which describe the evolution of perturbations in the photon, baryon, dark matter, and neutrino fluids. Publicly available codes such as CAMB [13] and

CLASS [15] perform this calculation efficiently, enabling precise comparisons between theoretical models and data.

1.3.3 CMB Polarization

So far, we have only discussed the temperature anisotropies of the CMB. Light, however, is an electromagnetic wave with both an intensity and polarization properties. What cosmological information is contained in the polarization of the CMB? In this section, we try to answer this question by starting with an intuitive description of how polarization arises in the CMB. Then, we proceed with the standard mathematical description of the CMB polarization field, exploring its patterns and how they relate to the underlying physics of the early Universe and inflation.

In general, scattering from a particular direction produces a net polarization in the scattered radiation. This is easily seen by the fact that the direction of the electric field of the scattered radiation has to be perpendicular to the direction of propagation. If the direction of scattering is at 90° with respect to the incident radiation, then the scattered radiation will be linearly polarized. However, if the incident radiation is isotropic (monopole) or has only a dipole pattern, the scattered radiation will be unpolarized. A quadrupole anisotropy, however, produces differential intensities along orthogonal directions, yielding a net linear polarization in the scattered photon. Via this mechanism polarization in the CMB arises from Thomson scattering of photons off free electrons in the presence of a local quadrupole temperature anisotropy. The quadrupole anisotropy can be thought of as ‘hot’ and ‘cold’ spots in the temperature field. The scattered radiation will be polarized in the direction of the hot spot and perpendicular to the cold spot.

Prior to recombination, this polarization is erased by the tight coupling between photons and baryons, which suppresses the overall quadrupole anisotropy. The CMB polarization that we observe today is generated during last scattering. Since the number density of free electrons drops during recombination, there is a brief window during which polarization is generated. Therefore, CMB polarization is a snapshot of the dynamics of the plasma at the surface of

last scattering, and, as we will see in this section, it would capture the unique signature of inflationary gravitational waves.

Stokes Parameters

We use a common parameterization of an electromagnetic wave propagating in the \hat{z} direction described by the Stokes parameters [22]. For electric field components $E_x = a_x \cos(\omega t - \xi_x)$ and $E_y = a_y \cos(\omega t - \xi_y)$, the Stokes parameters are defined as:

$$\begin{aligned}
 I &= \langle a_x^2 \rangle + \langle a_y^2 \rangle, \\
 Q &= \langle a_x^2 \rangle - \langle a_y^2 \rangle, \\
 U &= 2\langle a_x a_y \rangle \cos(\xi_x - \xi_y), \\
 V &= 2\langle a_x a_y \rangle \sin(\xi_x - \xi_y),
 \end{aligned} \tag{1.37}$$

where I is the total intensity, Q and U describe the two orthogonal states of linear polarization, and V describes circular polarization. Thomson scattering cannot produce circular polarization, so $V = 0$ for the CMB. Stokes Q and U are coordinate-dependent. Under a rotation of the coordinate axes by an angle ψ , they transform as a spin-2 field:

$$(Q \pm iU)' = e^{\mp 2i\psi} (Q \pm iU). \tag{1.38}$$

E-Mode and B-Mode Decomposition

To get around the coordinate dependence of Q and U , we decompose the polarization field into coordinate-independent quantities. To develop intuition, on how these new parameterization should be defined, we expand the polarization in spin-weighted spherical harmonics ${}_{\pm 2}Y_{\ell m}$:¹².

¹²The scalar spherical harmonics used to describe the temperature anisotropies, cannot fully describe a spin-2 field.

$$(Q \pm iU)(\hat{n}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} {}_{\pm 2}a_{\ell m} {}_{\pm 2}Y_{\ell m}(\hat{n}). \quad (1.39)$$

From the spin-2 harmonic coefficients, one constructs the coordinate-independent E-mode and B-mode coefficients ¹³:

$$a_{\ell m}^E \equiv -\frac{1}{2} ({}_{+2}a_{\ell m} + {}_{-2}a_{\ell m}), \quad (1.40)$$

$$a_{\ell m}^B \equiv -\frac{i}{2} ({}_{+2}a_{\ell m} - {}_{-2}a_{\ell m}). \quad (1.41)$$

The E-mode component is a parity-even scalar field whose polarization pattern is curl-free, analogous to an electrostatic field. The B-mode component is a parity-odd pseudo-scalar field exhibiting a divergence-free, curl-like pattern, analogous to a magnetic field. Under a parity transformation, E-modes remain unchanged while B-modes change sign. Figure 1.4 shows a visualization of the E-mode and B-mode patterns for a radial wave.




Figure 1.4. Visualization of the E-mode and B-mode polarization patterns for a radial wave in the xy plane with direction of photon propagation coming out of the page. Figure from [4].

This decomposition is crucial because different physical processes generate parity-even or parity-odd (or both) polarization, as we describe in the next section.

¹³This can be verified by applying a rotation to Stokes Q and U and calculating the resulting transformation of the E-mode and B-mode coefficients.

Angular Power Spectra of Polarization

With three fields defined on the sky: T , E , and B , we calculate six total angular power spectra by taking the two-point correlation of the spherical harmonic coefficients:

$$C_\ell^{XY} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m}^{X*} a_{\ell m}^Y, \quad X, Y \in \{T, E, B\}. \quad (1.42)$$

In a parity-conserving Universe, any cross-correlation involving one scalar and one pseudo-scalar field must vanish, so $C_\ell^{TB} = C_\ell^{EB} = 0$. The four non-vanishing spectra are C_ℓ^{TT} , C_ℓ^{EE} , C_ℓ^{BB} , and C_ℓ^{TE} .

1.3.4 Sources of E-Mode and B-Mode Polarization

The physical distinction between E-modes and B-modes is of fundamental importance because different physical processes generate different combinations of E-modes and B-modes. Specifically, scalar perturbations (density fluctuations) in the photon-baryon fluid can only generate E-mode polarization, while tensor perturbations (gravitational waves) can generate both E-mode and B-mode polarization.

E-Modes from Scalar Perturbations

As discussed at the beginning of this section, only quadrupole temperature anisotropies can generate polarization. Scalar (density) perturbations in the photon-baryon fluid create local quadrupole temperature anisotropy related to the velocity gradients of the fluid. Expanding the spatial dependence of the polarization anisotropy field in Fourier modes $\Theta(\hat{\mathbf{n}}, \mathbf{k})$ and express in terms of Legendre polynomials $P_\ell(\mu')$ gives:

$$\Theta(k, \mu') = \sum_{\ell} i^\ell (2\ell + 1) \Theta_\ell(k) P_\ell(\mu'), \quad (1.43)$$

where $\mu' \equiv \hat{k} \cdot \hat{n}$ represents the angle between the Fourier wavevector \mathbf{k} and the line-of-

sight direction \hat{n} , and $\Theta_\ell(k)$ are the multipole moments. Only the $\ell = 2$ term in this expansion contributes to the quadrupole anisotropy denoted by $\Theta_2(k)$, which is the source of polarization. By solving the Thomson scattering problem for a single Fourier mode in the flat sky, and projecting the resulting polarization onto the E-mode and B-mode basis, it can be shown that [23]:

$$E(\hat{n}, \mathbf{k}) \propto \frac{4\pi}{5} \left[1 - (\hat{n} \cdot \hat{k})^2 \right] \Theta_2(k), \quad (1.44)$$

$$B(\hat{n}, \mathbf{k}) \propto 0. \quad (1.45)$$

Therefore, scalar perturbations produce only E-mode polarization.

The E-mode signal was first detected by the DASI experiment [24], roughly four decades after the initial detection of the CMB.

B-Modes from Tensor Perturbations

The presence of gravitational waves at last scattering sourced by inflation would produce a unique signature in the CMB polarization field and its form can be determined by the geodesic equation¹⁴ for photons in the perturbed spacetime metric. Integrating the geodesic equation along the line of sight $\Theta^t(\hat{\mathbf{n}})$ from the last-scattering surface to today gives the theory expectation for the temperature anisotropy power spectrum exclusively due to inflationary gravitational waves. A tensor perturbation h_{ij} ¹⁵ defined as:

$$h_{ij} = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (1.46)$$

¹⁴The geodesic equation determines how inhomogeneities in the spacetime affect the photon energy, not just from redshift but also due to the expansion of the Universe. See section 7.2 in [4].

¹⁵The perturbations themselves are also functions of the spacetime coordinates.

to the FLRW metric at inflation can be represented as (see equation 1.2):

$$ds^2 = a^2(\eta) [-d\eta^2 + (\delta_{ij} + h_{ij}) dx^i dx^j], \quad (1.47)$$

where $a(\eta)$ is the scale factor, will generate a quadrupole anisotropy in the photon temperature field at last scattering. The presence of this anisotropy affects the free streaming of photons from last scattering to today. The line of sight solution to the geodesic equation is given by:

$$\Theta^{(t)}(\hat{n}) = -\frac{1}{2} \int_{\eta_*}^{\eta_0} d\eta h'_{ij}(\eta, \mathbf{x} = \eta \hat{n}) \hat{n}^i \hat{n}^j, \quad (1.48)$$

where:

- the superscript (t) denotes tensor perturbations, η is conformal time, with η_0 the conformal time today and η_* the conformal time at last scattering (the surface of last scattering),
- h_{ij} is the transverse and traceless perturbation to the metric tensor at inflation and prime denotes a derivative with respect to conformal time $h'_{ij} \equiv \partial h_{ij} / \partial \eta$,
- \hat{n}^i is the unit vector along the line of sight in comoving coordinates, and the product $\hat{n}^i \hat{n}^j$ projects the tensor onto the photon propagation direction,
- the argument $\mathbf{x} = \eta \hat{n}$ indicates evaluation of the perturbation along the unperturbed photon path ¹⁶.

Linearizing the geodesic equation for photons yields the conformal time derivative of the $\Theta^{(t)}(\hat{n})$:

$$\frac{d\Theta^{(t)}(\hat{n})}{d\eta} = -\frac{1}{2} \sin^2 \theta' \left(\frac{\partial h_+}{\partial \eta} \cos 2\phi' + \frac{\partial h_\times}{\partial \eta} \sin 2\phi' \right), \quad (1.49)$$

¹⁶The factor of $-1/2$ follows from linearizing (keeping only the first-order terms) the projection of the metric perturbation in the geodesic equation. See problem 7.6 in [4].

where the direction of the incoming photon \hat{n} is written in terms of comoving spherical coordinates as $\hat{n} = (\sin \theta' \cos \phi', \sin \theta' \sin \phi', \cos \theta')$. Equation 1.49 shows the angular dependence of the gravitational wave contribution that induces the quadrupole anisotropy (gravitational waves are anisotropic by definition). In other words, along the expanding direction photons are colder and along the contracting direction photons are hotter, due to blue/redshift. Since tensor fluctuations correspond to spatial distortions along *two* directions, they fundamentally source both E-mode and B-mode polarization. This is directly reflected in equation 1.49 as the $\cos 2\phi'$ term can be shown to source E-modes and the $\sin 2\phi'$ to source B-modes. The predicted tensor power spectrum $C_\ell^{(t)}$ is then [4]:

$$C_\ell^{(t)} = 4\pi \int d \ln k \left| \Theta_\ell^{(t)}(k) \right|^2 P_t(k), \quad (1.50)$$

where $P_t(k)$ is the primordial tensor power spectrum (equation 1.28) and the integral is over all Fourier modes k contributing to angular scale ℓ .

1.3.5 Gravitational Lensing of the CMB

Although scalar perturbations do not source B-modes at last scattering, they can appear as B-modes due to gravitational lensing. As the CMB photons traverse the large scale structure of the Universe on their way to Earth, their E-mode patterns get deflected and made to look like B-modes. These are ‘lensing B-modes’ and correspond to a spectrum that peaks at $\ell \sim 1000$ [25] in D_ℓ^{BB} , which corresponds to smaller scales than the expected peak of the inflationary B-mode amplitude. This weak gravitational lensing remaps the CMB temperature and polarization fields according to:

$$\tilde{X}(\hat{n}) = X(\hat{n} + \nabla\phi(\hat{n})), \quad (1.51)$$

where $X \in \{T, Q, U\}$, \tilde{X} is the lensed field, and $\phi(\hat{n})$ is the lensing potential projected along the line of sight. Lensing B-modes were first detected by the SPT [26]. Using the CMB itself

or external measurements of large scale structure and its distribution we can estimate the lensing potential and partially undo the lensing of E-modes. This process is called ‘delensing’ and is a necessary part of current and future precision measurements of the CMB polarization that aim to detect inflationary B-modes. Figure 1.5 shows the theoretical power spectra of TT, EE, tensor BB associated with $r = 0.1$, and lensing BB.

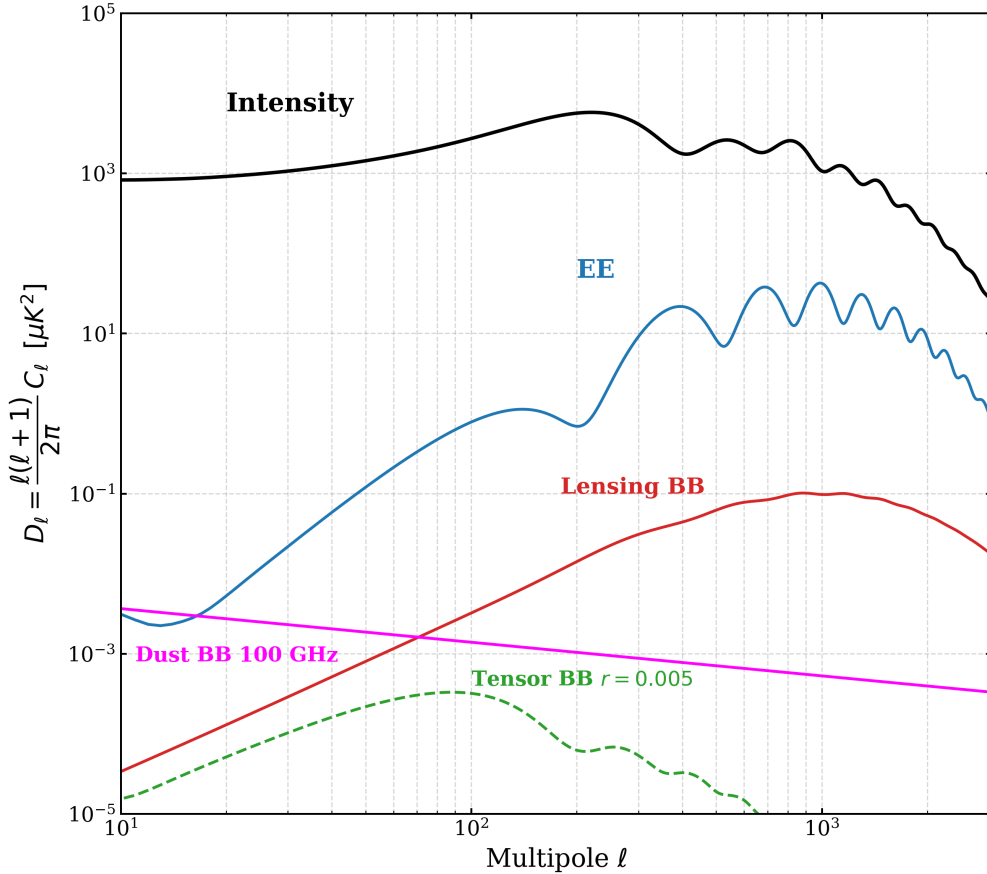


Figure 1.5. Theoretical CMB power spectra of TT, EE, tensor BB with $r = 0.005$, lensing BB, and spectrum of galactic dust B-modes at 100 GHz, all plotted as a function of multipole ℓ . Measurements of the tensor BB spectrum are extremely challenging due to the faint signal and the presence of lensing B-modes and foregrounds.

Figure 1.5 illustrates the mounting difficulty in detecting inflationary B-modes. At the current stage of observations, the limiting factor is not simply instrumental sensitivity (i.e. pushing down measurement noise by adding more detectors and more observation time). Rather, one of the main challenges is quantifying and removing lensing B-modes. In other words, cosmological B-mode detection is lensing-limited. Therefore, the correct approach when looking for the inflationary B-mode signal is to look for an excess of B-mode power in the measured BB power spectrum beyond the lensing contribution, and (as we will see in

section 1.4) beyond the contribution from foregrounds.

1.3.6 Current Observational Constraints on r

Primordial gravitational wave detection associated with inflation has been the hot topic in experimental cosmology for decades. Early constraints on the tensor-to-scalar ratio were derived from the temperature power spectrum, since gravitational waves add power at large angular scales in C_ℓ^{TT} through the Sachs-Wolfe effect. The *Planck* constraint from TT and low- ℓ temperature spectrum was $r_{0.002} < 0.10$ at 95% confidence [12]. This limit is not substantially improved further due to cosmic variance and difficulty constraining foregrounds accurately at large scales.

The most sensitive constraints on r come from dedicated B-mode polarization measurements at degree angular scales from the BICEP/Keck program. Using a series of small-aperture polarimeters operating from the South Pole, and including all data through the 2018 observing season (BK18) across 95, 150, and 220 GHz and along with joint analysis of the *Planck* and WMAP data, yields an upper limit on the tensor-to-scalar ratio of:

$$r_{0.05} < 0.036 \quad (95\% \text{ confidence}), \quad (1.52)$$

$$\sigma(r) = 0.009. \quad (1.53)$$

where $\sigma(r)$ is the uncertainty on the upper limit on r from simulations (Section 2.7). In the context of assessing instrumental sensitivity and performance, $\sigma(r)$ is the more relevant quantity to consider than the upper limit itself. This work is focused on probing how systematic effects evolve as we push $\sigma(r)$ down for our upcoming publication—BK24. This constraint rules out several prominent inflationary models, including monomial potentials like $V(\phi) \propto \phi^2$, and natural inflation (purple region in Figure 1.6) at greater than 95% confidence. Ongoing and planned observations with BICEP3 and the BICEP Array are projected to reach $\sigma(r)$ of a factor of 5-8 lower than the BK18 constraint. Figure 1.6 summarizes the

r constraints from BK18 and previous results from *Planck*.




Figure 1.6. Constraint on the tensor-to-scalar ratio r from the BK18 analysis (blue) overplotted with previous results from Planck (green) in parameter space of the scalar spectral index n_s and the tensor-to-scalar ratio r (at pivot scale of $k = 0.002 \text{ Mpc}^{-1}$). The upper limit on r from BK18 is $r_{0.05} < 0.036$ at 95% confidence, which strongly disfavours prominent inflationary models such as $V(\phi) \propto \phi$, $V(\phi) \propto \phi^2$, and natural inflation (purple). The orange region corresponds to the range of e-folds $N_e = 50 - 60$ of exponential expansion needed for inflation to explain the observed flatness and homogeneity of the Universe. The convex/concave regions distinguish between different classes of inflationary potentials. Figure from [27], courtesy of the BICEP/Keck collaboration.

1.4 Foregrounds

Lensing B-modes are not the only source of unwanted B-mode polarization we worry about. Anything that produces B-mode polarized emission between the last-scattering surface and the Earth can appear as inflationary B-modes. The two main sources of foreground contamination are Galactic dust and synchrotron radiation from relativistic electrons in the Galactic magnetic field. Although relevant at degree angular scales, these foregrounds are expected

to have a different frequency dependence than the CMB, which allows us to separate them using multi-frequency observations.

In this section, we briefly describe the origin and characteristics of these foregrounds, as well as foreground separation techniques. A more detailed discussion of foregrounds relevant to the BICEP/Keck program can be found in [28].

1.4.1 Thermal Dust Emission

Dust grains in the interstellar medium (ISM) are typically non-spherical and have temperatures of $T_d \sim 19.6$ K [29]. These grains tend to align such that their long axes are perpendicular to the local Galactic magnetic field, producing thermal emission that is partially polarized along the field direction. The projection of this magnetic field direction onto the plane of the sky determines the observed dust polarization. The frequency dependence of the polarized dust intensity is well described by a modified blackbody spectrum:

$$I_d(\nu) \propto \nu^{\beta_d} B_\nu(T_d), \quad (1.54)$$

where $B_\nu(T_d)$ is from equation 1.31 evaluated at the dust temperature T_d , and β_d is the dust emissivity spectral index. Full-sky observations by *Planck* constrain the dust spectral index to $\beta_d = 1.53 \pm 0.02$ in polarization.

At frequencies above ~ 100 GHz, thermal dust emission is the dominant polarized foreground. The amplitude of dust emission varies significantly across the sky, being the strongest near the Galactic plane and weakest at high Galactic latitudes. We target one of the cleanest regions of the Southern sky, but even in this field there is dust emission as can be seen in Figure 2.10. Accurate characterization and subtraction of the dust contribution is essential and achieved through multi-frequency observations.

1.4.2 Synchrotron Radiation

Cosmic-ray electrons and other charged relativistic particles spiraling in the Galactic magnetic field emit polarized synchrotron radiation. The synchrotron intensity follows a power-law spectrum:

$$I_s(\nu) \propto \nu^{\beta_s}, \quad (1.55)$$

where β_s is the synchrotron spectral index. Observations indicate a mean polarization spectral index of $\beta_s \approx -3.1$ [30], with some variation across the sky due to differences in the cosmic-ray electron energy spectrum and magnetic field configuration.

Synchrotron emission dominates the polarized foreground at frequencies below ~ 70 GHz and decreases rapidly at higher frequencies, making it subdominant to dust in the 90-270 GHz range. Inclusion of lower-frequency data (e.g., from WMAP at 23 GHz and BICEP Array at 30 and 40 GHz) in multi-component analyses helps constrain the synchrotron contribution and reduce its impact on the final r constraint.

1.4.3 Multi-Frequency Foreground Separation

The key to disentangling the CMB signal from Galactic foregrounds lies in their distinct frequency dependencies. The CMB follows a perfect blackbody spectrum (Equation 1.31) and is therefore the same brightness in thermodynamic temperature units at all frequencies. Thermal dust emission rises steeply with frequency above ~ 100 GHz, while synchrotron emission falls steeply at frequencies above ~ 30 GHz. By observing at multiple frequencies we can exploit this spectral dependence to separate the three components.

The BICEP/Keck program employs observations at 30, 40, 95, 150, 220, and 270 GHz across BICEP3 and BICEP Array instruments, covering a broad frequency range that covers both foregrounds. These multi-frequency maps are combined within a parametric likelihood framework (see section 2.8) that simultaneously fits for r , dust, and synchrotron amplitudes and

spectral parameters, yielding robust constraints on r that are marginalized over foreground uncertainties.

1.5 Thesis Outline

This thesis presents work on the calibration, beam characterization, and systematic error assessment of the BICEP/Keck series of CMB polarimeters, all in the pursuit of improving constraints on the tensor-to-scalar ratio r . The remainder of this thesis is organized as follows:

Chapter 2 provides an overview of the BICEP/Keck experimental program, describing the instrument design, detector technology, observing strategy, and the series of telescopes deployed at the South Pole from BICEP1 through BICEP Array. It ends with a description of our CMB data analysis pipeline (sometimes referred to as ‘mainline pipeline’), along with a brief description of the multi-component likelihood framework used to separate the CMB signal from foregrounds and derive constraints on r .

Chapter 3 describes the suite of calibration measurements performed on the BICEP/Keck instruments, including the various setups and techniques used to characterize the beam response of the telescopes. It concludes with a description of the beam analysis pipeline, used (with small modifications) for all beam characterization results presented in this thesis.

Chapter 4 presents the beam characterization results for the BICEP/Keck telescopes, detailing the far-field beam mapping campaigns, further analysis steps, and the resulting beam parameters that feed into the systematic error assessment of our instruments. It also includes a detailed description of the pair-difference beams, used to quantify potential sources of systematic contamination.

Chapter 5 focuses on systematics studies performed with a linearly polarized calibration source for BICEP3, directly quantifying the level of cross-polar response in the BICEP3 detectors and the resulting systematic contamination in the B-mode power spectrum in the form of polarization-to-polarization leakage.

Chapter 6 presents direct, multi-frequency calibration measurements of the ground emission at the South Pole site. The end goal is to quantify the level of potential azimuth-synchronous (or azimuth-fixed) contamination in CMB data, and characterize the temporal fluctuations of the ground emission.

Appendix A describes direct calibration measurements of the time it takes for our 30 and 40 GHz detectors to respond to changes in a calibration signal, referred to as the ‘time constant’ of the detector.

Appendix B showcases results from one of the exploratory analyses using our far-field calibration data to uncover unexpected optical features coupling to our 150 GHz detectors. These are found to be related to reflections from the tile frames that hold the detector tiles in place in the focal plane of the receiver.

Appendix C summarizes the noise modeling work used to quantify the noise in our ground data.

Chapter 2

BICEP/Keck Experiments

2.1 Program History

The first generation of the BICEP/Keck experiments was BICEP1, which observed from the geographic South Pole from 2006 to 2008. It was the first instrument to place constraints on $r_{0.05} < 0.70$ at 95% confidence [31]. Equipped with neutron transmutation-doped germanium (NTD Ge) thermistors, BICEP1 had a total of ~ 100 detectors at 100, 150, and 220 GHz that were optically coupled to the sky through feedhorns.

The next generation of the program was BICEP2 and Keck Array, which operated from 2010 to 2012 and 2012 to 2019 respectively. Both instruments utilized a completely new detector design that utilizes transition-edge sensor (TES) bolometers (see section 2.3). The detectors consist of phased-array antenna-coupled TES bolometers [32], cooled to mK temperatures. BICEP2, hosted at the Dark Sector Laboratory (DSL), was a single receiver at 150 GHz with ~ 500 detectors, observing through a 260 mm aperture. It was the first receiver to make a detection of degree-scale B -modes consistent with galactic dust foregrounds [33] and place a constraint of $r_{0.05} < 0.12$ at 95% confidence in combination with Planck data [2].

Keck Array, the successor to BICEP2, consisted of five BICEP2-class receivers hosted at the Martin A. Pomerantz Observatory (MAPO) at the South Pole, observing at 95, 150, 220, and 270 GHz with a total of ~ 2500 detectors. It utilized the same optical design as

2. BICEP/Keck Experiments

BICEP2 and its main goal was to expand multi-frequency observations to improve foreground separation (section 1.4). The historical evolution of the BICEP/Keck program is illustrated in Figure 2.1.

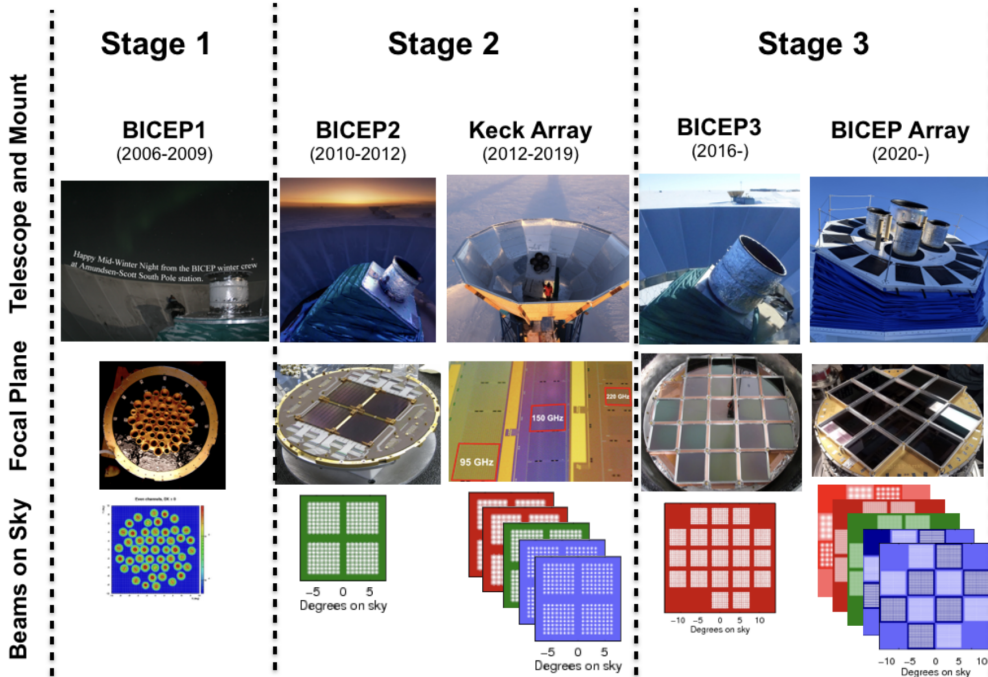


Figure 2.1. Historical evolution of the BICEP/Keck program. The top panel shows photographs of the instruments’ exterior, the middle panel shows their focal planes, and the bottom panel shows the corresponding projections onto the sky, color-coded by frequency.

In this section we review the design and observing strategy of the third generation experiments (stage-3) hosting $\mathcal{O}(10^4)$ detectors BICEP3 and BICEP Array (BA), which are the instruments used in the calibration measurements and systematics studies presented in this thesis. Throughout the rest of this thesis, the abbreviation ‘BA’ will refer to the entire instrument, while the notation ‘BAXXX’ will refer to an individual observing frequency of the instrument (e.g. BA150 refers to the 150 GHz receiver of BICEP Array). Table 2.1 summarizes the history of the stage-3 BICEP/Keck receivers era and their corresponding frequencies of operation.

Year	BICEP3	BA 30/40	BA 150	BA 220/270
2015	[95 GHz]			
2016	95 GHz			
2017	95 GHz			
2018	95 GHz			
2019	95 GHz			
2020	95 GHz	30/40 GHz		
2021	95 GHz	30/40 GHz		
2022	95 GHz	30/40 GHz		
2023	95 GHz	30/40 GHz	(150 GHz)	
2024	95 GHz	30/40 GHz	(150 GHz)	
2025	95 GHz	30/40 GHz	150 GHz	(220/270 GHz)
2026	95 GHz	30/40 GHz	150 GHz	(220/270 GHz)

Table 2.1. Receiver history of the BICEP/Keck program during the stage-3 era. Receivers in brackets denote years not included in CMB analysis, while parentheses denote a partially populated focal plane.

2.2 Instrument Overview

BICEP3 and BICEP Array are compact, on-axis refracting telescopes designed to maximize sensitivity to degree-scale B -mode polarization of the CMB. They share a common design philosophy inherited (and improved) from their predecessors, BICEP2 and Keck Array. They are small-aperture cold refractors with antenna-coupled transition-edge sensor (TES) bolometers mounted on three-axis mounts to allow for motion in azimuth (Az), elevation (El), and boresight rotation (“deck” or Dk). The frequencies of operation for BICEP3 and BICEP Array are 95 GHz and 30/40, 95, 150, and 220/270 GHz respectively. Figure 2.2 shows the bandpasses of the stage 3-era receivers with the atmospheric transmission at the South Pole as predicted by the AM atmospheric modeling software [34].

In this section we give a brief description of the unique features of each instrument in their individual subsections, and conclude with a discussion of their shared features.

2.2.1 BICEP3

BICEP3 is a single on-axis refracting telescope with a 520 mm aperture, a 27.4° instantaneous field of view, and $f/1.6$ optics. It observes at a center frequency of 95 GHz from the Dark Sector Laboratory (DSL) at the Amundsen-Scott South Pole Station, where it has operated since 2015. Its optical throughput is equivalent to all five Keck Array receivers combined,

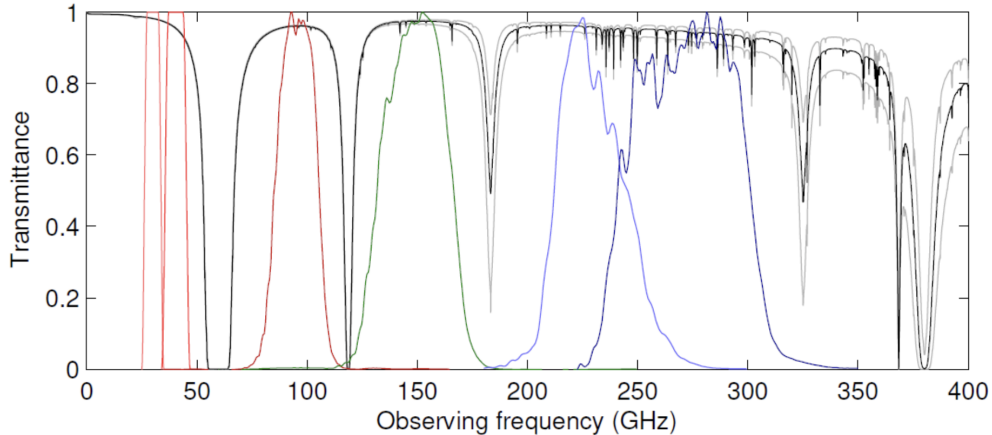


Figure 2.2. Bandpasses of BICEP3 and BA receivers with the atmospheric transmission at the South Pole as predicted by the AM atmospheric modeling software [34]. The bandpasses are normalized to their peak value and color-coded by frequency. The atmospheric transmission is shown in black and various levels of precipitable water vapor (PWV) are shown in gray.

representing a large increase in per-receiver mapping speed [35].

Light enters the instrument through (as of 2023) an extremely thin composite polymer window [36] with an anti-reflection coating. A stack of Zotefoam HD30 infrared-blocking filters at the 50 K stage, together with a 10 mm alumina ceramic filter, reduce thermal loading before the light couples to the cold optics. Two alumina refractive lenses cooled to 4 K focus different sky directions onto specific locations on the focal plane as seen in Figure 2.3.

Additional nylon IR-blocking filters are placed between and below the lenses. A metal-mesh low-pass edge filter [37] with cutoff at 120 GHz controls out-of-band response. The optics chain is housed inside a three-layer aluminum cryostat. A Cryomech PT415 pulsetube cryocooler provides continuous cooling to the 50 K and 4 K stages. A three-stage ($^4\text{He}/^3\text{He}/^3\text{He}$) helium sorption refrigerator [38] further cools the focal plane to its ~ 280 mK operating temperature. The mount partially protrudes from the roof of DSL, allowing the telescope to slew in elevation. BICEP3 was installed in BICEP2’s former mount (after some modifications), which provides full azimuth rotation capability and elevation range from 48° to 110° . The pulsetube cryocooler and surrounding accessories (colloquially known as the ‘dog house’ DH) are mounted at the base of the cryostat, which prevents the full 360° of Dk rotation as shown in Figure 2.4. However, it achieves the necessary coverage for all our CMB observing angles and their 180° complements. The limited Dk range of BICEP3 is something we have to work




Figure 2.3. Schematic of the optical path through the BICEP3 instrument. The radial directions on the sky are colored. The temperature of the various optical elements and filters is noted on the left. Image from BICEP/Keck collaboration [35].

around during our calibration measurements, but only has a minor impact on our calibration data products as discussed in section 4.2.

The focal plane consists of 20 detector modules (“tiles”) arranged in a 5-by-5 grid with the corners removed. Each tile houses 64 orthogonal, co-located, detector pairs, yielding a total of 2560 optically active detectors. A unique design choice for BICEP3 is the choice to orient 1/4 of the detector tiles at 90° relative to the others, to provide additional modulation of potential systematics that depend on the focal plane orientation and detector polarization angles. The relative tile orientations are shown in Figure 2.5.




Figure 2.4. Photograph of the BICEP3 cryostat from BICEP/Keck collaboration [35]. The entire aluminum shell is held under vacuum with subsequent shells cooled to 50 K and 4 K by the pulsetube cryocooler. Various IR filters ensure smooth transitions from ambient temperature to the ~ 280 mK focal plane.

2.2.2 BICEP Array

BICEP Array is the Stage-3 successor to Keck Array and represents the program’s current flagship instrument [39]. It employs an array of up to four BICEP3-class receivers observing at complementary frequency bands spanning 30/40, 150, and 220/270 GHz (grouped by their corresponding receiver, meaning 30/40 GHz is one receiver, while 150 GHz is a receiver that only houses 150 GHz detectors), thereby greatly expanding multi-frequency capability for foreground separation. The instrument is housed in a brand-new mount assembled *in-situ* at the Martin A. Pomerantz Observatory (MAPO), where it replaced Keck Array during the 2019–2020 Austral summer deployment. The new custom mount allows for full Dk and Az rotation with new helium rotary joints that allow for continuous rotation without the need for a cable wrap. This is a major improvement over BICEP3, eliminating any concerns of limited Dk coverage during calibration measurements.

Each BICEP Array receiver is an independent on-axis refracting telescope with a 550 mm aperture (slightly larger than BICEP3’s 520 mm), $f/1.6$ optics, and approximately a 29.6° field of view. The receivers share the same fundamental optical architecture as BICEP3:




Figure 2.5. Schematic of the BICEP3 focal plane showing the orientations of the 20 detector tiles from BICEP/Keck collaboration [35]. Notice that tiles read out by MCE0 (red) are rotated by 90° relative to the other tiles. The tile in position 21 (bottom left) is covered with a reflective material and it is not optically active, something that leads to increased sidelobe response for detectors exactly across from it. (see section 5.2.1 for more details). PT stands for pulsetube and the various number designations in each tile refer to its readout row/column as well as orderings of the general control program (GCP).

an HDPE vacuum window, Zotefoam IR-blocking filter stacks and an alumina IR filter at 50 K, two cryogenically cooled refractive lenses at 4 K (HDPE for BICEP Array, as opposed to the alumina used in BICEP3), nylon IR-blocking filters, and metal-mesh low-pass edge filters above the focal plane. Each receiver is independently cryogenically cooled via its own Cryomech pulsetube cryocooler and a three-stage helium sorption refrigerator to reach operating temperatures of ~ 250 mK at the focal plane. An annotated cutout from [38] is shown in Figure 2.6.

The receivers sit on a common three-axis mount that provides azimuth, elevation, and bore-sight rotation capability. The receivers are arranged on the Dk stage at 90° increments, and sit at a radius of ~ 1.0 m from the center of the drum. This will become important for our calibration measurements since we will have to account for parallax effects when mapping terrestrial sources a finite distance away from the telescope (see section 3.4 for a more in depth discussion and section 2.6 for a schematic of the BA focal planes projected onto the sky).

The focal planes of BA receivers consist of 12 detector modules arranged in a 4-by-4 grid with the corners removed. The number of detectors per module varies by frequency, as the

2. BICEP/Keck Experiments

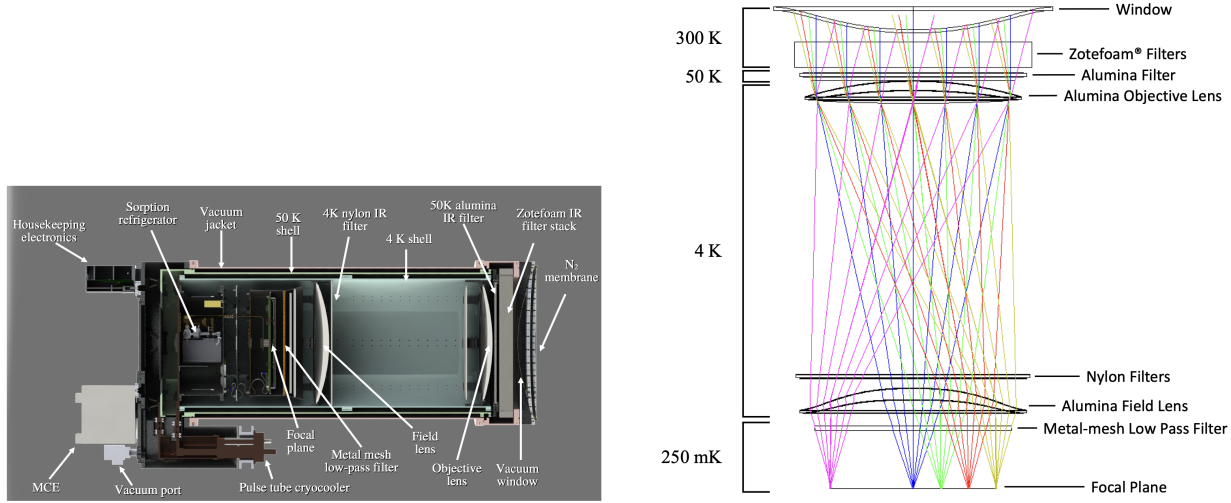


Figure 2.6. (Left) Cutout of the BICEP Array cryostat from [38] meant to illustrate the internal architecture of the optics and cryogenic components. The four aluminum shells are held under vacuum with subsequent shells cooled to 50 K and 4 K by the pulsetube cryocooler. Various IR filters ensure smooth transitions from ambient temperature to the ~ 250 mK focal plane. (Right) Schematic of the optical rays through the BICEP Array optics.

limiting factor is the density of the antenna array required to achieve the desired beam size at each frequency. For example, BA150 and BA220/270 have capacity for 7776 optically active detectors, while BA30/40 has roughly 1/10 of that capacity. Figure 2.7 shows a schematic of the BA30/40 focal plane.

BICEP Array’s first receiver, BA30/40 was deployed in January 2020, with seven 40 GHz modules, four 30 GHz modules and 1 dichroic module housing both frequencies. Since then, various tile swaps and upgrades have been made to the focal plane arrangement with tiles moved to new locations while others replaced with new better performing ones. The most notable change was inserting 45 GHz and 120 GHz low-pass Cardiff filters [37] in 2022 to control out-of-band response. BA30/40 is still observing with the goal of placing constraints on the synchrotron foreground. The second receiver, BA150, was deployed in January 2023 with only 5 out of 12 detector modules installed. The following year five more modules were installed and in 2025 the final two modules fully populated the focal plane. Given the enormous optical throughput of BA150 (equivalent to more than 3 times all of Keck Array), it is expected to quickly become a major contributor to the overall sensitivity. Hence a big portion of the optical characterization of this thesis shows results from BA150 (see sections 4.1 and 6.2). The third receiver, BA220/270, was deployed in January 2025 with




Figure 2.7. Schematic of the BA30/40 focal plane. Both BA150 and BA220/270 have the same layout but with different antenna array architecture that reflects the observing frequency. Plot from [39].

three 220 GHz modules. The following year three more 220 GHz modules were added along with one 270 GHz module. The goal is to fully populate BA220/270 by 2027 and use it to place constraints on the dust foreground.

2.3 Detectors and Readout

BICEP3 and BICEP Array both utilize antenna-coupled TES bolometers detector technology. The modular focal plane of BICEP3 houses 20 modules of size $4'' \times 4''$, while that of BICEP Array is composed of 12 modules of size $6'' \times 6''$ that allows for denser packing. Detailed descriptions of the detector design and fabrication can be found in [40] and [38] for BICEP3 and BICEP Array respectively. For a closer look at the frame design of the BA150 modules see appendix B. The difference between the $4''$ and $6''$ wafers is shown in Figure 2.8.

The detectors are read out through Multi-Channel Electronics units (MCEs) [41] using 2-stage, time-domain multiplexed, superconducting QUantum Interference Device (SQUID) amplifiers made by NIST [42]. Taking advantage of flux activated superconducting switches (‘row-selects’) [43], the need of an intermediate second-stage SQUID (previously referred to




Figure 2.8. Photographs of the 4'' and 6'' wafers used for BICEP3 and BICEP Array respectively from [39]. The larger wafers allow for more detector density and reduced fabrication time. Image from [38].

as SQ2) is eliminated, simplifying the multiplexing architecture substantially compared to previous generation experiments. SQ1s are arranged in rows and columns and the row-selects are used to sequentially bias each row of SQ1s at high frequency to their operating point. Each SQ1 column (sitting at FPU temperature) is coupled to a summation coil. Since SQ1s are biased per row, there is no ambiguity onto which detector is being read out at any given time, resulting in a large decrease in the total number of wires needed for readout.

The signal read by the SQ1s is minuscule and is amplified through the second stage of readout chain—the SQUID series array (SSA)—before being sent to the MCEs at room temperature. The SSAs offer impedance matching and low noise amplification, crucial for readout by the MCE units at room temperature. The MCEs control the row-selects, biasing, the flux-locked loop (FLL) feedback of the SQ1s, and digitize the feedback signal [44].

This architecture opens itself up to a systematic contamination seen in calibration measurements (see section 4.1.4) where neighboring detectors (read out by adjacent SQ1s) inductively couple to each other, resulting in part of the signal from one detector being picked up by its

upstream and downstream neighbors in the readout chain. This contamination is characterized in detail in section 4.1.

All of our cryostats are populated with a wide range of heaters and thermometers that offer precise monitoring and control of the temperature of various components of our instruments. In fact, we rely on NTDs to monitor the temperature of the focal plane using the same technology that was employed in BICEP1 detectors, but now used for temperature monitoring. Control is achieved through a custom-built ‘backpack’ that houses the necessary electronics and achieves precise temperature control through PID [45].

The instruments are controlled, and their data acquired by the Generic Control Program (GCP) [46], which orchestrates mount motion, receiver housekeeping, and timestream recording. Raw detector data are downsampled and filtered before being archived. GCP also keeps track of the labeling of the detectors as polarization ‘A’ or ‘B’ where $A \perp B$. In section 2.6 we will introduce a more robust way of labeling the detectors based on the angle of the detector’s polarization axis with respect to the focal plane orientation as opposed to the circuit-based labeling of the GCP.

2.4 Baffling

The South Pole is the driest and most stable site on Earth for CMB observations. However, at mm wavelengths everything is a source of emission, including the ground, the atmosphere, and the surrounding buildings. For this reason, careful attention is paid to the baffling of all our instruments in order to minimize the coupling of our detectors to unwanted sources of emission. We employ a two-stage baffling scheme to ensure that each ray originating from surrounding sources must diffract at least twice before reaching the vacuum window (the two-diffraction constraint).

The first stage is the use of roughly conical ground shields that surround our instruments and block direct line of sight to the ground and surrounding buildings. They are mounted at the roof of the DSL and MAPO buildings (for BICEP3 and BICEP Array respectively)

and are made of aluminum panels with a reflective surface to redirect stray reflections to the cold sky. They stay fixed at their nominal position during CMB observations, and are only tilted/lowered during calibration measurements during the Austral summer calibration seasons. The ground shield was redesigned for BICEP3 and therefore meets the two-diffraction constraint for all of BICEP3’s observing angles. However, BICEP Array inherited the ground shield from Keck Array, which has raised concerns on whether the two-diffraction constraint is met for all of BICEP Array’s observing angles. This is investigated in detail in chapter 6.

The second stage of baffling is the use of co-moving absorptive ‘forebaffles’ which consist of aluminum cylinders lined with Eccosorb HR-10 microwave-absorptive foam that are attached to each receiver’s window. Inevitably, the forebaffles add to the overall loading of the instrument, but significantly reduce the coupling of the detectors to sidelobes (see section 4.3 for measurements with forebaffles on/off, fully characterizing the level of sidelobe power attenuation). They are heated to a temperature slightly above ambient to prevent snow and ice accumulation. The forebaffle offers a natural boundary which we use to define the various realms of the angular response of our detectors (see section 3.1). Knowing the diameter of the forebaffle allows us to calculate the ‘forebaffle cutoff angle’ beyond which the forebaffle starts to attenuate the coupling of the detectors to the sky. For BICEP3 this cutoff is at $\sim 15^\circ$ away from the boresight pointing ray, while for BICEP Array this is at $\sim 18^\circ$ ¹. A visual representation of the two-stage baffling scheme for BICEP3 is shown in Figure 2.9.

2.5 Data Taking and Observation Strategy

There are many advantages to observing from the South Pole. Besides the existing infrastructure and support of the Amundsen-Scott South Pole Station through the United States Antarctic Program (USAP), the South Pole site offers unique atmospheric and environmental conditions that are ideal for CMB observations (low PWV, stable atmosphere, etc.). The BICEP/Keck series of experiments have been observing the same patch of the sky (varies

¹These will become important when we characterize forebaffle loading and attenuation using calibration data.




Figure 2.9. Schematic of the two-stage baffling scheme for BICEP3. The receiver is surrounded by the fixed ground shield that reflects off-axis rays to the cold sky. Atop BICEP3’s window is the co-moving, absorptive, forebaffle which attenuates power coming from sidelobes.

from experiment to experiment depending on the field of view) from the South Pole since BICEP1, which is accessible year-round. Figure 2.10 shows the location of the BICEP/Keck observing field in the southern sky, which is centered at $RA = 0h$, $Dec = -57.5^\circ$ and covers a sky fraction of $f_{sky} \sim 0.02$.

The fundamental unit of our observations is a ‘scanset’, which consists of a set of roughly 50 back-and-forth scans in azimuth with an Az throw of $\sim 60^\circ$, at a fixed El and Dk angle. The scan speed is fixed at $2.8^\circ/s$, which makes a scanset last for about 50 minutes. Internally, a scanset is also known as a ‘tag’, and a single scan in azimuth is referred to as a ‘halfscan’. A single tag is the basis for all our data reduction and mapmaking process (described in detail later in this chapter). Part of our observing strategy is to let the sky drift across the field of view because we do not adjust the azimuth zero-point to track the sky. This has been shown to remove terrestrial contamination that is fixed with respect to the ground and not the sky [38]. The scanset is decorated at the beginning and end by a set of elevation nods (“elnods”). These are small excursions ($\mathcal{O}(1^\circ)$) from the El center of the scanset that are used to modulate the atmospheric loading on the detectors and probe their responsivities. Various data quality cuts and the relative detector gain procedure rely on the elnods, making them an integral part of our data reduction process (see sections 2.7.1 and 2.7.1). Partial

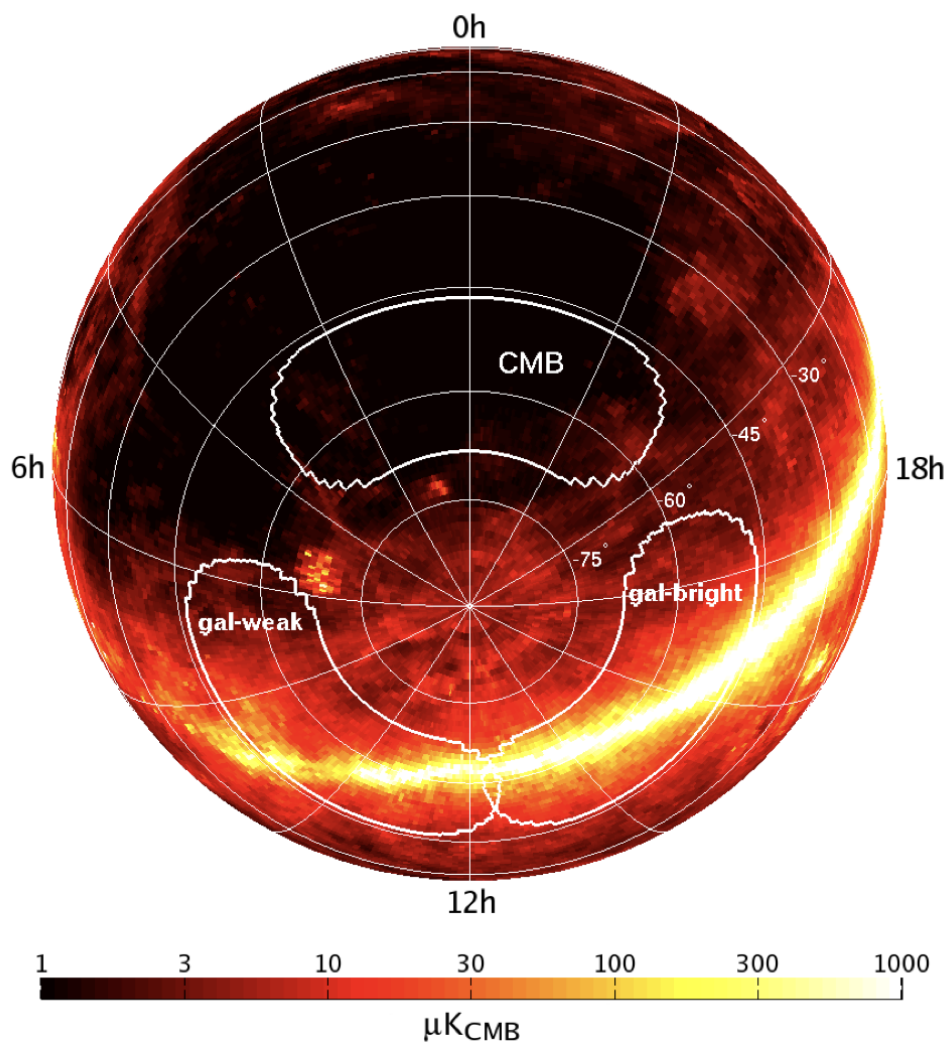


Figure 2.10. Location of the BICEP/Keck observing field in the southern sky overplotted the FDS Model 8 prediction of dust emission in [47]. Image from [48].

load curves (PLCs) are also taken at the beginning of each scanset, which are used to monitor any excursions from linearity in the detector response.

Each tag is repeated twice prior to slewing to a new Az and El position. The applied offset in elevation is $\sim 0.25^\circ$ in the Declination (Dec) direction. 10 repeated tags (5 El and Az offsets) belong to a group called a ‘phase’. The total number of phases that BICEP3 or BA can continuously observe depends on how long the sorption refrigerator can hold its base temperature prior to cycling. The typical hold time during CMB observations for BICEP3 is ~ 3 days, while for BA it is ~ 2 days. The total number of tags that can be observed consecutively belong to a single ‘CMB schedule’. After the schedule completion the fridge is cycled and Dk is rotated by a fixed angle. BICEP3 due to its limited Dk range observes at 4 evenly spaced Dk angles: 23° , 68° , 203° , and 248° , while BA observes at 8 evenly spaced Dk angles: 23° , 68° , 113° , 158° , 203° , 248° , 293° , and 338° . The phases along with the Az and El ranges of the individual tags that comprise them are shown in Figure 2.11.




Figure 2.11. Schematic of the BICEP Array observing strategy showing phases of CMB observations. Each tag consists of a set of back-and-forth scans in azimuth at a fixed elevation and deck angle. Each phase consists of 10 tags with small offsets in Az and El .

2.5.1 Starpointing

Overall telescope pointing of both BICEP3 and BICEP Array is monitored by small optical star cameras attached to each receiver. Starpointing observations are conducted approximately every two weeks, during which a predefined list of stars is observed at multiple deck

angles and fit to a multi-parameter model of the mount alignment. For BICEP Array, two additional parameters beyond the standard seven are included to account for sagging of individual receivers under gravity relative to the common mount boresight. For more details refer to [49].

2.6 Pointing Model

Exact knowledge of the instrument’s pointing and orientation at all times during observations is integral to accurately reconstructing the CMB sky. The pointing model allows us to map the receiver’s focal plane onto the sky. The goal is to utilize knowledge of the instrument’s mount parameters (e.g., encoder zero positions, tilt terms, etc.) to convert from raw encoder readings (Az_0, El_0, Dk_0) to topocentric horizontal coordinates (Az, El, PA) and then to celestial coordinates (RA, Dec, PA) . The input parameters for the pointing model are empirically measured on a roughly biweekly cadence throughout the calendar year, and are obtained through optical starpointing.

For CMB observations from the South Pole we are only concerned about the pointing direction of the instrument. A Cartesian basis is chosen in horizon coordinates (x, y, z) . The z -direction is chosen to be perpendicular to the groundshield floor and pointing towards zenith. The y -axis is then defined in a right-handed fashion as $\hat{y} = \hat{z} \times \hat{x}$, and points in the direction of $Az = -90^\circ$. We define a few quantities we keep track of: the instrument’s boresight pointing direction \hat{B}_3 and the focal plane orientation axis \hat{B}_1 :

$$\hat{B}_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad \hat{B}_3 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad (2.1)$$

With *a priori* knowledge of the telescope’s mount and starpointing parameters (i.e., encoder zero positions, flexure and tilt terms, etc.), we apply a series of translations and rotations to the boresight basis vectors to obtain the instantaneous position and orientation of the

boresight. These rotations are parameterized by the 3 Euler angles α, β, γ with the full solid rotation matrix $R(\alpha, \beta, \gamma)$ given by:

$$R(\alpha, \beta, \gamma) = \begin{bmatrix} \cos \alpha \cos \beta \cos \gamma - \sin \alpha \sin \gamma & -\cos \alpha \cos \beta \sin \gamma - \sin \alpha \cos \gamma & \cos \alpha \sin \beta \\ \sin \alpha \cos \beta \cos \gamma + \cos \alpha \sin \gamma & -\sin \alpha \cos \beta \sin \gamma + \cos \alpha \cos \gamma & \sin \alpha \sin \beta \\ -\sin \beta \cos \gamma & \sin \beta \sin \gamma & \cos \beta \end{bmatrix} \quad (2.2)$$

In the case of BICEP Array once we have the orientation axis \hat{B}_1 we apply 3 successive 90° rotations around the \hat{B}_3 axis to obtain the orientation axes of each receiver's focal plane. Figure 2.12 shows the layout of the BA receivers on the Dk drum along with their respective focal plane orientation vectors.

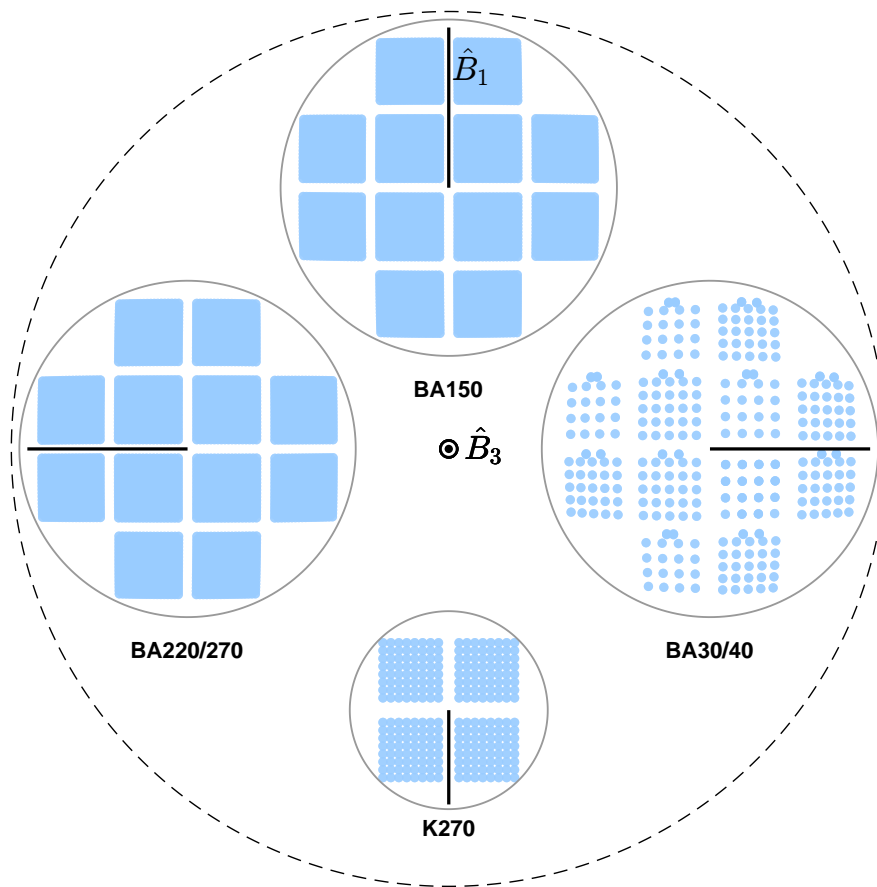


Figure 2.12. BICEP Array receiver orientations on the Dk drum in its 2026 configuration. Plotted in blue are the detector pointings for each receiver calculated from equation 2.4 and the orientation of the $\hat{B}_1 = 0^\circ$ axis that defines the orientation of each focal plane. The focal plane sizes are to scale. Receiver distances from the center of the drum have been scaled down by a factor of 2.5 for visualization purposes.

The topocentric horizontal coordinates of the boresight can be read off Figure 2.13 by projecting the boresight basis vectors onto (Az, El, PA) :

$$\begin{aligned}
 AZ &= \tan^{-1} \left(-\frac{\hat{B}_3 \cdot \hat{y}}{\hat{B}_3 \cdot \hat{x}} \right), \\
 EL &= \sin^{-1} \left(\hat{B}_3 \cdot \hat{z} \right), \\
 PA &= \tan^{-1} \left(-\frac{\hat{B}_2 \cdot \hat{z}}{\hat{B}_1 \cdot \hat{z}} \right)
 \end{aligned} \tag{2.3}$$

With the topocentric coordinates in hand, and knowledge of the instrument's latitude and longitude, we can convert to celestial coordinates (RA, Dec, PA) . Note that in topocentric coordinates, PA is defined relative to zenith as opposed to the celestial coordinates where PA is defined relative to the north celestial pole. Therefore, in the case of the South Pole site, they differ by 180° . With the position of the boresight fully characterized we want to get the coordinates and orientations of individual detectors on the focal plane. Suppose we want to calculate the horizontal coordinates of a detector located at position P . The orthonormal basis at the boresight can be rotated via parallel transport along the sphere to the detector's location. We use optical and empirical modeling of the focal plane geometry [50] to get the form of the detector's pointing vector \hat{P}_3 in our $(\hat{x}, \hat{y}, \hat{z})$ basis to define boresight basis r, θ, χ as:

- r : the projection angle of \hat{P}_3 onto \hat{B}_3 .
- θ : the bearing angle from the \hat{B}_1 axis toward $-\hat{B}_2$.
- χ : the angle between the detector's co-polar axis and r at P .

The bearing refers to the positive direction of θ . It is defined in the plane spanned by the orthonormal basis vectors \hat{B}_1 and \hat{B}_2 and is positive starting from \hat{B}_1 towards $-\hat{B}_2$ also shown in Figure 2.13. Picking the zero-axis of θ is somewhat arbitrary, keeping with the BICEP/Keck convention, it is referenced to the meridian at $Dk = 0$. Figure 2.13 shows a schematic defining the (r, θ, χ) coordinate system and all relevant quantities for the pointing

model. To rotate to each detector's location, we parallel transport the boresight basis vectors using r and θ . This solid rotation is applied in the following order:

- Rotate by θ around the \hat{B}_3 axis.
- Rotate by $-r$ around the newly formed \hat{B}_2 axis.
- Rotate by $-\theta$ around the newly formed \hat{B}_3 axis.

The resulting solid rotation ($R(\theta, -r, -\theta)$ in equation 2.2) can be expressed in matrix form as:

$$\begin{bmatrix} \hat{P}_1 \\ \hat{P}_2 \\ \hat{P}_3 \end{bmatrix} = \begin{bmatrix} \sin^2 \theta + \cos r \cos^2 \theta & (1 - \cos r) \sin \theta \cos \theta & -\sin r \cos \theta \\ (1 - \cos r) \sin \theta \cos \theta & \sin^2 \theta \cos r + \cos^2 \theta & \sin r \sin \theta \\ \sin r \cos \theta & -\sin r \sin \theta & \cos r \end{bmatrix} \begin{bmatrix} \hat{B}_1 \\ \hat{B}_2 \\ \hat{B}_3 \end{bmatrix} \quad (2.4)$$

Finally, we can use equation 2.3 to convert the detector's pointing vector into topocentric coordinates but instead of the boresight vectors $\{B_1, B_2, B_3\}$, we replace with the ones rotated to the detector location $\{P_1, P_2, P_3\}$ accordingly:

$$\begin{aligned} AZ_d &= \tan^{-1} \left(-\frac{\hat{P}_3 \cdot \hat{y}}{\hat{P}_3 \cdot \hat{x}} \right), \\ EL_d &= \sin^{-1} \left(\hat{P}_3 \cdot \hat{z} \right), \\ PA_d &= \tan^{-1} \left(-\frac{\hat{P}_2 \cdot \hat{z}}{\hat{P}_1 \cdot \hat{z}} \right) \end{aligned} \quad (2.5)$$

The prescription is then the same for converting to celestial coordinates.

Having empirical knowledge of a detector's \hat{P}_3 has allowed us to pinpoint its exact location on the sky. To keep track of its orientation, we require knowledge of its co-polar axis \hat{P}_{pol} . Focusing on the plane defined by \hat{P}_1 and \hat{P}_2 , the angle from the \hat{P}_1 axis to \hat{P}_{pol} with bearing towards $-\hat{P}_2$ gives the orientation of the detector's co-polar axis. Reading off from Figure 2.13 we have:




Figure 2.13. Schematic of the setup used to calculate the detector locations and orientations on the sky, illustrating the definitions of the r, θ, χ basis. Updated from Yoon 2007 [50]. The boresight basis vectors are denoted by $\hat{B}_1, \hat{B}_2, \hat{B}_3$, the detector pointing is \hat{P}_3 , and its co-polar axis is \hat{P}_{pol} . PA is the parallactic angle defined relative to zenith given by equation 2.5. We also define the orientation angle of the detector ϕ_d relative to the focal plane's orientation axis \hat{B}_1 , used to designate detectors as 'H' or 'V' (see text for details).

$$\phi_d = \tan^{-1} \left(-\frac{\hat{P}_{pol} \cdot \hat{P}_2}{\hat{P}_{pol} \cdot \hat{P}_1} \right) \quad (2.6)$$

Our detector pairs, labeled as ‘A’ and ‘B’, are defined relative to the circuitry of the readout system. However, depending on the orientation of the module that houses the detectors, the ‘A’ detector of one module may not necessarily have the same orientation as the ‘A’ detector of another module. Here we adopt a convention first introduced in [51] where detectors with ϕ_d nearly 0° are denoted as ‘H’ and detectors with ϕ_d nearly 90° are denoted as ‘V’. The designations stand for *horizontal* and *vertical* orientations with respect to the focal plane’s orientation axis \hat{B}_1 . This definition allows us to decouple the readout-based A/B nomenclature from the actual physical orientation of the detectors on the focal plane. We can now write the relationship between the detector’s co-polar axis and r as:

$$\chi = \phi_d - \theta \quad (2.7)$$

This definition completes the set of quantities we need to map a detector’s location and orientation on the sky in terms of the focal plane coordinates (r, θ, χ) . The pointing model can now be used to project the per-detector polarization angles onto the sky α_d . A location on the sky in celestial coordinates with reference to the north celestial pole is observed with polarization angle:

$$\alpha_d = PA_d + 180^\circ + \phi_d \quad (2.8)$$

used in mapmaking of CMB polarization maps (section 2.7.2) discussed in the next section.

2.7 From Timestreams to Power Spectra

BICEP3 and BICEP Array data are saved to disk after being filtered by the MCEs and GCP as described in section 2.2. We use binary files, called *arcfiles*, to store the raw timestreams

along with relevant information that give us a complete picture of the state of the instrument at each sampling instance. Arcfiles are transferred to North America via the South Pole Tracking and Data Relay Satellite System² (SPTR), staged through the United States Antarctic Program (USAP) Data Center in White Sands, CO. Finally, raw data are transferred to the Harvard University Faculty of Arts and Sciences (FAS) Research Computing (RC³) high-performance computing (HPC) cluster, where the data are stored and processed. In this section we give an overview of our CMB analysis pipeline, including mapmaking and power spectrum estimation, along with systematics mitigation techniques. The final goal is to transform the raw timestreams into CMB power spectra that are fit to a multi-component, foreground-separated model to extract cosmological and foreground parameters. This parameterized model is then used (along with specialized simulations) in our maximum likelihood (ML) search to establish the statistical significance of a measurement of the tensor-to-scalar ratio r .

2.7.1 Data reduction

As of the writing of this thesis, the BICEP/Keck CMB analysis pipeline is carried through in the MATLAB⁴ software. Since our instruments are set up for 24/7 operation, preprocessing the raw timestreams into data quality monitoring products is critical. This preliminary data reduction process is fully automated and conducted by the graduate student members of the BK collaboration (the ‘*Reduc Czars*’). Time ordered data (TOD) are first deconvolved to undo the effects of the MCE and GCP filtering operations. At this stage, individual detector pointings are calculated (as described in the previous section) and known artifacts (e.g., glitches, dropped samples, etc.) are corrected or flagged. Starpointing model parameters are also calculated at this stage if a starpointing schedule was run during the reduction period. Elnod calibrations and corresponding data quality metrics are calculated prior to downsampling the TOD and saving to disk. For a detailed description of the TOD reduction

²<https://www.usap.gov/technology/1971/>

³<https://www.rc.fas.harvard.edu>

⁴<https://www.mathworks.com>

process see [38]. Figure 2.14 shows an example of raw BICEP3 timestreams used in data quality monitoring.

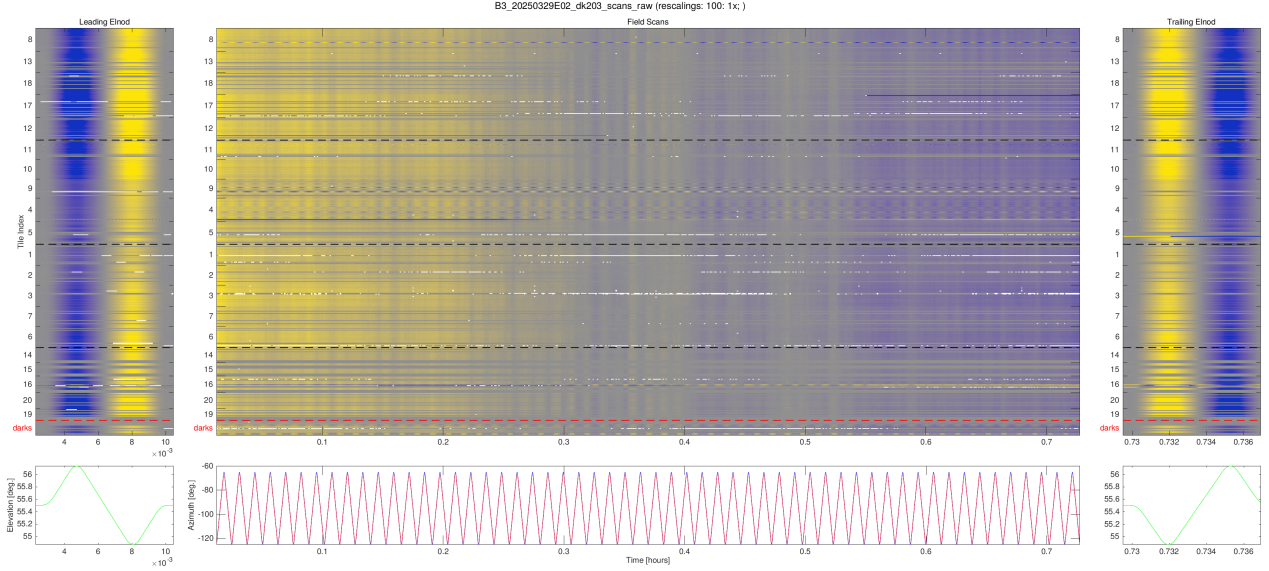


Figure 2.14. BICEP3 raw TOD waterfall plot for a single 50 minute scanset. From left to right, the top row of sub-panels shows the individual detector response to the leading elnod, then that of the azimuth scans, and finally of the trailing elnod. The bottom row shows the elevation and azimuth motion of the instrument during the scanset. The horizontal axis indicates the temporal duration of the entire measurement (internally referred to as ‘tag’). The vertical axis indicates individual detectors ordered by tile location on the focal plane. Horizontal dashed lines indicate MCE readout boundaries so each block of detectors corresponds to a single MCE. Some scan synchronous signal (likely associated with the atmosphere) is visible in the azimuth scans, along with an overall temperature gradient.

Relative gain calibration and filtering

Among a variety of data quality metrics calculated during the reduction process, we apply a series of calibration and filtering processes to the TODs to remove known contamination. We start by accounting for the relative gain between detectors within a pair by fitting their respective elnod responses to a $\csc(El)$ model of airmass. Without relative gain calibration pair-sum/difference quantities would be meaningless. The atmosphere is optically thin in the BK observation bands (as shown in Figure 2.2), hence the atmospheric emission recorded by a detector $d(t)$ is proportional to the model of airmass:

$$\Delta d(t) = g_{elnod} \cdot \Delta \csc(El) + n(t) \quad (2.9)$$

where g_{elnod} is the elnod gain factor to be fit, and $n(t)$ is the noise contribution (negligible in this case). Since the atmospheric temperature can change with weather or over time, we use the band-averaged elnod gain to calibrate the TODs between scansets. The key assumption here is that the band-averaged sensitivity of our instruments is constant [52]. The relative gain procedure is therefore:

$$d_{cal}(t) = \frac{1}{2}d_{uncal}(t) \times \left(\frac{\text{med}_{band}(g_{elnod})}{g_{elnod}} \right) \quad (2.10)$$

where the *cal* and *uncal* subscripts refer to calibrated and uncalibrated quantities respectively, and $\text{med}_{band}(g_{elnod})$ is the median elnod gain across all detectors in the frequency band.

A waterfall plot of BICEP3's elnod gains normalized across an entire season is shown in Figure 2.15. Each row corresponds to a detector and each column to a scanset. The color indicates the relative gain of the detector as measured by the elnod. The x-axis ticks indicate changes in Dk and correspond to tags that pass round 2 cuts. Excluding the prominent vertical stripes that denote periods of bad weather, and dead detector channels denoted as horizontal stripes, one notices the repeated Dk -dependent pattern. This influences our detector relative calibrations and needs to be understood. This was first noticed in BICEP3 and later confirmed to be present in both BICEP2 and Keck Array. We will come back to this in detail in section 5.2

Practically a few problems arise with this method. The median elnod value can shift depending on detector bias changes during the CMB observing season. It can also shift when some blocks of detectors are not working. During periods of bad weather some detectors go into saturation (reduced gain) before others resulting in a biased median elnod gain value. Finally, any Dk dependent systematics in the relative gains would further complicate the calibration (described in detail in section 5.2 for BICEP3).

To address the above issues, we implement per-pair relative gain calibration during the initial data reduction stage, but defer the pair-to-pair relative gain correction until the pairmap coaddition step (see section 2.7.2).




Figure 2.15. Observed elnod gains for BICEP3 during the 2018 observing season. The y-axis corresponds to individual detectors. The x-axis spans all measurements during the season and each tick corresponds to a single Dk observation. The elnod gain of each detector, extracted from a $\text{csc}(El)$ model of airmass, is normalized over the entire season and then plotted. The oscillatory behavior as a function of scanset number associated with boresight rotation is discussed in detail in section 5.2

The next filtering step is azimuth-fixed signal subtraction. As the name suggests, for each scan, we remove the mean signal as a function of azimuth. For each scan i , this process takes the form:

$$d_i^*(Az) = d_i(Az) - \text{mean}(d_i(Az)) \quad (2.11)$$

This process removes known scan-synchronous contamination, such as satellite communication signals from the nearby station, or any other azimuth-fixed signals. However, any time-varying azimuth-synchronous signal, such as coupling to the ground emission, may not be fully removed by this process, something we explore in detail in Chapter 6. An example of the azimuth-fixed signal removed from BICEP3’s data over the course of a single tag, is shown in Figure 2.16 for right- and left-going scans separately.

Finally, we apply a 3rd order Legendre polynomial filter to each scan to remove large scale variations primarily associated with atmospheric emission. TOD weights are calculated prior to Legendre polynomial filtering as the reciprocal of the power spectral density (PSD) from 0.1 – 0.3 Hz ⁵. Post filtering, and relative gain calibration, the pair-sum and pair-difference

⁵Frequencies corresponding to angular scales most relevant to the science goal around $\ell = 80$.

2. BICEP/Keck Experiments

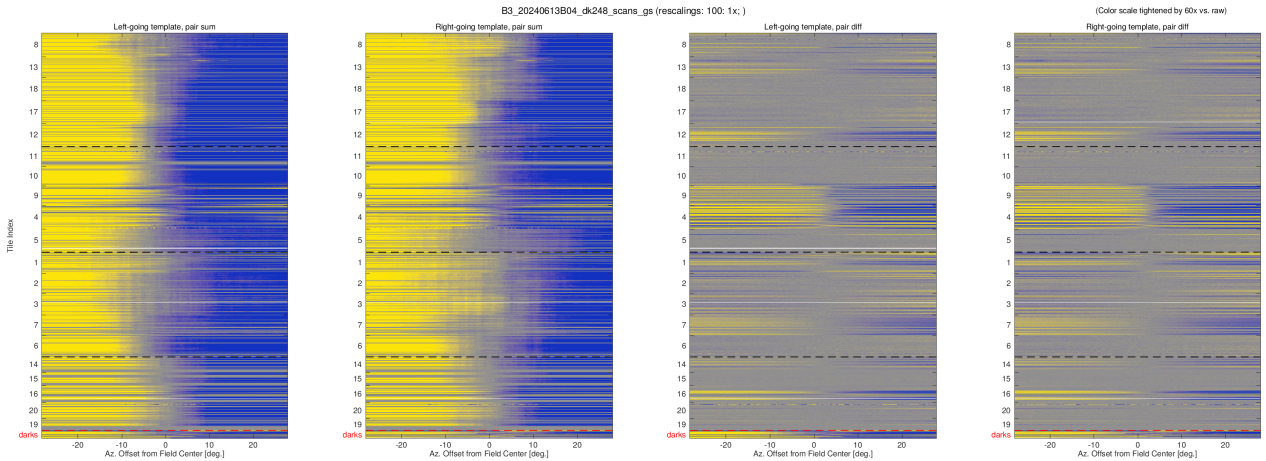


Figure 2.16. Pair-sum and pair-difference azimuth-synchronous signal templates removed from BICEP3 data shown for right-going and left-going scans separately. The vertical axes correspond to detector number and the colorscale to the azimuth-synchronous signal amplitude for each detector (ordered by tile location on the focal plane as in Figure 2.14). This process robustly removes azimuth-synchronous contamination [38]. Any time varying signal (e.g., ground coupling) is not fully removed by this process. Such coupling is explored in Chapter 6.

TODs are calculated, to be used in the mapmaking process along with the pair’s polarization

$$\text{angle } \alpha_{pair} = \frac{\alpha_V + \alpha_H - \pi/2}{2}.$$

Data quality cuts

An integral part of the data reduction process is data selection based on diagnostic parameters that define quality data. These parameters describe (among other things) the measurement and instrument stability, the atmospheric conditions, and the detector performance during a scanset. A subset of these parameters is then compared against empirically defined thresholds to produce logical masks that exclude low-quality data from downstream analysis. An example would be the exclusion of known, non-functioning detectors. We keep track of *good light* detectors (‘gls’) that are known to be functioning well. In the case of a good A/B pair, the list is further reduced to *real good lights* (‘rgls’). The level of curation of these lists has fluctuated substantially over the years with a steady decline proportional to detector count. The data cutting procedure involves two categories of cuts:

- **Round 1:** per-halfscan cuts related to pathologies at the raw timestream level (e.g., glitches, flux jumps, dropped samples, etc.) that are fairly uncommon.

- **Round 2:** per-scanset cuts related to noise stationarity, elnod fit quality, weather conditions and detector performance. This stage is where the majority of data cuts occur.

A comprehensive list of data quality cuts used in analysis can be found in Chapter 4 of [38]. In their current implementation, Round 2 cuts are compared against threshold values and can be turned on or off depending on those thresholds. We note that scansets that pass cuts are proportionally downweighted in analysis depending on the overall achieved cut fraction. That way, scansets that have suffered substantial data cuts are downweighted accordingly. We apply the cuts prior to mapmaking which is what we turn to next.

2.7.2 Mapmaking

With the initial reduction of the TODs completed, we have all the ingredients to bin the data into CMB maps. From 2022 and onwards, binning CMB maps is done in the Hierarchical Equal Area isoLatitude Pixelization⁶ (HEALPix) scheme [53] with $N_{side} = 256$. This choice over the previously used equidistant cylindrical projection, considerably reduces the power in a given ℓ that is sourced from other multipole bins (discussed in section 2.8), allowing for a more accurate comparison between data and theory. After correcting for the telescope mount pointing and calculating the per-detector pointings as described in section 2.6, the Zemax⁷ modeling software is used to correct the per-detector pointing centroids that (up to this point) have only been based on location on the focal plane. Using ray-tracing we propagate the detector pointings through the optics chain, defining our ‘ideal’ detector pointings on sky. The ideal pointings are then further adjusted by cross-correlating a fine resolution CMB temperature map of our instrument against the nearest frequency Planck temperature map [54]. These define the ‘CMB-derived’ pointings in celestial coordinates (RA, Dec) that are used in mapmaking.

To sum up, we have the pair-sum and pair-difference, filtered TODs denoted by z_s, z_d . These

⁶<https://healpix.jpl.nasa.gov/>

⁷<https://www.ansys.com/products/optics/ansys-zemax-opticstudio>

are recorded in (RA, Dec) at an angle α (dropping the subscript for ease) on the sky. We also have the corresponding weights w and variances v_s, v_d . The mapmaking process is carried out in the so called ‘one-shot’ fashion, where the above quantities are accumulated over all data as opposed to combining individual maps from detector pairs. This way, additional data and weight/cuts are incorporated dynamically. For a set of observations i that fall into a given pixel in (RA, Dec) we accumulate the following quantities:

$$w = \sum_i w_i, \quad (2.12)$$

$$T = \frac{1}{w} \sum_i w_i z_{s,i}, \quad (2.13)$$

$$v_T = \frac{1}{w^2} \sum_i w_i^2 v_{s,i}, \quad (2.14)$$

$$wcz = \frac{1}{w} \sum_i w_i \cos(2a_i), z_{d,i}, \quad (2.15)$$

$$wsz = \frac{1}{w} \sum_i w_i \sin(2a_i), z_{d,i}, \quad (2.16)$$

$$wcc = \frac{1}{w} \sum_i w_i \cos^2(2a_i), \quad (2.17)$$

$$wss = \frac{1}{w} \sum_i w_i \sin^2(2a_i), \quad (2.18)$$

$$wcs = \frac{1}{w} \sum_i w_i \cos(2a_i) \sin(2a_i), \quad (2.19)$$

$$wwccv = \frac{1}{w^2} \sum_i w_i^2 \cos^2(2a_i), v_{d,i}, \quad (2.20)$$

$$wwssv = \frac{1}{w^2} \sum_i w_i^2 \sin^2(2a_i), v_{d,i}, \quad (2.21)$$

$$wwcsv = \frac{1}{w^2} \sum_i w_i^2 \cos(2a_i) \sin(2a_i), v_{d,i}. \quad (2.22)$$

These quantities are sufficient to calculate the Q, U maps and their corresponding variance maps ⁸. The first three quantities in the above equations correspond to the pair sum TODs,

⁸This mapmaking process will be revisited in section 4.2 in the context of calibration measurements with a polarized source.

and they already give us the observed temperature map T_{obs} . The remaining quantities involve projecting the pair difference TODs along the Stokes Q, U basis and keep track of the corresponding weights and variances. To make life easier, we define a vector corresponding to each measurement i of sky pixel as:

$$\mathbf{s}_i \equiv \begin{bmatrix} \cos(2a_i) \\ \sin(2a_i) \end{bmatrix}^T \quad (2.23)$$

These vectors live in \mathbb{R}^2 . We calculate the observed polarization maps Q_{obs}, U_{obs} as:

$$\begin{bmatrix} Q_{obs} \\ U_{obs} \end{bmatrix} = \mathbf{m} \begin{bmatrix} wcz \\ wsz \end{bmatrix} = \frac{1}{w} \sum_i w_i \mathbf{m} \mathbf{s}_i v_i, \quad (2.24)$$

where the 2×2 matrix \mathbf{m} (sometimes referred to as the *QU decorrelation matrix*) is:

$$\mathbf{m} = \begin{bmatrix} wcc & wcs \\ wcs & wss \end{bmatrix}^{-1} = \begin{bmatrix} e & f \\ f & g \end{bmatrix}. \quad (2.25)$$

which isn't invertible in the case of a single observation at angle α . This emphasizes our scan strategy of observing at multiple equally spaced boresight angles. The result is:

$$Q_{obs} = e wcz + f wsz, \quad U_{obs} = f wcz + g wsz. \quad (2.26)$$

If all observations of pixel i were made at the same angle α , then \mathbf{m} would be singular and analytically non-invertible, showcasing the importance of our scanning strategy that involves multiple boresight rotation angles.

Our small aperture receivers have large field of view. Given the scanning strategy we inevitably observe the center regions of the sky patch more compared to the edges. For this reason we construct variance maps to keep track of the per-pixel coverage of the observed sky. These variance maps are the elements of the 2×2 matrix:

$$\begin{bmatrix} v_Q & v_{QU} \\ v_{QU} & v_U \end{bmatrix} = \frac{1}{w^2} \sum_i w_i^2 \mathbf{m} \mathbf{s}_i \mathbf{s}_i^T \mathbf{m}^T v_i. \quad (2.27)$$

Binning is first applied to individual detector pairs, producing ‘pairmaps’ over a double scanset, at which time the second round of cuts is applied. The pairmaps are then array-averaged before further coadding over all scansets in a given season. The result is a full season coadded map, and an example for the BA150 receiver is shown in Figure 2.17.

2.7.3 Deprojection

In the time-reverse sense, one can think of the illumination pattern of a detector as its two-dimensional angular response projected onto the sky instead of receiving light. The angular response pattern is referred to as the *beam* of the detector (also known as the point spread function (PSF) in optics). In other words, the beam describes how the power received by a detector is distributed on the sky as a function of angle away from its central pointing ray. The mapmaking formalism of the previous section assumes ideal detectors (perfectly orthogonal and co-located pairs), with perfectly circular beams. If non-idealities within a detector pair are ignored, beam mismatch will leak part of the bright temperature anisotropies into the much fainter polarization anisotropies ($T \rightarrow P$ leakage). The leakage signal in terms of the pair-difference beam at direction \hat{n} , is expressed as:

$$\begin{aligned} d_{T \rightarrow P} &= T(\hat{n}) * [B_A(\hat{n}) - B_B(\hat{n})] \\ &\equiv T(\hat{n}) * B_\delta(\hat{n}) \end{aligned} \quad (2.28)$$

with B_δ denoting the ‘difference beam’.

Deprojection is a technique used to remove the leading order contamination due to beam mismatch. We model the difference beam as the difference between two elliptical Gaussian

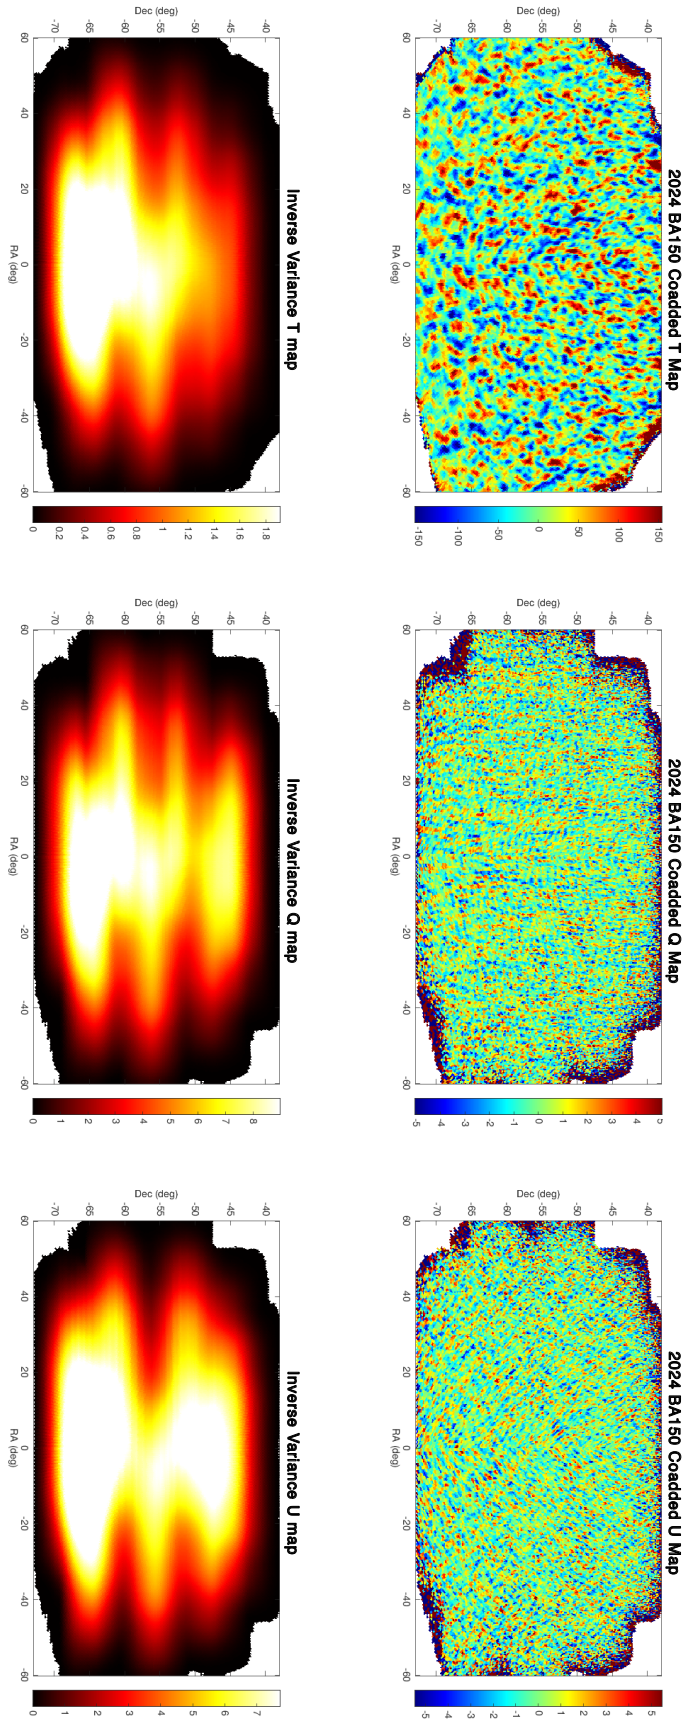


Figure 2.17. BA150 Single year coadded map and inverse variance map.

beams (see section 3.5.2 and BK-III systematics paper [55]). Four non-idealities are considered resulting in six total parameters used to describe the difference beam:

- **Relative gain:** $B_{\delta g}(\hat{\mathbf{n}})$ arising from the difference in peak amplitude of the two beams.
- **Differential pointing:** $B_{\delta x}(\hat{\mathbf{n}}), B_{\delta y}(\hat{\mathbf{n}})$ modeled as a centroid offset in focal plane coordinates.
- **Differential beamwidth:** $B_{\delta\sigma}(\hat{\mathbf{n}})$ resulting from beamwidth mismatch between the two beams.
- **Differential ellipticity:** $B_{\delta p}(\hat{\mathbf{n}}), B_{\delta c}(\hat{\mathbf{n}})$ modeled as intra-pair differences purely in the ‘+’ and ‘×’ orientations, defined with respect to the focal plane’s $\theta = 0^\circ$ axis.

Figure 2.18 shows the six elliptical Gaussian modes used to model the difference beam. Differences in elliptical Gaussians would couple to different combinations of the temperature field and its derivatives (by the convolution theorem). Therefore, analytic knowledge of the beam mismatch modes would allow us to construct templates of the leakage signal from the temperature map and distinct linear combinations of its derivatives. We filter out part of the leakage signal by regressing these templates against our data and subtracting them. The biggest advantage of this approach is that it does not require precise knowledge of the actual mismatch amplitude.

The analytic form of the leakage templates can be derived by perturbing parameter δk with $k \in \{g, x, y, \sigma, p, c\}$ of the elliptical Gaussian beam model and Taylor expanding keeping terms up to second order where appropriate. They are derived in [56] and are found to be:

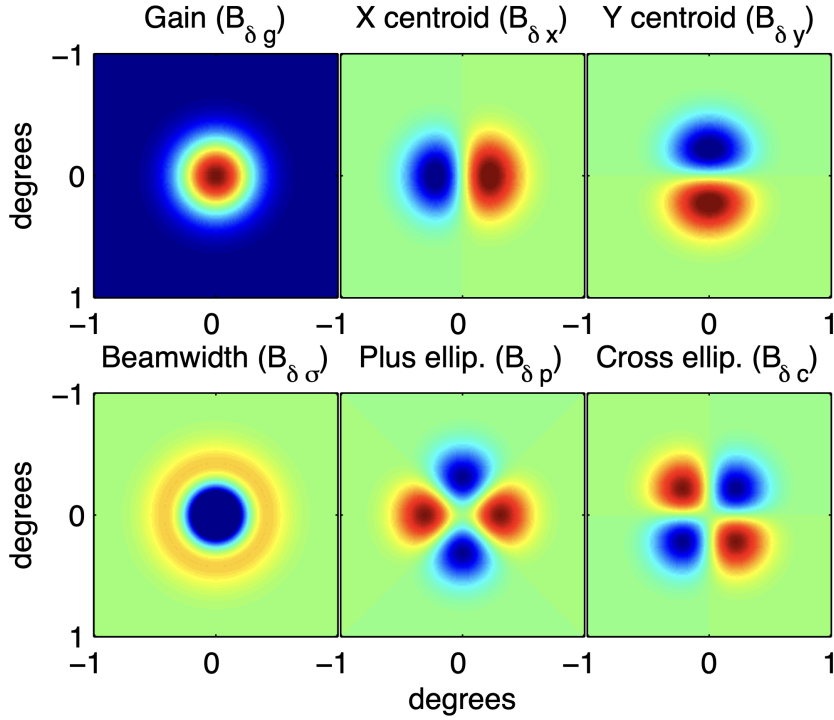


Figure 2.18. Leading six elliptical Gaussian modes used to model the difference beam $B_\delta(\hat{\mathbf{n}})$. From left to right, top to bottom: relative gain, differential pointing in x and y , differential beamwidth, and differential ellipticity in ‘+’ and ‘x’ orientations. Figure from the BK-III systematics paper [55].

$$B_{\delta g} \simeq \delta g B(\hat{\mathbf{n}}) \quad (2.29)$$

$$B_{\delta x} \simeq \delta x \nabla_x B(\hat{\mathbf{n}}) \quad (2.30)$$

$$B_{\delta y} \simeq \delta y \nabla_y B(\hat{\mathbf{n}}) \quad (2.31)$$

$$B_{\delta \sigma} \simeq \sigma \delta \sigma (\nabla_x^2 + \nabla_y^2) B(\hat{\mathbf{n}}) \quad (2.32)$$

$$B_{\delta p} \simeq \frac{\sigma^2}{2} \delta p (\nabla_x^2 - \nabla_y^2) B(\hat{\mathbf{n}}) \quad (2.33)$$

$$B_{\delta c} \simeq \sigma^2 \delta c (\nabla_x \nabla_y) B(\hat{\mathbf{n}}) \quad (2.34)$$

where $B(\hat{\mathbf{n}})$ is a perfectly circular elliptical Gaussian centered at the origin with $p = c = \delta x = \delta y = 0$, and $g_A = g_B$. The process starts with Planck maps of the CMB temperature field $T(\hat{\mathbf{n}}, t)$ (the ‘template maps’) and its derivatives⁹ calculated at each time sample t along the

⁹If a frequency band is not a direct match to the Planck bands, we take the nearest frequency.

pointing trajectory of each detector pair during a scanset. We interpolate off the template maps along the pointing trajectory to produce derivative timestreams $d_{i,j}$ in rectangular focal plane coordinates $j = x, y$. The template maps have their instrument beam deconvolved and then convolved with the nominal beam width of a given BK frequency band¹⁰. The resulting derivative timestreams are:

$$d_{i,j}(\hat{\mathbf{n}}, t) = \left. \frac{\partial^i}{\partial j^i} T(\hat{\mathbf{n}}, t) \right|_{\hat{\mathbf{n}}=\hat{\mathbf{n}}(t)} \quad (2.35)$$

The derivative timestreams are combined using equations 2.34 to produce ‘leakage templates’ $d_{\delta k}(t)$ for each beam mismatch mode. The leakage templates are computed for each scanset and coadded over each phase. The net leakage is therefore (dropping the explicit dependence on $\hat{\mathbf{n}}$ for clarity):

$$d_{\delta}(t) = \sum_k a_k d_{\delta k}(t) \quad \text{with} \quad k \in \{g, x, y, \sigma, p, c\} \quad (2.36)$$

During the pairmap coaddition stage, the leakage templates are regressed against the pair’s bolometer data to obtain the ‘deprojection coefficients’ a_k . Subsequently, the net leakage template is added up and removed from the data. We have the option to scale the templates $d_{\delta k}(t)$ by their corresponding deprojection coefficient a_k , or to fix a_k to a known value (e.g., from beam map measurements see section 3.5.2) before subtracting from the data. In practice, we deproject differential pointing and differential gain, subtract differential ellipticity, and leave differential beamwidth uncorrected as it is subdominant [55]. The nomenclature used to describe the deprojection choices consists of a 0 for no deprojection/subtraction, a 1 for deprojection using the deprojection coefficients, and a 2 for subtraction using fixed coefficients from independent beam measurements. For example the nominal deprojection scheme used in BK18 is denoted as *dp1102*. The inclusion of additional deprojection templates motivated by the far-field characterization of the beams are explored in section 4.1.4.

¹⁰Pre-convolving involves the measured beam window functions discussed in section 4.1.1

2.7.4 Absolute calibration

The bolometer data are recorded in arbitrary analog-to-digital converter units ADU (also sometimes referred to as raw feedback units FBU) and have been equalized within each pair with the relative gain calibration procedure. Prior to the pairmap making stage, the units of the bolometer data are in atmospheric median emission, before dividing by the individual elnod gains that brings the units back to ADU equalized per pair. An absolute calibration is performed to relate these units to μK_{CMB} . This is done by cross-correlating the season coadded maps (post pair-to-pair relative gain calibration) against external temperature maps from Planck similar to the calculation of the CMB-derived detector pointings described in the previous section 2.7.2. Specifically, the maps used for cross-correlation are coadded per-tile and per- Dk angle, to capture tile-to-tile gain, bandpass, and optical efficiency variations, as well as gain systematics related to Dk rotation (see section 5.2). This means that a full season must first be coadded for the absolute calibration factors to be calculated. The external maps have been sampled, filtered, and weighted to match the BICEP/Keck analysis pipeline (a process referred to as ‘reobservation’). For a given multipole bin ℓ the absolute calibration factor α_ℓ is:

$$\alpha_\ell = \frac{\hat{P}_{ref} \hat{P}_{cal}}{\hat{P}_{ref} \hat{K}} \quad (2.37)$$

where \hat{P}_{ref} is the reference Planck map at the nearest frequency band, \hat{P}_{cal} is the ‘calibrator’ Planck map of the second-nearest frequency band ¹¹, and \hat{K} is the power of BICEP/Keck full season coadded map at the given multipole bin. All the Planck maps used in this process have been reobserved through the BICEP/Keck pipeline.

2.7.5 Power spectra

We transform the measured Q/U maps to the coordinate-independent E/B basis in harmonic space using equation 1.41. A few complications arise trying to implement this in

¹¹Used to avoid self-correlating the noise in the Planck maps in the numerator

practice. First, since transforming to E/B basis is done in harmonic space, we must ensure the maps go to zero at the boundaries to avoid ringing. This is accomplished by applying an apodization¹² mask, calculated by the geometric mean of the inverse variance maps across all frequency bands. Since the pixels near the map boundaries are naturally noisier, apodization significantly downweights them which is a bonus. Second, the partial sky coverage results in ambiguous modes that mix E and B. Reconstructing E and B modes hinges on the orthogonality of the chosen basis. Ground-based experiments can only observe a fraction of the sky and are therefore unable to fully and orthogonally separate E and B modes¹³. Third, filtering operations like the Legendre polynomial filtering, deprojection, apodization, and azimuth-synchronous signal subtraction inevitably remove modes which introduces additional E/B mixing. At the current levels of sensitivity, these effects are very relevant and must be accounted for.

Matrix-based Map Purification

To address all the above complications, we take advantage of the linearity of our filtering operations. In principle, they can all be expressed in terms of matrices that operate on the data. When these matrices are multiplied together, they form a single *observation matrix* that encapsulates all filtering operations. Therefore, an external map is *observed* by applying the observation matrix to it. The starting point is to calculate the pixel-pixel covariance matrix of the observed map, which would have all the filtering operations baked in. A matrix-based purification approach that utilizes the eigendecomposition of the observed pixel covariance is used to define ‘pure-E’ and ‘pure-B’ subspaces. It is achieved by solving a regularized eigendecomposition problem (equation 70 in [58]). This approach was originally proposed by Bunn et al. [59] and later extended to BICEP2 and Keck Array analysis [58]. With the eigenvectors in hand (which might not span the space), matrix operators referred to as ‘purification matrices’ are then constructed to project the observed map onto the pure-E and pure-B subspaces. The process is described in detail in [60], which shows how matrix

¹²A Greek word that translates to ‘remove leg’

¹³An example of an estimator aiming to correct E-to-B leakage due to partial sky can be found here [57]

purification reduces E-to-B leakage power due to filtering to subdominant levels at the current sensitivity. The computational cost of this process is high for maps with $N_{side} = 256$, so we are exploring approaching the problem using GPU acceleration. However, it is still work in progress and (as of the writing of this thesis) is still unclear how helpful it will be.

Transforming to Spherical Harmonic Space

At this stage, the actual binning procedure is carried out, along with application of the apodization mask. Then the pure-B projection operators are applied to the Q/U maps. For the reanalysis (post 2022 era) the power spectrum estimation is done by using the HEALPix `map2alm` routine implemented via the `healmex` package for MATLAB compatibility.¹⁴ The auto and cross spectra C_ℓ^{XY} are calculated from the resulting spherical harmonic coefficients $a_{\ell,m}^{XY}$ by using equation 1.42. C_ℓ 's are then averaged in ℓ when binned into 5 bins from $\ell = 37.5$ to $\ell = 177.5$ equally spaced by $\Delta\ell = 35$ which is the BK ‘science band’. Notably, prior to the reanalysis the transformation to ℓ space was computed using the flat-sky approximation implemented via 2D discrete Fourier transforms (DFTs) returning the maps in terms of angular scales ℓ_x, ℓ_y with the radial bin defined as $\ell = \sqrt{\ell_x^2 + \ell_y^2}$ and the angular coordinate $\varphi = \arctan(\ell_y/\ell_x)$. In this basis, the filtering operations discussed in the previous sections that remove atmospheric and large scale effects would manifest as a substantial power suppression at low ℓ . Figure 2.19 shows the EE power in the 2D Fourier domain from a single season coadded BICEP3 map, showcasing the effect of beam suppression at high ℓ and filtering at low ℓ .

It is worth noting that most of the modes in the lowest bin we observe (centered at $\ell = 17.5$), are removed due to the timestream filtering operations (like Legendre polynomial filtering) and therefore contains little signal. For this reason, it is excluded from analysis.

¹⁴<https://healpy.readthedocs.io/en/latest/generated/healpy.sphtfunc.map2alm.html>

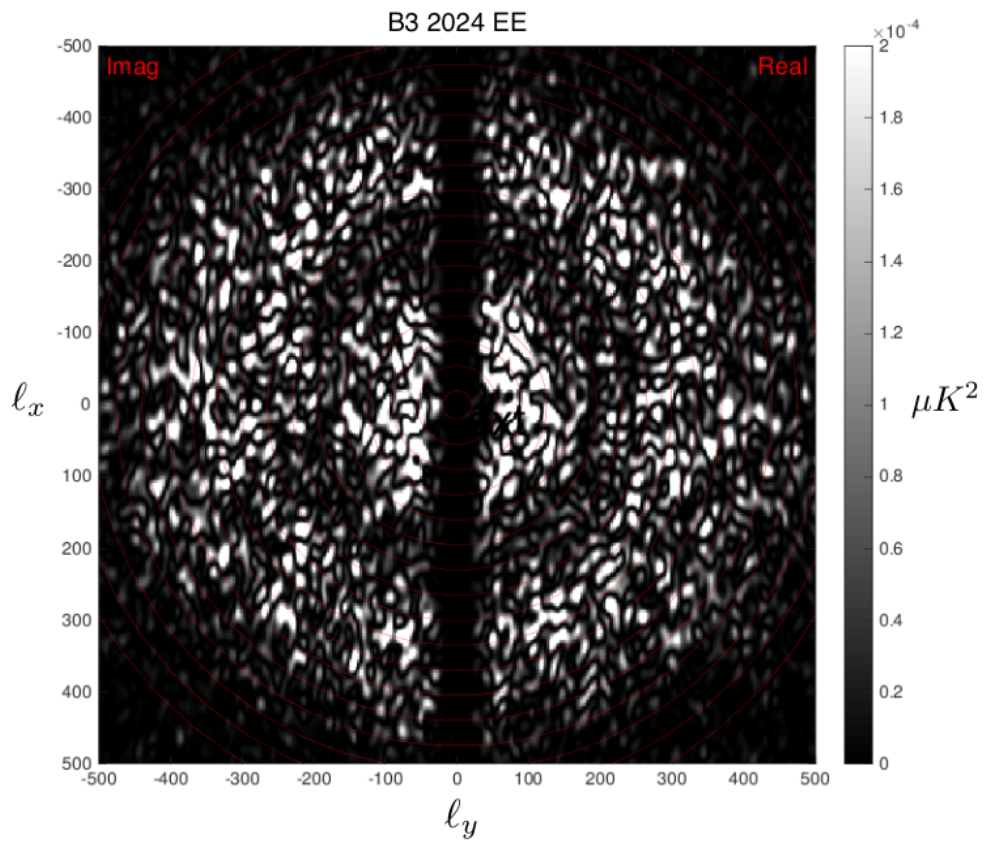


Figure 2.19. BICEP3 2024 EE power spectrum in the Fourier domain ($\ell = \sqrt{\ell_x^2 + \ell_y^2}$) showing the characteristic ‘trench’ at low ℓ due to our filtering operations, shown purely for illustrative purposes. The bright rings correspond to the low- ℓ modes of the EE power spectrum, while the power suppression at high ℓ is the result of the beam window function suppressing power at small angular scales that the instrument is not sensitive to.

Simulations framework

Our CMB analysis follows the pseudo- C_ℓ ¹⁵ approach where we estimate the power spectra from the maps directly, and correct for the effects of filtering, beam smoothing, and sky coverage via simulations. We generate a set of signal and noise simulations using a timestream-based reobservation approach that does not use the observation matrix described in the previous section. Realizations of the noise are generated using the ‘sign-flip’ procedure during the per-phase coaddition stage. The data chunks are randomly multiplied by ± 1 and coadded while paying attention that the data associated with $+1$ and -1 have equal weights.

BK signal simulations include lensed and unlensed Λ CDM, Gaussian dust, and $r = 0.1$ B-mode only simulations. 499 realizations of the Λ CDM sky are generated using CAMB software [13] constrained by the power spectrum of the Planck 2013 best fit cosmological parameters [3]. The lensed variants are produced using the Lenspix software [61]. The dust simulations are generated by assuming Gaussian statistics with a thermal power law spectrum:

$$C_{\ell,\text{dust}}^{BB} = A_d \left(\frac{\ell}{80} \right)^{-0.4}, \quad (2.38)$$

$$C_{\ell,\text{dust}}^{EE} = 2 \times C_{\ell,\text{dust}}^{BB}.$$

with the value of A_d taken from Planck at 353 GHz to be $A_d = 3.75 \times \mu K_{CMB}^2$ ¹⁶ The $r = 0.1$ simulations are also generated using CAMB but explicitly setting $C_\ell^{TT} = C_\ell^{EE} = 0$. In an ongoing effort to enable delensing with SPT3G¹⁷, a final set of $a_{\ell m}$ ’s are generated that include lensed- Λ CDM but have the B-modes explicitly set to zero (internally referred to as type-2 simulations), that are also utilized for systematics simulations in Chapter 5.

Armed with a set of input $a_{\ell m}$ ’s, we simulate the instrument’s observation by smoothing with the measured, array-averaged beam window functions (see section 4.1.1 for details). We build per-detector timestreams by using second order Taylor interpolation off the smoothed

¹⁵The term ‘pseudo’ refers to accounting for the effects of observation and filtering post power spectrum estimation.

¹⁶Alternative dust models have been put forward for example in [62]

¹⁷<https://spt3g.ncsa.illinois.edu>

input maps along the pointing trajectory of each detector, keeping track of their polarization orientation and timing information. The simulated timestreams are then filtered and binned into maps the same way as the real data following the prescription of section 2.7.2. The computational cost of this procedure is minimized by not considering every scanset within a season. Rather, only simulating the unique set of scans in El, Dk (since the azimuth scan range is the same for each scanset) within a given season(s) and reusing them for each repeated scanset. Irrespective of the number of repeated scansets, each scanset's weights and cuts are applied to the timestreams. Therefore, even though the scan patterns are reused, the act of observation is still unique to each scanset. The resulting reduction in data volume is roughly a factor of 50, although this varies quite a bit from season to season. The list of unique El, Dk scansets is called the 'minimum tag list' (MTL).

Once we have the per-season signal and noise simulations, we add them and spherical harmonic transform the sum to produce signal+noise power spectra. Finally, to establish consistency between the simulations and the real data, we correct each bandpower by the ratio of real to the mean of simulated bandpowers. No mention up to this point has been made of the theory bandpowers and how to measure cosmological parameters like r since this framework is only concerned with producing realizations of the data that help us understand the effects of filtering and noise under various input map parameters. We turn to that in section 2.8.

2.7.6 Null Tests

The end goal of our analysis is to compare the observed bandpowers to the theory model and infer cosmological parameters like r . Establishing data consistency by comparing the real data with itself and with simulations is imperative to ensure the robustness of our results. Null tests are a powerful tool for this assessment (also referred to as *jackknives*). During the pairmap coaddition stage, the data are split under various criteria based on the type of systematic contamination we are trying to probe. For example, if we are trying to probe for systematics related to observing conditions we split the data in the first half and second half of the observing season. After the final map coaddition stage, but prior to power spectrum

estimation, the various data splits are differenced to produce ‘null maps’ or ‘split maps’. The power spectrum estimation is carried out on both the real data and the null maps, and the resulting spectra compared with simulations. The full list of the 15 different null tests used in the BICEP/Keck analysis is summarized in [38]. Here we list the ones relevant to this thesis

- **Scan Direction:** Final maps are constructed from all left-going scans and compared with one made from all right-going scans. Probes fluctuations along the temporal direction and is relevant to the time it takes for a detector to respond to changes in the measured signal (see appendix A).
- **Tile:** Final maps are constructed from detectors from certain tiles and compare with those from other tiles. Relevant to tile-depended systematics based on the location of the tiles in the focal plane unit (FPU). Relevant in exploring edge-tile behavior that might couple differently to the sky (see section 5.2.1).
- **Tile inner/outer:** Final maps are constructed by considering detectors in the inner half of each tile versus the outer half. Relevant to consistency checks within a tile along the detector dimension. Important when examining the impact of crosstalk in residual beam in section 4.1.4.

For each bandpower, we calculate three quantities used to assess the consistency of the null spectra with simulations. The first is the ‘bandpower deviation’ δ_i where i is the index over our science bandpowers excluding the lowest one centered at $\ell = 17.5$, defined as:

$$\delta_i = \frac{d_i - \langle s_i \rangle}{\sigma(s_i)} \quad (2.39)$$

where d_i is the observed bandpower, $\langle s_i \rangle$ is the mean of the simulated bandpowers, and $\sigma(s_i)$ is the standard deviation of the simulated bandpowers. The next two quantities are the χ and χ^2 statistics defined as:

$$\chi = \sum_i \delta_i, \quad (2.40)$$

$$\chi^2 = (\mathbf{d} - \langle \mathbf{s} \rangle)^T \mathbf{C}^{-1} (\mathbf{d} - \langle \mathbf{s} \rangle) \quad (2.41)$$

where \mathbf{d} is the vector of observed bandpowers, $\langle \mathbf{s} \rangle$ is the vector of mean simulated bandpowers, and \mathbf{C} is the covariance matrix of the simulated bandpowers, all calculated for a given null test. The χ and χ^2 statistics are separately accumulated for each null test and are subject to the Kolmogorov-Smirnov (KS) [63] test, after which we extract a *p-value*. The passing criterion for a null test is such that $0.01 < p < 0.99$. For a detailed discussion of the distributions of the χ and χ^2 for each null test refer to [38] and [60].

2.8 Multi-Component Likelihood Analysis

The lensed- Λ CDM+foregrounds+ r model used in BICEP/Keck analysis is of the form [54]:

$$\mathcal{D}_{\ell, \text{total}}^{BB, \nu_1 \times \nu_2} = A_L \mathcal{D}_{\ell, \text{lensing}}^{BB} + r_{0.05} \mathcal{D}_{\ell, \text{tensor}}^{BB} + \mathcal{D}_{\ell, \text{foreground}}^{BB, \nu_1 \times \nu_2}. \quad (2.42)$$

where A_L is the effective lensing amplitude relative to the Λ CDM prediction, $r_{0.05}$ is the tensor-to-scalar ratio at the pivot scale of $k = 0.05 \text{ Mpc}^{-1}$, and $\mathcal{D}_{\ell, \text{foreground}}^{BB, \nu_1 \times \nu_2}$ is the foreground contribution originally proposed in [64]:

$$\begin{aligned} \mathcal{D}_{\ell, \text{foreground}}^{BB, \nu_1 \times \nu_2} &= A_d \Delta'_d f_d^{\nu_1} f_d^{\nu_2} \left(\frac{\ell}{80} \right)^{\alpha_d} \\ &+ A_{\text{sync}} \Delta'_s f_s^{\nu_1} f_s^{\nu_2} \left(\frac{\ell}{80} \right)^{\alpha_s} \\ &+ \epsilon \sqrt{A_d A_{\text{sync}}} (f_d^{\nu_1} f_s^{\nu_2} + f_s^{\nu_1} f_d^{\nu_2}) \left(\frac{\ell}{80} \right)^{(\alpha_d + \alpha_s)/2}. \end{aligned} \quad (2.43)$$

and contains the contributions from dust and synchrotron emission with amplitudes A_d and

A_{sync} respectively, defined at the characteristic scale of $\ell = 80$ where B-mode power is expected to peak for primordial gravitational waves. Their frequency scalings (denoted by f_d^ν and f_s^ν) are defined at 353 GHz for dust and 23 GHz for synchrotron and contain frequency spectral indices β_d and β_s respectively. The spatial power law indices are α_d and α_s and are assumed to be Gaussian-distributed. The last term describes the correlation between dust and synchrotron emission with correlation coefficient ϵ . Finally, Δ'_d and Δ'_s model potential suppression of correlated dust and synchrotron power between frequencies, hence their name: ‘decorrelation parameters’.

With the model in hand, we employ a Bayesian approach to simultaneously infer the contributions and parameters of the distinct components that together best fit the observed data. To do so, we calculate the joint likelihood of the ensemble of auto- and cross-spectra from our data and external maps from Planck and WMAP. These include the Planck 30, 44, 143, 217, and 353 GHz maps [65], and the WMAP K (23 GHz) and Ka (33 GHz) maps [66], used primarily to help constrain foreground components in their respective frequency ranges (as described in section 1.4). These external maps go through the reobservation and purification procedure described in section 2.7.5 after having their respective instrument beams deconvolved¹⁸. 66 *BB* auto- and cross-spectra are calculated from BK data and external maps. The final preparation step, referred to as ‘noise debiasing’, involves subtracting the mean of the spectra of the 499 noise realizations from our timestream-level simulations. To compare the observed bandpowers to the theory model, we need to account for the bin-to-bin correlations introduced by filtering, apodization, beam effects, and partial sky coverage. Specifically, we need to know how power mixes between multipole bins. We start by generating 100 signal-only realizations with power sourced from a single ℓ bin from $\ell = 0 - 999$. We then apply the observation matrix to each realization and calculate the resulting bandpowers. The result is a direct estimate of the bandpower window functions (BPWFs) that describe how power from a given input multipole bin ℓ contributes to all other output multipole bins ℓ' . The resulting BPWFs are shown in Figure 2.20. The resulting spectra are shown in Figure

¹⁸The 3-year BICEP3 observation matrix is used for the reobservation process. Reobservations includes beam smoothing, filtering, and apodization applied the same way as the BK data.

2 of [27].

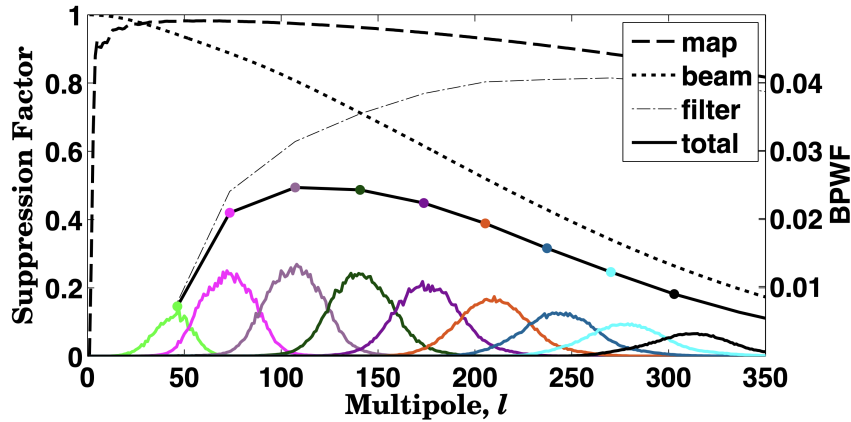


Figure 2.20. Total BB Suppression factor for each multipole bin for the BK18 analysis, illustrating the effects of beam smoothing, filtering, and apodization on how power is sourced from different ℓ bins. Figure from [58].

We are now ready to compare with theory models using our Bayesian framework. Because the joint likelihood of the ensemble of bandpowers is in general non-Gaussian¹⁹, we employ the Hamimeche-Lewis (HL) approximation [67] which requires a fiducial theory model²⁰ $C_{\ell,f}$ producing a fixed set of bandpowers used to define the non-linear transformation that makes the distribution of the bandpowers approximately Gaussian. The HL likelihood \mathcal{L} is expressed as:

$$\begin{aligned}
 -2 \log \mathcal{L}(\hat{\mathbf{C}} | \mathbf{C}) &= \mathbf{X}^T \mathbf{M}_f^{-1} \mathbf{X} \ , \\
 \mathbf{X}_\ell &= \text{vecp} \left[\mathbf{C}_{\ell,f}^{1/2} g \left(\mathbf{C}_{\ell,f}^{-1/2} \hat{\mathbf{C}}_\ell \mathbf{C}_{\ell,f}^{-1/2} \right) \mathbf{C}_{\ell,f}^{1/2} \right] .
 \end{aligned}
 \tag{2.44}$$

where $\hat{\mathbf{C}}_\ell$ are the observed bandpowers, M_f is the bandpower covariance matrix (BPCM) of the fiducial model, X_i is the non-linear transformation of the observed bandpowers that diagonalizes M_f , and $g(x) = \text{sign}(x - 1) \sqrt{2(x - \ln x - 1)}$. is a matrix function applied to a symmetric positive-definite matrix with the understanding that g is applied to its eigenvalues. Finally, vecp is an operator that vectorizes the distinct elements of a symmetric matrix. With the ensemble likelihood in hand, we perform maximum likelihood search (for details refer to Chapter 4 of [45]) to get the best-fit parameters and use a Markov Chain Monte Carlo

¹⁹Due to partial sky coverage, filtering, and presence of correlated and uncorrelated noise.

²⁰Usually taken to be an $r = 0$ standard Λ CDM model from Planck derived parameters, effectively defining the null hypothesis.

(MCMC) using the `CosmoMC` software [68] (for more details refer to Chapter 5 of [54]). The current constraints on r from the BK18 analysis are $r_{0.05} < 0.036$ at 95% with sensitivity to $\sigma(r) \sim 0.009$ [27] and illustrated in Figure 2.21.

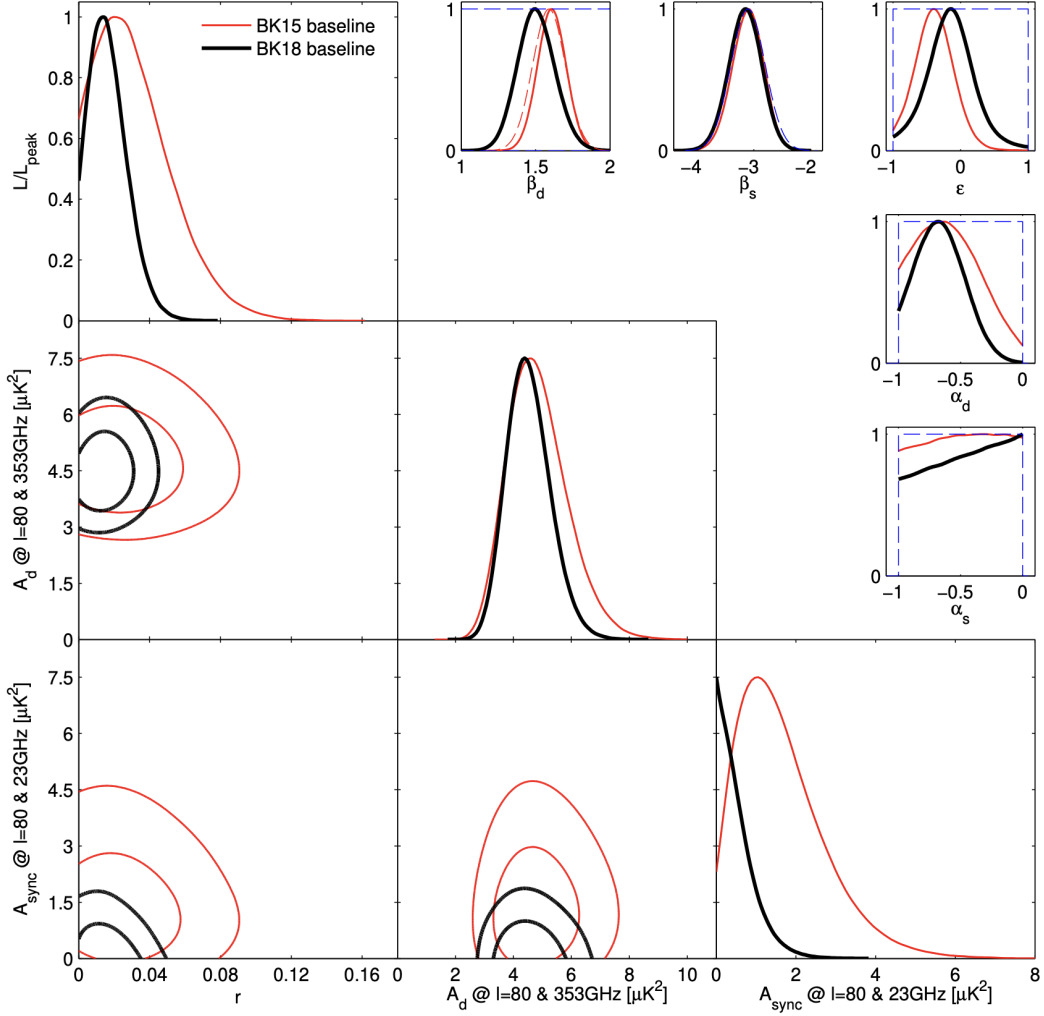


Figure 2.21. MCMC results from the BK18 multi-component, multi-frequency analysis showing cosmological parameter posterior distributions of r , A_{dust} , and A_{sync} . Dashed lines correspond to informed priors. Figure from [27].

Chapter 3

Calibration Measurements Overview and Analysis

For all our small aperture receivers, we build detectors and optics to be as close to ideal as possible, but imperfections are inevitable. We want to capture any imperfections and account for them in our analysis. In addition, our CMB analysis pipeline requires knowledge of the as-built characteristics of our instruments. One of these characteristics is the beam pattern of our detectors.

When observing the CMB sky, we measure the convolution of the true sky signal with the beam pattern (equation 4.2). To detect B-mode excess power at degree angular scales, we need the deepest understanding and highest control of the detector beam patterns. In fact, many design choices for BICEP/Keck receivers are made with such control in mind (i.e. forebaffle and ground shield design, optics materials choices, anti-reflection coatings etc.), to mitigate optical coupling with unwanted surrounding microwave emitters. Hence, characterizing the beam patterns of our deployed detectors is crucial for mitigating systematic errors in our CMB polarization measurements. We take advantage of the fact that convolving the beam pattern with a delta function source (i.e. a point source) yields the beam pattern itself. We call this process ‘beam mapping’. There are various sources to measure the beam patterns of our detectors including some astronomical sources like planets (e.g. Jupiter, Saturn) and

the Moon. A notable example of a polarized astronomical source is Taurus-A (Crab Nebula) [69] that could in principle be used for mapping the polarized beam response of our detectors. However, such sources are only available during certain times of the year (if at all) and their brightness and angular size may not be optimal for detailed beam characterization. In addition, their polarization signature could be complicated, uncertain, and hard to reconstruct to the precision required for our systematics studies.

We require stable, bright, and well-characterized sources that can be deployed at the South Pole site during the Austral summer season when all our calibration activities take place. Since we rely on calibration data products for our main CMB analysis (e.g. beam window functions—see section 2.7.5), we prioritize beam mapping during the months of the Austral summer when the South Pole Station is accessible. In this chapter, we describe the far-field beam mapping (FFBM) calibration measurements, the sources used, and the data analysis pipeline. The main focus is on establishing the framework for FFBM, which requires deep understanding of the instrument pointing model and its extensions for calibration applications.

3.1 Beam Mapping

When beam mapping, there are typically two regimes to consider: the near field (Fresnel region) and the far field (Fraunhofer region). Loosely speaking, the near field is the region close to the aperture of the telescope where the angular beam pattern depends on both angle from the beam center and distance traveled. We perform near field beam mapping (NFBM) to characterize the illumination of the telescope’s optics and check for misalignment in the optical system. It is a crucial measurement when commissioning new receivers to identify any major issues in the construction of the optics chain. These are typically performed by scanning a thermal source just above the window of the receiver. This thesis focuses on systematics related to the far-field beam patterns, hence we defer the discussion of NFBM to Chapter 4.3 of [51].

3. Calibration Measurements Overview and Analysis

To get to the far-field regime we require that the phase difference of the wavefront across the aperture is small. Defining the aperture diameter to be D and the wavelength λ . The path difference between the center and the edge of the aperture at a distance d along the optical axis is

$$\Delta L = \sqrt{d^2 + (D/2)^2} - d. \quad (3.1)$$

For $d \gg D$, we can use a binomial expansion:

$$\sqrt{d^2 + (D/2)^2} \approx d \left(1 + \frac{(D/2)^2}{2d^2} \right) = d + \frac{D^2}{8d}. \quad (3.2)$$

The far field is defined as the region where the phase variation across the aperture is sufficiently small e.g.,

$$\Delta L \lesssim \frac{\lambda}{16}. \quad (3.3)$$

Substituting, we obtain

$$d_{\text{far}} \gtrsim \frac{2D^2}{\lambda}. \quad (3.4)$$

Table 3.1 summarizes the far-field distances for BICEP3 and BICEP Array receivers at their band-center frequencies.

Receiver	Band (GHz)	d_{far} (m)
BICEP3	95	171
BA30/40	30	61
BA30/40	40	79
BA150	150	303
BA220/270	220	444
BA220/270	270	545
K270	270	129

Table 3.1. List of far-field distances for BICEP3 and BICEP Array receivers at the configuration as of the writing of this thesis. The aperture diameter is 52 cm for BICEP3, 55 cm for BICEP Array-style receivers, and 26.4 cm for the Keck-style receiver.

It is clear that the transition boundary from near-field to far-field is continuous and depends

on the criterion of acceptable phase variation across the aperture. The far-field distance for our small aperture telescopes (SATs) is on the order of a few hundred meters, which is manageable compared to larger aperture telescopes (LATs) for which the far-field distance could be orders of magnitude larger depending on their design. We take advantage of the South Pole site facilities that are spaced roughly 200 meters apart to mount our calibration sources. However, this distance is well below the criterion for some BA receivers (e.g. BA220/270). The effects of being in the quasi-far field are explored in [51] and show that a curvature in the wavefront changes the beam window function $B(l)$ by less than 1% at $\ell \sim 200$. Notably, similar analysis needs to be carried out to verify this claim for the high frequency BICEP Array receivers (i.e. BA150 and BA220/270).

Beam Mapping Regions

In addition to the near- and far-field regimes, we also need precise terminology to differentiate between different regions of the detector’s angular response. The division proposed in figure 3.1 is loosely based on the BICEP1 beam characterization results (Figure 17 in [48]), and is meant to define the naming scheme of the calibration measurements discussed throughout the rest of this thesis.

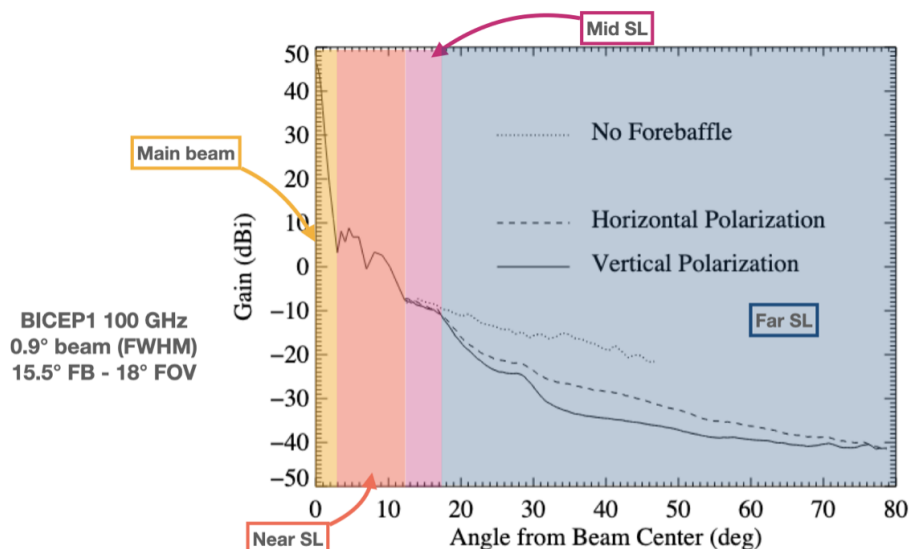


Figure 3.1. Schematic azimuthally averaged beam profile of the BICEP1 instrument. Adapted from Figure 17 in [48] to define the naming scheme of the beam mapping regions used in this thesis (see text for more details).

The angular response of our detectors is composed of three main regions:

- **Main beam:** the region starting at the detector pointing direction out to the first null of the Airy pattern response. This varies from receiver to receiver depending on the aperture illumination and taper ¹, along with observing frequency. It generally extends up to $\mathcal{O}(1^\circ)$ for example it is $\sim 0.5^\circ$ for BICEP3.
- **Near and mid sidelobes:** the region between the main beam and the forebaffle cutoff (see section 2.4). It is split between the near sidelobes that include ‘optically adjacent’ beam features, and the mid sidelobes that represent the intermediate region between the near sidelobes and the forebaffle cutoff.
- **Far sidelobes:** The region beyond the forebaffle cutoff, where the beam response is dominated by spillover and scattering from the forebaffle and ground shield.

There are various terms used to refer to different regions of an antenna’s angular response, and the naming scheme presented here may not necessarily be consistent with the rest of the literature. Regardless, it has become the norm within the BICEP/Keck Collaboration and is used throughout this thesis. We use far-field beam mapping with a variety of sources to characterize the beam response of our detectors in each region, which is what we turn to next.

3.2 Far Field Beam Mapping Setup

Deploying a source ~ 200 meters away from the telescope on top of a 10-meter tall mast mounted at the roof of the opposite building, puts the source near the horizon as shown in Figure 3.2.

BICEP3 and BICEP Array are housed in large reflecting ground shields to mitigate ground coupling and sidelobe pickup, that sometimes interrupt the line of sight to the source at

¹A tapered aperture truncation substantially suppresses Airy rings, which is the case for BA150. However, the main beam region is still well-defined.




Figure 3.2. Schematic showing the South Pole Station facilities and instrument setup for far-field beam mapping. Image slightly modified from BK-IX [70], courtesy of the BICEP/Keck Collaboration.

low elevation angles. In addition, due to the design of the cryogenic system our telescopes cannot maintain thermal control when tipped below 45° in elevation for extended periods. To overcome these challenges, we deploy redirecting flat mirrors mounted such that when the instruments are pointed at zenith the mirror reflects the source signal into the telescope. In the case of BICEP3, the ground shield is tilted in the direction of the source to ensure that all the beams clear it, since the mirror is not high enough above the window to both cover the entire field of view and fully clear the ground shield. For both instruments, the mirror is mounted on the Azimuth and Elevation stage to ensure tandem movement with the mount motion in those directions but independent of boresight rotation Dk . Separate mirrors are used for BICEP3 and BICEP Array instruments, with the former being $(1.7 \times 2.5m)$ and the latter $(1.8 \times 2.7m)$ in size. The mirror used for BICEP Array is the same as the previous generation instrument (Keck Array) so we will be referring to it as the *Keck mirror*. The mounting scheme involves setting the mirror on a set of carbon fiber struts attached to a steel triangular frame at the back of the mirror, with a set of pulleys to adjust its orientation. Utmost attention is paid to ensure lateral symmetry. Once in place, we use the mount motion in elevation (El), to lift the mirror into position. Figure 3.3 shows the mirror mounted on BICEP3 pointed at zenith ready for FFBM measurements. For BICEP Array, the mounting scheme allows for forward-backward movement along the line connecting the mirror and the source, such that lateral symmetry is preserved.



Figure 3.3. BICEP3 mirror setup for FFBM, courtesy of the BICEP/Keck Collaboration

BICEP3’s mirror covers each detector’s beam out to a radius of $> 2^\circ$ for any rotation of the Dk stage. Its beam mapping strategy has stayed unchanged after being optimized in 2018 [51]. The scan strategy is summarized in Table 3.2, with the caveat that BICEP3’s ground shield is tilted roughly 15° in the direction of the source. This will become important in chapter 6 when discussing ground mapping measurements. Having a mirror large enough to ensure similar coverage for each BICEP Array instrument is unfeasible. Discussions of sourcing a larger mirror optimized for BICEP Array have paused given the amount of resources needed to design, build, and ship it to the South Pole site. The Keck mirror only partially covers the field of view of any given BICEP Array-style instrument. To ensure a beam coverage $> 2^\circ$ we implement four different mirror positions ‘*Backest*’, ‘*Back*’, ‘*Back-ish*’, and ‘*Forward*’. These somewhat obscure names refer to the position of the mirror-center with respect to the middle of the Dk stage as viewed from above. The combination of multiple mirror positions and 16 uniformly-spaced Dk angles allows us to achieve the 2° coverage goal. In the limit of infinite calibration time, we could map all BICEP Array detectors out to this radius. However, the tight scheduling constraints of the Austral summer season and the ever-decreasing personnel allowed on site, require us to optimize the scan strategy to meet these constraints while maximizing the number of detectors being mapped. The nominal series of beam mapping schedules has been extended in both Azimuth and Elevation for BICEP Array, to cover a

total of 34° in Az and 33.8° in El and last ~ 9 hours (on average) for a given Dk angle.

Instrument	Az Scan Speed (deg/s)	Az Range (deg)	El Step (deg)	El Range (deg)
BICEP3	0.4	30°	0.05°	30°
BICEP Array	0.9	34°	0.05°	33.8°

Table 3.2. Far-field beam mapping schedule parameters for BICEP3 and BICEP Array instruments.

The final ingredient for our scan strategy is the type of scan. For BICEP3 we perform *out-and-back* scans where the telescope scans in azimuth at a fixed elevation from one edge of the azimuth range to the other and then back to the starting point. Once complete, a 0.05° elevation step is taken and the process repeats. Such scans allow deep measurements of the beam patterns, and the construction of a *scan direction ‘jackknife’* (see section 2.7.6) maps by subtracting right-going from the left-going scans. The result is a good estimate of the measurement noise discussed in further detail in chapter 5. For BICEP Array, we perform *snake raster* scans where the telescope steps in elevation at the end of each azimuth scan. Due to the small step in elevation the impact on the final beam quality is negligible, although measurement noise estimation via scan-direction jackknife maps is not reliable.

3.3 Sources

Depending on the calibration measurement goal, we have the option of deploying either polarized or unpolarized sources. Bright and stable microwave sources are specifically designed for BICEP3 and BICEP Array calibrations. In this section, we give a brief overview of the sources used in our calibration measurements.

3.3.1 Thermal chopper

The unpolarized sources nominally used for FFBM, yielding data products that directly interface with our CMB analysis, were built in 2015. They are chopped thermal sources deployed atop 10 meter tall masts raised at the roof of the facilities shown in Figure 3.2. They consist of a cylindrical enclosure, housing a two-spoke chopper blade made of a composite

material of carbon fiber, low-density foam, and Nomex honeycomb fiber, that is coated in a microwave-absorbing material (HR-10). The enclosure is further encased in aluminum tape to scatter unwanted rays from surrounding microwave emitters as shown on the left of Figure 3.4. Behind the chopper blade is a 45° aluminum reflector that redirects to zenith ($\sim 10K$). The chopper frequency is nominally set to 16 Hz (motor frequency 8 Hz), alternating between ambient temperature (~ 250 K) and the cold sky. The 24-inch chopper aperture is sealed to prevent wind from entering the enclosure and affecting the stability of the chopper rotation.

3.3.2 Broad Spectrum Polarized Source

Additional sources used in FFBM (and sometimes in other configurations) are linearly polarized broad spectrum noise sources, at various frequencies. They consist of a quasi-thermal, bandpass filtered, broad spectrum noise source (BSNS) that is amplified to the desired output frequency range. To achieve the broad spectrum output, a 50Ω SMA terminator resistor is used as a source of thermal radiation. Its output is fed into a frequency multiplier chain and is bandpass filtered. It is coupled to free space via feedhorn, after passing through a fine mesh polarizer grid ensuring off-axis emission of $< 0.01\%$. We add a degree of freedom by mounting the feedhorn and wire grid on a rotation stage that registers the angle of the output polarization with respect to gravity via a high-precision tilt meter. A detailed description of the rotating polarized source (RPS) design and construction is in [40]. A set of attenuators allows for precise control of the BSNS peak output power that is adjusted depending on the calibration measurement performed (further discussed in Chapter 4). The source construction is shown on the right of Figure 3.4.

The 100 GHz BSNS is currently at the South Pole and has been used for detector polarization angle measurements (ϕ_d in equation 2.8) [72] and cross-polar response characterization for BICEP3 explored in Chapter 4.2. An additional BSNS source centered at 40GHz output, was built during the summer of 2019 for sidelobe mapping of the BA30/40 GHz detectors first deployed during Austral summer of 2020 (discussed in Chapter 4.3). A third BSNS source centered at 150GHz was built during the Fall of 2025, awaiting on-site deployment for




Figure 3.4. Calibration sources. Left: thermal chopper used for unpolarized beam mapping. Modified from Figure 1. in BK-IX [70]. Right: rotating polarized source (RPS) used for polarized beam mapping. Figure from [71]. (A) BSNS housing enclosure, (B) rectangular feedhorn, co-linear with the rotation stage’s axis of rotation, (C) wire grid polarizer location, (D) high-precision tilt meter.

BICEP Array-150 GHz receiver calibrations during the 2026-2027 Austral summer season. For the purposes of this thesis, a calibration measurement taken with the 100 GHz linearly polarized source at a fixed source polarization angle (no rotation of the source’s stage) will be referred to as ‘BSNS campaign’, while a measurement taken with the source at multiple source polarization angles (various rotations of the source’s stage) will be referred to as ‘RPS campaign’.

3.4 Coordinate Systems

During CMB observations we can safely assume that the source is effectively at infinity. Therefore, parallax effects arising from observing at different boresight orientations can be ignored. On the contrary, our calibration sources are located on the opposite building from the receiver being calibrated roughly ~ 200 meters away (see Figure 3.3). Therefore, parallax effects become important. This is especially relevant for BA since each receiver is offset from the central boresight ray (see Figure 2.12). In this section we describe the process of correcting for parallax during beam mapping, introducing the various coordinate systems utilized in calibration data analysis, along the way. Since the parallax correction calculation involves knowledge of the finite distance between the instrument and the source being mapped we need to pick the origin for the horizon coordinate system. It is defined at the intersection

3. Calibration Measurements Overview and Analysis

point of the three mount axes, with the x-axis pointing in the direction of zero azimuth. For BICEP Array, there is no point that all three mount axes intersect, so the zero point is defined (somewhat arbitrarily) at the height of the ground shield floor.

To account for the mirror reflection and parallax effects, we need to know the location(s) of each receiver's aperture(s). Vector \vec{A} (in units of meters) is defined to be the location in 3D space that all the detector pointings intersect before emerging out of the receiver. It is different from the location of the boresight \vec{B} since the location of the aperture is not necessarily unchanged with Dk rotation.

$$\vec{A} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad (3.5)$$

We apply a series of translations and solid rotations based on mount drawings to locate the exact position of a given receiver's aperture with respect to the boresight ray in 3D space. All the information to pinpoint the aperture location and orientation for each receiver is now contained in $(\vec{A}, \hat{B}_3, \hat{B}_1)$ with $\hat{B}_2 = \hat{B}_3 \times \hat{B}_1$ (see equation 2.1). For BICEP Array we locate four apertures corresponding to each receiver on the Dk drum by noting that each receiver sits on the drum 90° apart from the next.

An optically equivalent way to think of the mirror reflection is to introduce an aperture behind the mirror at the same distance as the vertical distance between the mirror and the physical aperture. This is the 'apparent aperture' position \vec{A}' which introduces a set of coordinates we use during calibration campaigns—the apparent topocentric coordinates $(Az_{ap}, El_{ap}, PA_{ap})$. To correctly locate the apparent aperture position we need information about the mirror mounting and orientation scheme.

We start with the mirror normal vector at the origin pointing towards the negative \hat{z} axis:

$$\hat{n} = \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \quad (3.6)$$

such that when the instrument is pointed at zenith the mirror normal vector points in the opposite direction of the boresight ray (i.e. towards the cryostat). In general, we need six parameters to fully describe the position of the mirror in 3D space, three for its orientation and three for its location. We use Euler angles to specify orientation: (*Tilt*, *Roll*, *Pitch*). *Tilt* is defined as the rotation of the mirror about an axis parallel to elevation, while *Roll* is the rotation about an axis along the plane of the mirror, perpendicular to both the tilt axis of rotation, and the mirror normal vector. *Pitch* is defined as the final rotation axis resulting from the right-handed cross product of the tilt and roll axes. Since the mirror mounting scheme does not allow for such movement, it is disregarded. In practice, we enforce lateral symmetry which sets *Roll* very close to zero. The tilt is set to 45° for most mirror positions, except for the *Forward* mirror position of BICEP Array which is set to 41° to appropriately match the optical path with the rest of the mirror positions. With the tilt and roll parameters in hand we apply the solid rotation per equation 2.2 as $R(0, \textit{tilt}, 0)$ and then $R(90, \textit{roll}, 0)$. These rotations do not commute, but since the roll parameter is so close to zero, the order in which they are applied makes negligible difference. Finally, we translate and rotate the normal vector in a similar way to the aperture basis vectors as described in § 2.6 and apply the mirror height along the \hat{z} direction. The mirror height h is defined as the distance along the center-axis of the *Dk* stage (i.e. \hat{B}_3 when pointing at zenith), from the plane of the aperture(s) to the intersection with the mirror plane. The exact distance comes from *in situ* measurements and cross-referenced with the SolidWorks model of the mirror. This defines the location that the boresight pointing vector intersects the mirror denoted by point D . Hence:

$$\vec{D} = \vec{A} + h\hat{z}. \quad (3.7)$$

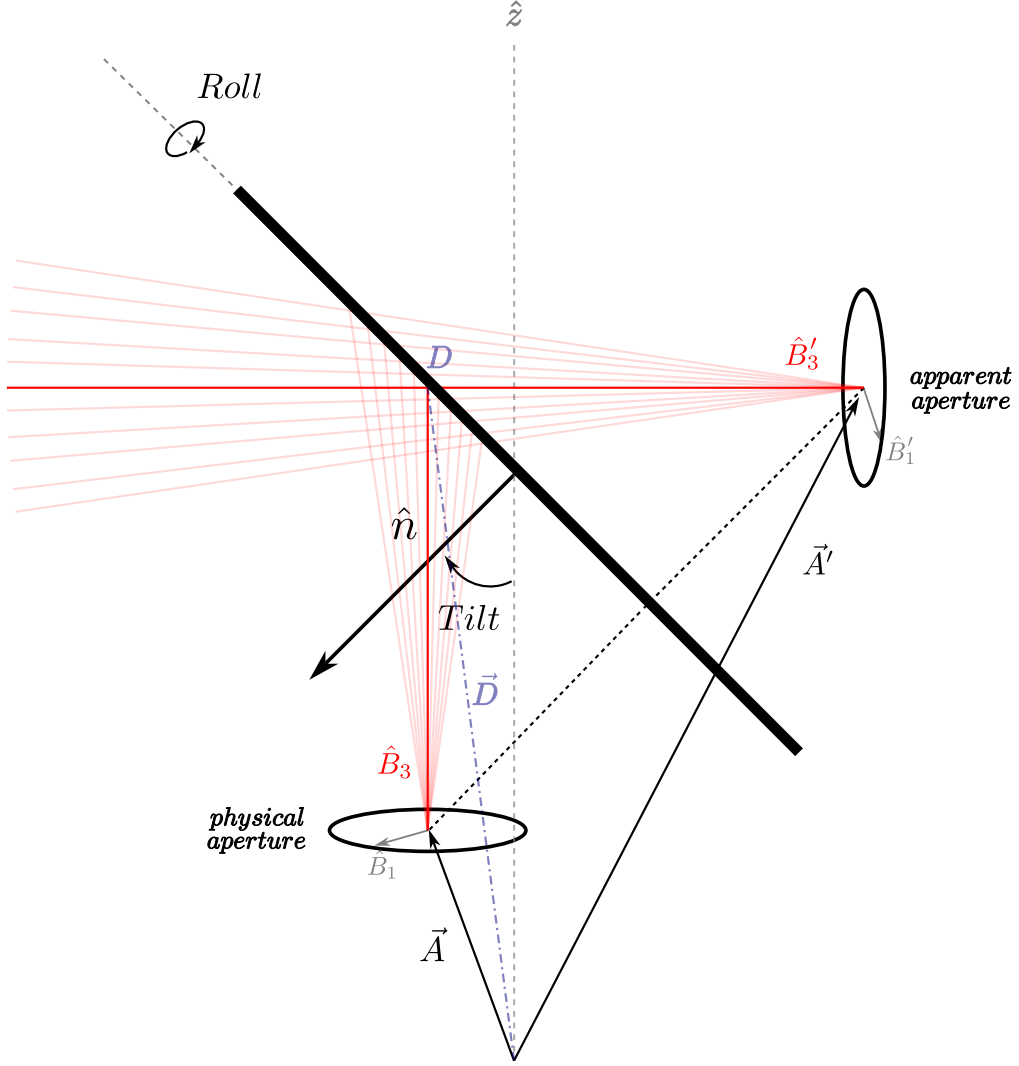


Figure 3.5. Mirror reflection setup.

With the help of Figure 3.5 which shows a sketch of the mirror reflection geometry, the reflected pointing and orientation (\hat{B}'_3, \hat{B}'_1) can be calculated as:

$$\hat{B}'_3 = \hat{B}_3 - 2(\hat{B}_3 \cdot \hat{n})\hat{n} \quad (3.8)$$

$$\hat{B}'_1 = \hat{B}_1 - 2(\hat{B}_1 \cdot \hat{n})\hat{n} \quad (3.9)$$

$$\hat{B}'_2 = - \left[\hat{B}_3 \times \hat{B}_1 - 2((\hat{B}_3 \times \hat{B}_1) \cdot \hat{n}) \hat{n} \right] \quad (3.10)$$

where the mirror reflection is accounted for by flipping the parity of the second orientation vector. Therefore, the apparent boresight basis vectors are now defined in the left-handed sense. Note that for the rest of this section, prime (') denotes apparent quantities. We also need to calculate the position of the apparent aperture with respect to the origin which is simply:

$$\vec{A}' = \vec{A} - 2 [(\vec{A} - \vec{D}) \cdot \hat{n}] \hat{n} \quad (3.11)$$

With the apparent boresight basis vectors and the apparent aperture position, we define a spherical coordinate system centered at the apparent aperture position, using the \hat{B}'_3 and \hat{B}'_1 directions. From there, we calculate each detector's pointing vector as in equation 2.4. The resulting detector-centered coordinate vectors are given by:

$$\hat{P}'_1 = (\sin^2 \theta + \cos r \cos^2 \theta) \hat{B}'_1 + ((1 - \cos r) \sin \theta \cos \theta) \hat{B}'_2 - (\sin r \cos \theta) \hat{B}'_3 \quad (3.12)$$

$$\hat{P}'_3 = (\sin r \cos \theta) \hat{B}'_1 - (\sin r \sin \theta) \hat{B}'_2 + (\cos r) \hat{B}'_3 \quad (3.13)$$

$$\hat{P}'_2 = -\hat{P}'_3 \times \hat{P}'_1 \quad (3.14)$$

and are calculated for each detector by only using θ and r . With the apparent position and orientation of each detector we can correct for parallax. The method used follows the procedure in [56]. The main parameters needed is the distance and height of the source in horizontal coordinates. A source a distance s_d , height s_h , and azimuth position AZ_S has a position vector:

$$\vec{S} = \begin{bmatrix} s_d \cos AZ_S \\ -s_d \sin AZ_S \\ s_h \end{bmatrix} \quad (3.15)$$

The source distance is approximately 200 meters and the height is $\sim 13m$. The corrected

3. Calibration Measurements Overview and Analysis

Az/El values are calculated by finding the intersection of a pointing vector (\hat{P}'_3) originating at the apparent aperture position (A') with a sphere of radius equal to the source distance (s_d), that is centered at the origin of the horizontal coordinate system.

The parallax-corrected, apparent topocentric azimuth and elevation coordinates for a given detector are:

$$\begin{aligned} AZ_{ap} &= \tan^{-1} \left(-\frac{\hat{P}'_3 \cdot \hat{y}}{\hat{P}'_3 \cdot \hat{x}} \right), \\ EL_{ap} &= \sin^{-1} \left(\hat{P}'_3 \cdot \hat{z} \right), \\ PA_{ap} &= \tan^{-1} \left(-\frac{\hat{P}'_2 \cdot \hat{z}}{\hat{P}'_1 \cdot \hat{z}} \right) \end{aligned} \quad (3.16)$$

This spherical coordinate system centered at each detector is useful for a number of diagnostic checks as well as some filtering and masking operations (to be described in the next section 3.5). It is also used for binning the ground calibration data that we delve into detail in chapter 6.

The next step is to examine how the source and its co-polar axis map onto the detector-centered coordinates. We define the unit vector from the apparent aperture to the source as:

$$\hat{R} = \frac{\vec{S} - \vec{A}'}{\|\vec{S} - \vec{A}'\|} \quad (3.17)$$

where \vec{A}' is the location of the apparent aperture and \vec{S} the location of the source in horizontal coordinates, both with units of meters (also drawn in Figure 3.8). The source can then be projected onto the detector-centered coordinate system as:

$$\begin{aligned} r_S &= \cos^{-1} \left(\hat{R} \cdot \hat{P}'_3 \right), \\ \theta_S &= \tan^{-1} \left(\frac{-\hat{R} \cdot \hat{P}'_2}{\hat{R} \cdot \hat{P}'_1} \right) \end{aligned} \quad (3.18)$$

3. Calibration Measurements Overview and Analysis

$$\begin{aligned} x' &= r_S \cos(\theta_S) \\ y' &= r_S \sin(\theta_S) \end{aligned} \tag{3.19}$$

If we wanted to get the coordinates of the source with respect to the boresight center, we would replace the detector basis vectors $(\hat{P}'_1, \hat{P}'_2, \hat{P}'_3)$ in equation 3.18 with the boresight basis vectors $(\hat{B}'_1, \hat{B}'_2, \hat{B}'_3)$ where appropriate. The resulting coordinates are the boresight centered coordinates denoted by (X, Y) and are useful to probe systematic effects that have to do with the instrument's pointing. The detector pointings onto the sky and the focal plane orientation are shown in Figure 3.7

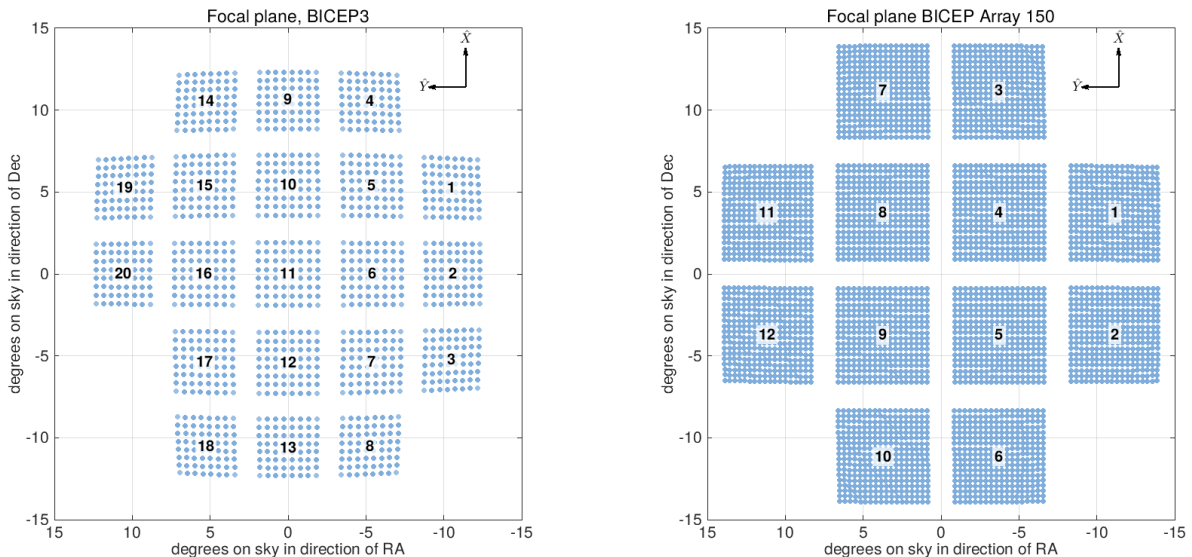


Figure 3.7. Pointings on sky for BICEP3 (left) and BICEP Array 150 GHz (right).

In the case of a polarized source, we need to specify the polarization orientation in addition to the source projection onto the focal plane. Consider the polarization orientation of the source \hat{S}_{pol} , and also its projection to the apparent focal plane \hat{S}'_{pol} —the apparent source polarization axis. To do this, we first need to identify the plane perpendicular to the source pointing vector \vec{S} denoted by unit vectors \hat{e}_b and project \hat{S}_{pol} onto it. The geometry is shown in Figure 3.8.

$$\hat{e}_b = \frac{\hat{S}_{pol} \times \hat{S}_3}{\|\hat{S}_{pol} \times \hat{S}_3\|} \quad (3.20)$$

$$\hat{S}'_{pol} = \hat{S}_3 \times \hat{e}_b$$

The source orientation with respect to the detector's orientation axis is then:

$$\phi_s = \tan^{-1} \left(\frac{-\hat{P}'_2 \cdot \hat{S}'_{pol}}{\hat{P}'_1 \cdot \hat{S}'_{pol}} \right) \quad (3.21)$$

where I have used the apparent detector-centered basis that accounts for the mirror reflection and is parallax-corrected. The main assumption baked in equation 3.21 is that \hat{S}'_{pol} points along \hat{P}'_3 , which is true to a good approximation due to our source mounting scheme [40].




Figure 3.8. Source coordinates geometry setup.

Finally, we define ϕ to be the angle of the apparent source orientation \hat{S}'_{pol} and the detector's

co-polar axis \hat{P}'_{pol} as:

$$\phi = \phi_d - \phi_s \quad (3.22)$$

where ϕ_d is the angle of the detector's co-polar axis with respect to the boresight orientation axis defined in equation 2.6. In practice, we calculate ϕ_s from the pointing model using the source information (source Az , distance, and height), and ϕ_d is a measured quantity from the RPS campaign [40]. Knowledge of ϕ is crucial for polarized beam mapping (section 4.2) since it allows us to accurately reconstruct the polarization field of a polarized source.

At this point we have defined all the relevant coordinate systems and transformations needed to locate a source (polarized or not) in the far-field using a reflecting mirror. A final coordinate system that has proven useful during analysis of calibration data is the mirror-centered coordinate system $(\hat{x}_m, \hat{y}_m, \hat{n})$. The motivation for looking at calibration data is to examine the case where a detector's pointing might completely or partially miss the reflecting mirror. This is not an issue for BICEP3 with the beam-filling mirror. For BICEP Array the mirror does not cover the entire field of view of all receivers at the same time, for any mirror position or Dk combination. The following description of the mirror-centered coordinate basis exclusively applies to BICEP Array receivers.

For this coordinate system, we are not concerned with the finite size of the detector beam on the mirror. Rather, we are only interested whether the ray along the detector's pointing direction (\hat{P}_3) intersects the mirror plane. I will be referring to this ray as the detector *pencil beam* that can be thought of as a delta function beam along the detector's pointing direction. The framework extends the existing pointing model for calibration measurements to return how the pencil beams map on the mirror plane as opposed to the source. Therefore, there are no parallax considerations. The geometry of this coordinate system is shown in Figure 3.9 where BICEP Array is pointed at some arbitrary azimuth (Az), elevation (El), and boresight orientation (Dk).

The mirror normal vector \hat{n} is defined the same way as in equation 3.6 and is rotated to the

3. Calibration Measurements Overview and Analysis

mirror position by using the pointing model. Next we find the plane of the mirror spanned by:

$$\hat{m}_x = \frac{\hat{n} \times \hat{z}}{\|\hat{n} \times \hat{z}\|} \quad \hat{m}_y = \frac{\hat{n} \times \hat{m}_x}{\|\hat{n} \times \hat{m}_x\|} \quad (3.23)$$

defined in the right-handed sense. This completes the orthonormal basis of the mirror coordinate system $(\hat{m}_x, \hat{m}_y, \hat{n})$. However, we still need to define the origin of this coordinate system on the mirror plane. We pick it (somewhat arbitrarily) to be at the center of the mirror. There were a few choices (i.e. any of the four corners of the mirror), but the center is the most intuitive and has rectangular symmetry. The only thing left is to locate the mirror center \vec{m}_c in horizon coordinates. To do this, notice that for a given mirror position, the distance from the center of the drum along the boresight ray to the mirror \vec{d} , does not change as we scan in Az , step in El , and rotate in Dk . Suppose \vec{d} intersects the mirror at point L . In terms of the perpendicular distance \vec{D} we used to get the mirror height in equation 3.7, we calculate \vec{d} in horizon coordinates as:

$$\vec{d} = |\vec{D}| [\sin(90^\circ - El) \cos(Az) \hat{x} + \sin(90^\circ - El) \sin(Az) \hat{y} + \cos(90^\circ - El) \hat{z}] \quad (3.24)$$

Such that when the instrument is pointed at zenith, $\vec{d} = \vec{D}$ and points L and D coincide. We take advantage of the mirror mounting scheme that enforces lateral symmetry (i.e. Roll ≈ 0) to claim that the vector along the mirror plane from point L to the mirror center \vec{m}_c is along the \hat{m}_x direction. This vector is denoted by \vec{L} . This somewhat unintuitive distance is measured from the SolidWorks model of the mirror and is the only parameter we need to fully specify the mirror center position (on top of the other mirror and mount parameters discussed this far). The mirror center position is then:

$$\vec{m}_c = \vec{B} + \vec{d} + |\vec{L}| \hat{m}_x \quad (3.25)$$

Where \vec{B} is the position of the center of the BICEP Array boresight in horizon coordinates

(also shown in Figure 3.9). The end goal is to find each detector's pencil beam intersection point with the plane of the mirror, or equivalently, its coordinates in the mirror-centered coordinate system. For a detector j , these will be denoted by $m_x^{(j)}, m_y^{(j)}$. Equation 2.4 gives us the detector pointing vector $\hat{P}_3^{(j)}$ and orientation $\hat{P}_1^{(j)}$ vectors. With those, we can read off the mirror-centered coordinates for this detector by following the procedure:

- Find the vector from the origin of the horizon coordinates $(\hat{x}, \hat{y}, \hat{z})$ to the aperture of a given BICEP Array receiver \vec{A} .
- Get the distance along the detector pointing to the mirror plane (i.e. the dashed line in Figure 3.9 that originates at \vec{A} and terminates at mirror plane), which is not necessarily the perpendicular distance, nor \vec{d} .
- Subtract the mirror center position to get a vector along the mirror plane, not necessarily along any of the mirror basis vectors.
- Project onto the mirror basis vectors (\hat{m}_x, \hat{m}_y) to get the mirror-centered coordinates of the detector's pencil beam.

This is summarized in the following set of equations:

$$\begin{bmatrix} m_x^{(j)} \\ m_y^{(j)} \end{bmatrix} = \left(\vec{A} + \frac{(\vec{m}_c - \vec{A}) \cdot \hat{n}}{|\hat{P}_3^{(j)} \cdot \hat{n}|} \hat{P}_3^{(j)} - \vec{m}_c \right) \cdot \begin{bmatrix} \hat{m}_x \\ \hat{m}_y \end{bmatrix} \quad (3.26)$$

To effectively interpret our calibration data, and identify which detectors are fully or partially illuminating the mirror, we apply the mirror reflection when visualizing the pencil beams. This effectively flips the sign of the m_y coordinate such that positive m_y points towards the side of the mirror closest to the Dk stage. Figure 3.10 shows the reflected m_y direction. With *in-situ* measurements of the mirror size and the mirror mounting parameters for different mirror positions we can plot the pencil beams on the mirror.

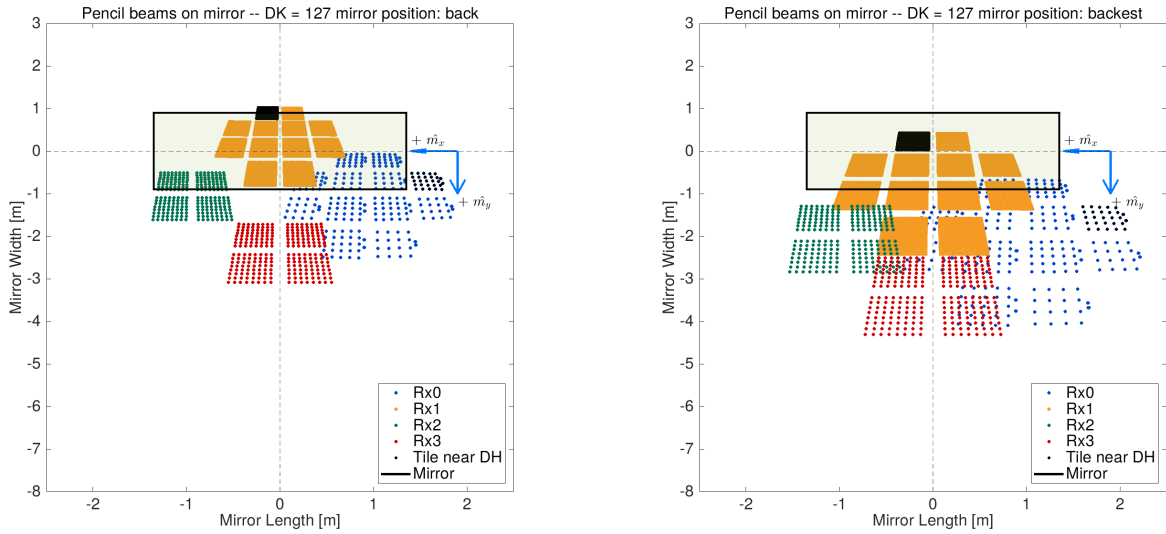


Figure 3.10. Pencil beam positions on the mirror in mirror-centered coordinates for all BICEP Array receivers at $Dk = 127^\circ$. These pointings correspond to the 2024 instrument configuration prior to the deployment of the BA220/270 receiver for visualization purposes. Left: ‘Back’ mirror position Right: ‘Backest’ mirror position.

3.5 Analysis Pipeline

This section documents the analysis of far-field beam mapping data with various calibration sources. The goal is to produce high fidelity calibration data products, (per-detector or per-array) that play an integral role in:

- Instrumental systematics analyses like $T \rightarrow P$ leakage due to beam mismatch, for which we use the thermal chopper source.
- Secondary systematics analyses like polarization-to-polarization ($P \rightarrow P$) leakage due to detector cross-polar response (to be covered in detail in Chapter 5), that use the BSNS source.
- Optical characterization, and instrument performance validation for which we also use the BSNS source in various observation configurations.

The initial analysis steps of far-field calibration measurements are similar. Hence, in this section we lay the foundation that is shared among them. In the next chapter (4) we describe how the various analyses diverge to produce final data products that achieve their respective

goals. In the limit of infinite time, one could unify the analysis of all calibration measurements into a single *calibration pipeline* that directly talks to our CMB analysis pipeline. This would eliminate the need for intermediate data products and reduce potential inconsistencies on data handling and input parameter choices. However, such project is not currently in the works given the limited number of human resources. At high level, the analysis steps of a large part of our calibration data is summarized as follows:

- Demodulation and chop reference cleaning
- Coordinate timestreams calculation
- Elliptical Gaussian beam fitting and pair-centroid centering
- Masking and binning into per-schedule maps (component maps)
- Automatic quality cuts based on fitted parameters

The analysis pipeline has been improved in the last five years to accommodate the ever-increasing number of detectors in BICEP Array, which (as of the writing of this thesis) has $\simeq 10,000$ optically active detectors. Namely, improvements in throughput, wall-clock runtime, automation, robustness, and computational efficiency have been prioritized to handle the large data volume. Aside from these optimizations and development of better visualization tools, the FFBM pipeline’s core functionality has seen little change from the description in [51]. This maintains consistency in beam characterization across nearly a decade of BICEP/Keck calibration measurements across multiple BK experiments. In this section we describe the details the above analysis steps, listing the various improvements along the way.

3.5.1 Timestream Demodulation

To isolate the source signal from surrounding microwave emitters we employ a widely used method in radiofrequency applications of modulating the source amplitude at a known frequency. The process of extracting the signal that is synchronous with the source modulation

3. Calibration Measurements Overview and Analysis

rate is called demodulation. By design our sources are *chopped* (blinking rapidly on and off) at a specified frequency. The main concern of contamination comes from the structures in the vicinity of our sources (mast, SPT/DSL structure, etc.). In practice, we record the chop frequency of the source parallel to the detector timestreams and use it to reconstruct the source signal. The sampling rate of the detectors ($\sim 150 - 200$ Hz) is much higher than the chop frequency (~ 16 Hz), allowing for accurate reconstruction of the chopped signal. The process starts with the recorded square wave chop reference signal $Q(t)$ with $\omega_c = 2\pi f_{chop}$, where f_{chop} is the chop frequency in Hz:

$$Q(t) = \frac{2}{i\pi} \sum_{k=1,3,5,\dots}^{\infty} \frac{1}{k} (e^{ik\omega_c t} - e^{-ik\omega_c t}) \quad (3.27)$$

The chopped signal $d_c(t)$ seen by a detector is the chop reference $Q(t)$ multiplied by the source signal $d_{sig}(t)$, plus noise $n(t)$:

$$d_c(t) = d_{sig}(t)Q(t) + n(t) \quad (3.28)$$

By defining the in-phase $C(t)$ and out-of-phase $S(t)$ signal components with respect to the chop reference square wave, we can write:

$$\begin{aligned} C(t) &= d_c(t) \cos(\omega_c t) \\ S(t) &= d_c(t) \sin\left(\omega_c t + \frac{\pi}{2}\right) \end{aligned} \quad (3.29)$$

In the case that the chop reference is perfectly aligned with the source signal, the in-phase component $C(t)$ contains the source signal while the out-of-phase component $S(t)$ contains only noise. The phase between the two is $\theta = \tan^{-1}[S(t)/C(t)]$. In the case of $\theta \approx 0$ $d_{sig} = C(t)$ and $S(t) = n(t)$ so the in-phase component is maximized. Figure 3.11 illustrates the demodulation process for two BICEP Array detectors at different frequencies. At the beginning of each calibration season we manually align the chop reference phase θ by inspecting the demodulated timestreams of a few detectors across the focal plane, and find the optimal

phase offset that maximizes $C(t)$ and minimizes $S(t)$. This phase offset is then applied to all detectors in the array for the entire season.

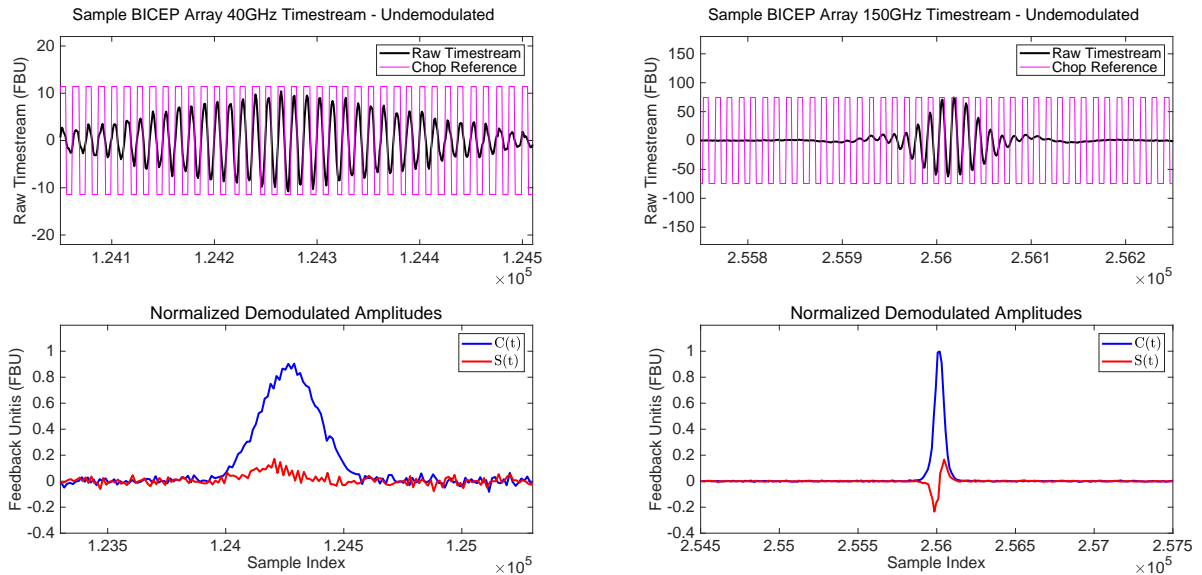


Figure 3.11. Sample beam map timestreams before and after demodulation for BICEP Array 40GHz detector (left) and 150GHz detector (right). For each, the top panel shows the raw timestream (black) with the chop reference (magenta) overlaid. The bottom panel shows the in-phase (blue) and out-of-phase (red) demodulated timestreams. The units are raw feedback units (FBU). The x-axis limits are picked to include the same number of chop cycles to illustrate the difference in beam size between the two detectors.

Albeit a straightforward process, demodulation in practice opens a window for various analysis-related systematics that might affect the reconstruction of the beam shape. Any jitter or imperfection during readout of the chop reference signal can introduce amplitude suppression, or distortion of beam shape. We require near sample-level precision when calculating the phase difference between $C(t)$ and $S(t)$, along with phase stability throughout the entire observation. To achieve this, we clean the chop reference signal by constructing an ideal, high-resolution square wave at the chop frequency (extracted from the Fourier transform of the recorded chop reference). This process has been shown to mitigate demodulation systematics related to phase jitter and drift, non-instantaneous chop transitions, and phase gradients due to the finite swing distance of the chopper blade across the source [51]. The final step in chop reference cleaning is to apply a low pass filter to the ideal square wave to

3. Calibration Measurements Overview and Analysis

remove high frequency noise that can leak into the demodulated timestreams, with negligible penalty to the noise at lower frequencies.

Outside of improvements in the algorithmic implementation and logic of the chop reference cleaning process, we implement chop reference repairs unique to BICEP Array data. Local network delays between the housekeeping electronics and GCP sometimes result in repeated chop reference samples. Such behavior introduces spurious phase anomalies between the chop reference and the detector timestreams, simultaneously seen by all detectors in the array. An example is shown in Figure 3.12.




Figure 3.12. Repeated samples in square wave chop reference resulting in phase differences between in-phase and out-of-phase components.

When such anomalies occur while a detector is scanning across the source, the demodulated signal is biased low, leading to systematic suppression of the beam peak. To identify repeated samples we use run-length encoding of the *source-on* and *source-off* states. States with more than 5 times the encoded length trigger the repair algorithm. We repair these repeated samples by extracting the chop frequency prior to the outage (via FFT) and inserting a reconstructed square wave reference at that frequency. When the chop reference recovers the inserted square wave naturally matches the measured reference, since the chances of a phase drift during the second-long network delays are slim. Such operation is not applied to BICEP3 since the readout scheme of the chop reference is through the mount's control computer that runs on the same clock as the control program, reducing the possibility of

network delays.

Finally, we calculate the coordinates of our chopped and demodulated timestreams via the pointing model described in section 3.4. The output data product from this step is referred to as the *demodulated timestream*.

3.5.2 Beam Parameterization and Pair-Centroid Centering

Although the pointing model for FFBM is calibrated to good precision, there are a number of uncertainties that enter through the measurement of its parameters (mount, mirror, source, etc.). If we want to produce high fidelity beam maps over multiple calibration schedules, we need to ensure that the beam maps are precisely centered in detector-centered coordinates. We achieve this by centering each detector's demodulated timestream to a common location, the pair centroid, defined as the point exactly in between the two detectors in a pair. A 2D elliptical Gaussian is employed to extract the fitted centroid center for each demodulated timestream in (x', y') . The mathematical description of fitting for the beam shape closely follows [73] and reflects the implementation in the FFBM pipeline.

For each calibration schedule we start by defining the 2D elliptical Gaussian model:

$$B(\mathbf{x}) = \frac{1}{\Omega} \exp \left[\frac{-(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})}{2} \right] + b \quad (3.30)$$

where \mathbf{x} is a 2D position vector (x', y') , $\boldsymbol{\mu}$ is the beam center (x_0, y_0) , Ω is the normalization constant, b is a background offset, and $\boldsymbol{\Sigma}$ is the covariance matrix defined as:

$$\boldsymbol{\Sigma} = \mathbf{R}^{-1} \mathbf{C} \mathbf{R} \quad (3.31)$$

where \mathbf{R} is a 2x2 rotation matrix parameterized by the angle γ of the major axis of the elliptical Gaussian with respect to x' counter-clockwise. They are parameterized as:

$$\mathbf{R} = \begin{pmatrix} \cos \gamma & -\sin \gamma \\ \sin \gamma & \cos \gamma \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} \sigma_{\text{maj}}^2 & 0 \\ 0 & \sigma_{\text{min}}^2 \end{pmatrix} \quad (3.32)$$

The 7-parameter fit consists of $[x_0, y_0, \sigma_{\text{min}}, \sigma_{\text{maj}}, \gamma, \Omega, b]$, which fully characterizes the beam.

We adopt an equivalent parameterization in terms of the ellipticity components p, c (for ellipticity plus and cross) and beam width σ transforming as:

$$e = \left(\frac{\sigma_{\text{maj}}^2 - \sigma_{\text{min}}^2}{\sigma_{\text{maj}}^2 + \sigma_{\text{min}}^2} \right), \quad (3.33)$$

$$p = e \cos 2\gamma, \quad (3.34)$$

$$c = e \sin 2\gamma, \quad (3.35)$$

$$\sigma^2 = \left(\frac{\sigma_{\text{maj}}^2 + \sigma_{\text{min}}^2}{2} \right). \quad (3.36)$$

which is more relevant to beam-related systematics analyses [55]. In particular, differential beam parameters between detectors within a pair are defined as detector V minus detector H and are of particular interest as they directly probe beam-related systematics. For example, for differential beamwidth this means: $d\sigma = \sigma_V - \sigma_H$.

Fitting for these parameters prior to binning avoids biasing due to pixelization. To ensure robustness of the fits we employ techniques including: deglitching, and applying amplitude + noise level cut thresholds. These are empirically derived based on the expected amplitude and noise of the data. A particularly sensitive choice when fitting is the initial guess for the parameters needed for the Nelder-Meads minimization procedure. The values for the elliptical Gaussian's widths is taken from the instrument design specifications (i.e. the beamwidth of BICEP Array 220/270GHz is roughly 0.06°) while the amplitude is empirically estimated by looking at the peaks of various demodulated timestreams across the focal plane. The fitted parameters serve as a powerful diagnostic used to inform our systematics mitigation strategies and identify trends across the focal plane. (more details in 3.5.5). The

results of these parameters for BICEP3 and BICEP Array are presented in section 4.1. At this point, the demodulated timestreams are shifted such that the point exactly between the fit centers (pair centroid) is at the origin of the detector-centered coordinate system.

3.5.3 Binning and Masking

With unbiased estimates of the per-detector beam parameters we proceed to bin the demodulated, pair-centered timestreams into maps. We refer to these as ‘component maps’ hinting that the end goal is to coadd multiple component maps into ‘composite maps’³. The pipeline supports binning in various coordinate systems including: apparent azimuth and elevation (Az_{ap}, El_{ap}) , and focal plane centered coordinates (X, Y) , which are used for instrument performance validation and exploratory analysis of beam maps. Nominally, we bin in detector-centered coordinates (x', y') that allow us to probe instrument-fixed detector characteristics.

The choice of bin size varies across different calibration measurements. The nominal bin size for most FFBM is 0.01° square pixels. Given the observing strategy described in section 3.2, we hit each pixel at least once per scan across the source for both BICEP3 and BICEP Array schedules. This ensures uniform coverage across the beam map and minimizes pixelization artifacts. However, receivers like BA220/270 with beamwidth of roughly 0.06° would require a smaller bin size and finer steps in elevation to meet the existing beam quality standards. As illustrated in the top right of Figure 3.14, we apply a mask to the binned maps to remove surrounding microwave emitters that have escaped demodulation. Such response is typically $\sim 25dB$ below the main beam peak and contaminates the beam map in the direction of the ground. These masks vary with instrument orientation and are calculated using the apparent azimuth and elevation coordinates calculated in section 3.4. For BICEP Array an additional mirror mask is applied to exclude detectors with beams that miss or partially miss the mirror. In principle, we could use the mirror-centered coordinates to identify such detectors. However,

³The process of compositing varies depending on the type of calibration measurement and will be discussed in detail in Chapter 4.

3. Calibration Measurements Overview and Analysis

this has not been implemented yet. Instead, the mirror mask is calculated by projecting the mirror onto the focal plane at different Dk angles as shown in Figure 3.13. The mirror height varies as a function of distance from the receiver (due to the 45° tilt). Hence, the mirror projection is trapezoidal, effectively covering more detectors as the beam size grows with mirror height.




Figure 3.13. Mirror projection onto the BA150 focal plane at three different Dk angles for a single mirror position in focal plane-centered coordinates (\hat{X}, \hat{Y}) . From left to right $Dk = 105, 127, 150$. The trapezoidal projection is due to the increasing beam size with mirror height, effectively ‘catching’ more detector beams.

3.5.4 Pipeline Performance

Important part of calibration measurement success is data quality reporting. Transferring data from the South Pole via satellite, results in a ~ 1 day (on average) delay between data acquisition and analysis⁴. For this reason, having a fast turnaround from raw data to beam maps is essential in assessing measurement quality and adjusting the measurement parameters in the case of unexpected problems. In this section we describe two ways used to accelerate data quality reporting, and overall pipeline performance. The aim is to accommodate the ever-increasing number of deployed detectors in BICEP Array, and meet the strict timing constraints of the Austral summer calibration season.

Pipeline Parallelization: the wall-clock time of reducing raw timestreams to beam maps used to scale with detector number. A single 8-hour BICEP3-style schedule nominally takes ~ 8 hrs to go from raw arcfiles to beam maps while serially loading all arcfiles in the array

⁴This assumes no data transfer delays and just reflects the partial SPTR coverage. Data transfer delays for BA have reached as long as 3 weeks during the 2024 season.

onto memory. Peak memory requirements would reach $\sim 67\text{GB}$ during the fitting step of the analysis. This includes analysis steps up to section 3.5.3 along with per-detector diagnostic plots similar to Figure 3.14. BICEP Array schedules are longer (typically closer to 9hrs). In 2024 BICEP Array had nearly 10,000 operational detectors across four receivers which is more than a factor of four increase in detector count compared to BICEP3. The clock time and memory usage would scale accordingly, resulting in unpractical runtimes and memory requirements. To address this issue, in 2022 the FFBM pipeline was restructured to split the per-schedule analysis across multiple compute nodes in the Cannon high-performance computing (HPC) cluster. The split is based on detectors read by a single MCE (as opposed to per-receiver or per-detector), since the number of detectors per MCE is roughly the same across experiments. This choice results in:

- Marginal increase in CPU-cost due to parallelization. Even though the number of CPU cores requested scales linearly with the number of MCEs, the per-core CPU time and memory requested is substantially reduced, and can be safely capped to accommodate the largest MCEs in the array.
- A large reduction in wall-clock time. The speedup is slightly larger than the number of MCEs in the array. For example, BICEP3 has 4 MCEs and the wall-clock time is reduced by a factor of ~ 5 .
- A reduction in peak memory requirement proportional to the number of MCEs in the array.

With this pipeline architecture we achieve data quality reports in less than 1 day (modulo HPC cluster availability), dominated by data transfer delays. The per-schedule wall-clock time is reduced to just $\sim 1.3\text{hrs}$ and peak memory is capped to $< 20\text{ GB}$ per CPU core for both BICEP3 and BICEP Array schedules. The main advantage of this architecture is its scalability since the achieved pipeline performance is now decoupled from detector count. Currently, the bottleneck is serially loading all the raw data from disk (loaded from multiple compute nodes simultaneously) for the demodulation step of the analysis. However, splitting

the analysis further on a per-arcfile basis (roughly 5 minutes of data) is not considered practical, since building full demodulated timestreams from multiple arcfiles analysed separately would be a bookkeeping nightmare.

On-site demodulation: FFBM analysis depends on the Cannon HPC cluster for parallelization. However, a serial version of the pipeline is still maintained, optimized to execute the demodulation step on-site with the memory and computing resources available at the Pole. The demodulated timestreams (which represent a factor of 10 reduction in data volume from raw data), would then be transferred to North America for the remaining analysis. However, this mode of operation has not been tested on new calibration data.

3.5.5 Automatic Cuts

So far we have described per-detector, per-schedule mapmaking. At this point the focus shifts from individual schedules to analysis across the entire calibration measurement campaign. Recall that the goal of a FFBM measurement is to provide the highest fidelity measurement of the detector beam in the regions of interest. For example when using the thermal chopper we are primarily interested in mapping the main beam and near sidelobes, while when using the bright BSNS source we are interested in mapping the near/mid sidelobe regions out to roughly 16° . During FFBM campaigns, the total number of calibration schedules for a given instrument varies from 10 all the way to 70 depending on the goal of the calibration measurement. For instance, when beam mapping with the thermal chopper we typically take 40 to 70 schedules per season. After all measurements are completed we have a large ensemble of them that can be used for statistics and to identify poorly behaved detectors. This section briefly describes the automatic cut selection criteria. For a more detailed description refer to [51]. The automatic cuts consist of the following criteria:

- Mirror Illumination Cut: Detectors that fully or partially miss the mirror are cut from the analysis, which avoids spillover systematics.
- Beam parameters must be within specified thresholds that are based on our extensive

library of beam mapping measurements, informed by the detector design and optical simulations.

- Both detectors within a pair must be responsive and pass the above cuts. In practice, we keep single detectors for visualization purposes only.
- Responsivity of a detector must be consistent throughout the schedule duration with its noise levels remaining stable.
- Manual ‘hand cuts’ are implemented post visual inspection of individual beam maps. Detectors that somehow pass the above criteria are cut here.

The resulting cuts are kept in a separate file and applied to all schedules at once during analysis of the entire season’s data. An example of the data products coming out of the FFBM pipeline analysis steps using the thermal chopper source is shown in Figure 3.14.

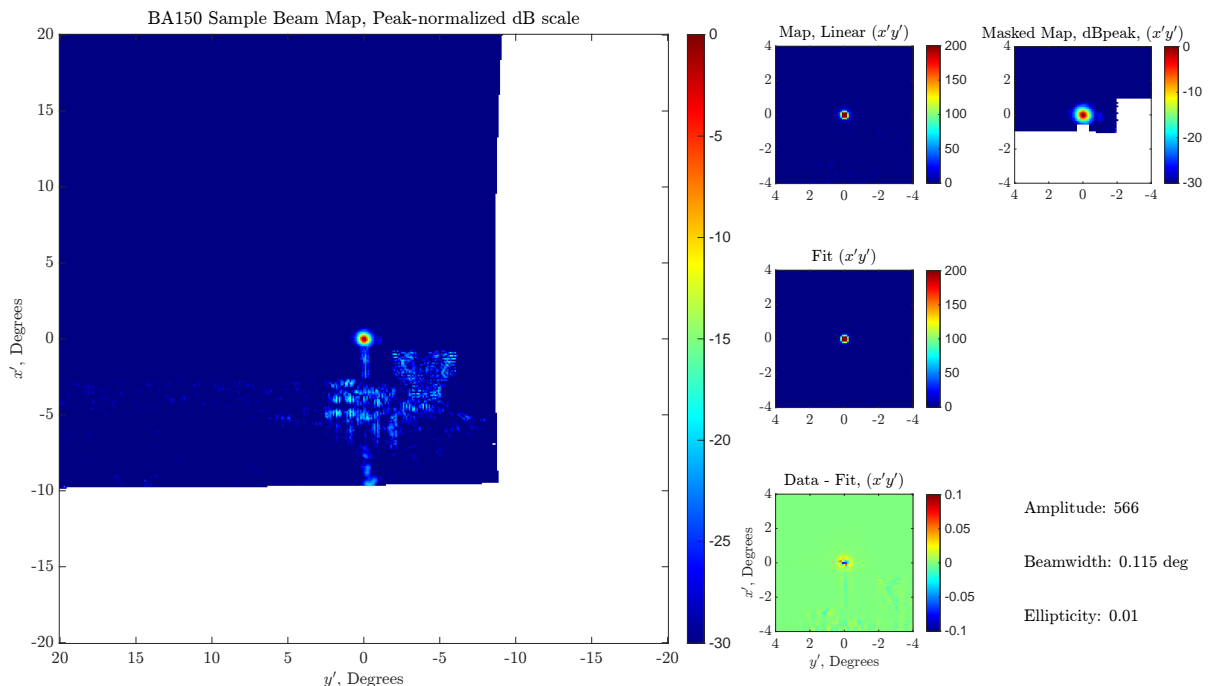


Figure 3.14. BICEP Array 150GHz FFBM beam map using the thermal chopper. The various plot windows illustrate the results of demodulation, pair-centroid centering, fitting, and masking.

The remaining analysis steps of the various far-field measurements vary depending on the source used (thermal chopper vs BSNS), the calibration configuration (with or without mirror, scan strategy), and the analysis goal (beam window functions, optical characterization,

3. Calibration Measurements Overview and Analysis

systematics analysis, instrument performance validation etc.). We defer these descriptions to the appropriate sections in Chapter 4.

Chapter 4

Beam Characterization Results

Control of instrument systematics is critical for placing constraints on cosmological parameters. To this end, we perform a variety of beam mapping measurements, each specifically designed to probe different aspects of instrument performance. This chapter focuses on three such measurements:

- **Unpolarized Far Field Beam Mapping** (section 4.1): Measurements of the main beam with the thermal chopper source to produce high signal-to-noise, per-detector composite beam maps, used to make array-averaged beam window functions B_ℓ 's. B_ℓ 's are used in CMB analysis to smooth external sky maps for our simulations framework, and to evaluate the beam suppression correction applied to the measured CMB power spectra. Most importantly, data products from these measurements are used to directly estimate the level of temperature-to-polarization leakage in real data.
- **Polarized Far Field Beam Mapping** (section 4.2): Measurements with the BSNS to characterize the per-detector polarized beam response, that inform our systematics studies. Primarily used for detector polarization angle calibration, and to estimate polarization-to-polarization ($P \rightarrow P$) leakage—a subdominant systematic in real data, due to cross-polar beam mismatch within detector pairs (described in Chapter 5).

- **Far Sidelobe Mapping**¹ (section 4.3): Measurements of the optical response at large angles away from the detector pointing ray, to probe sidelobe pickup and validate forebaffle loading performance.

The calibration measurements described in this chapter along with their observation setup are summarized in Table 4.1. Results from each of these measurements are presented within their respective sections.

Name	Calibration source	Location	Mirror	Region mapped	Polarized beams
Unpolarized FFBM	24' thermal chopper	Mast on opposite building	Yes	Main + near sidelobes	No
Polarized FFBM	BSNS/RPS	Mast on opposite building	Yes	Main + near/mid sidelobes	Yes
Far Sidelobe Mapping (FSL)	BSNS	Mast on same building	No	Far sidelobes	No

Table 4.1. Summary of calibration measurements and their observation/setup details. The unpolarized and polarized FFBM setup is the same except for the source used. Note that for Far Sidelobe Mapping, we observe the source on the mast located on the *same building* with the receiver, effectively bringing it closer to the near field of the instrument.

4.1 Unpolarized Far Field Beam Mapping

The only FFBM measurement that directly interfaces with our CMB analysis uses the thermal chopper source to map the unpolarized beam response of each detector. Data products from these measurements (beam parameters, composite beam maps, B_ℓ 's) are used in our simulations framework and to inform our systematics studies. Moreover, unpolarized FFBM is a very powerful way to validate the optics performance of the instrument and identify any unwanted optical response coupling to our detectors (see for example appendix B). For these reasons, we devote the majority of the available calibration time to these measurements with the goal to map each detector out to a radius of $> 2^\circ$. We take anywhere from 40 to 70 schedules worth of data per instrument, at 16 different Dk angles (and 5 mirror positions

¹Also known as ‘amplified sidelobe mapping’—we will not be using this name in this thesis

in the case of BICEP Array) to achieve this goal. This section describes the analysis steps following those described in section 3.5 that apply exclusively to FFBM measurements with the thermal chopper.

4.1.1 Composite Maps and Beam Window Functions

Our simulations framework discussed in section 2.7 relies on accurate knowledge of the instrument's beam to reconstruct the observed sky from an external input map. For this reason, we construct the highest signal-to-noise beam maps possible for each BICEP3 and BICEP Array detector by coadding each map across all schedules. This process results in per-detector *composite maps*. In principle, coadding would simply be averaging. In practice, we take the median across all component maps to reject outlier measurements and avoid the scenario of amplifying the noise due to a few unusually noisy component maps that pass automatic cuts. An example of coadding over component maps to produce a composite map for a single BA150 detector is shown in Figure 4.1.

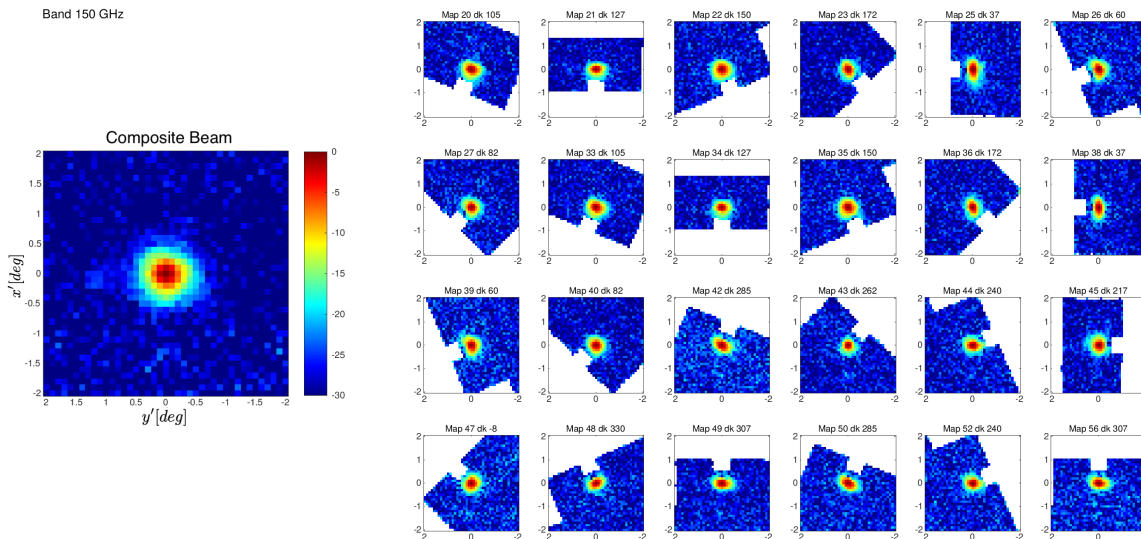


Figure 4.1. BA150 sample composite map (left) from a single detector, produced by coadding 24 component maps (right) taken at different Dk angles and different mirror positions.

The resulting composite maps are then used to calculate the array-averaged beam window functions B_ℓ 's for each observing frequency band. These mathematical objects describe how

much power at angular scale ℓ is suppressed by the instrument beam. To simulate the observed sky, we convolve external sky maps with the beam window function B_ℓ in harmonic space for a given frequency band. Hence, accurate knowledge of B_ℓ is critical for unbiased observed CMB skies. To calculate them, we first array-average the per-detector composite maps and take an azimuthal average at each ℓ bin (effectively asserting that to a good approximation the instrument's beam is circularly symmetric). The Fourier transform of the azimuthally averaged beam profile is the beam window function B_ℓ . In Chapter 3, we asserted that the beam maps are a representation of the detector's point spread function assuming the thermal chopper is a point source in the far field. However, the chopper aperture has a non-negligible angular size of $\theta_s \sim 0.086^\circ$ as seen from our receivers that needs to be corrected for accurate B_ℓ estimation. By the Green's function method [22] and using the Fraunhofer approximation for the far field, the measured beam is the convolution of the true beam with the chopper aperture function. Since convolution in real space is multiplication in Fourier space, we can divide the B_ℓ by the Fourier transform of the chopper aperture function to recover the true beam window function of the instrument $B(\ell)_{true}$ as:

$$B(\ell)_{true} = \frac{\ell\theta_s B(\ell)}{2J_1(\ell\theta_s)}, \quad (4.1)$$

where J_1 is the Bessel function of the first kind. The resulting beam window functions are shown in [51]. The array-averaged maps for BICEP3 and BICEP Array are shown in Figure 4.2.

All BICEP Array and BICEP3 difference beams have substantially reduced amplitude compared to previous generation experiments [74]. However, the dominant structure is complex and not well described by a single beam mode like differential gain or pointing. More on this in section 4.1.3.




Figure 4.2. Array-averaged beams for BICEP3 (B95) and BICEP Array (BAXXX where XXX denotes different frequencies) updated from [73] to the 2024 FFBM campaign with the thermal chopper. From left to right, the columns correspond to different observing frequencies. The top and middle rows show the array-averaged V and H polarization beams, and the bottom row the pair-difference beams. The BA30/40 receivers are much noisier than B95 and BA150 with mostly noise present in the difference beam. B95 and BA150 show more complex structure further discussed in section 4.1.3.

4. Beam Characterization Results

4.1.2 Beam Parameter Results

As discussed in section 3.5.2, the beam parameters extracted from the component maps are a powerful diagnostic and validation tool. Here we present a summary of these parameters for BICEP3 and BICEP Array receivers from 2019 to 2024. Table 4.2 summarizes the beam parameters of interest with the different color codes, corresponding to different years.

Parameter	BA30	BA40	B95	BA150
$\sigma(^{\circ})$	-	-	0.166 ± 0.004 ± 0.002	-
	0.478 ± 0.018 ± 0.015	0.364 ± 0.017 ± 0.009	-	-
	-	-	0.165 ± 0.004 ± 0.002	-
	0.488 ± 0.024 ± 0.012	0.355 ± 0.013 ± 0.007	0.162 ± 0.003 ± 0.002	0.105 ± 0.006 ± 0.005
p	0.478 ± 0.022 ± 0.011	0.357 ± 0.013 ± 0.008	-	0.102 ± 0.004 ± 0.005
	-	-	-0.023 ± 0.026 ± 0.018	-
	0.027 ± 0.046 ± 0.043	0.029 ± 0.069 ± 0.028	-	-
	-	-	0.013 ± 0.025 ± 0.023	-
c	0.015 ± 0.072 ± 0.035	0.022 ± 0.053 ± 0.028	0.010 ± 0.029 ± 0.020	0.006 ± 0.054 ± 0.061
	0.009 ± 0.066 ± 0.049	0.018 ± 0.062 ± 0.040	-	-0.002 ± 0.048 ± 0.077
	-	-	-0.026 ± 0.031 ± 0.017	-
	-0.015 ± 0.035 ± 0.022	0.007 ± 0.030 ± 0.036	-	-
$d\sigma(^{\circ})$	-	-	-0.020 ± 0.023 ± 0.018	-
	-0.003 ± 0.019 ± 0.036	-0.010 ± 0.029 ± 0.029	-0.009 ± 0.026 ± 0.024	0.006 ± 0.058 ± 0.057
	0.008 ± 0.016 ± 0.036	0.004 ± 0.027 ± 0.034	-	0.024 ± 0.039 ± 0.069
	-	-	2.4e-04 ± 0.001 ± 0.000	-
dp	0.003 ± 0.010 ± 0.008	1.3e-4 ± 0.008 ± 0.003	-	-
	-	-	-4.8e-04 ± 0.001 ± 0.000	-
	0.003 ± 0.007 ± 0.005	-1.9e-4 ± 0.007 ± 0.002	-5.0e-04 ± 0.001 ± 0.000	-5.9e-04 ± 0.004 ± 0.001
	1.6e-03 ± 0.014 ± 0.006	-8.5e-04 ± 0.009 ± 0.002	-	4.8e-04 ± 0.002 ± 0.001
dc	-	-	0.004 ± 0.012 ± 0.002	-
	0.014 ± 0.018 ± 0.019	0.014 ± 0.018 ± 0.019	-	-
	-	-	-0.003 ± 0.012 ± 0.003	-
	0.001 ± 0.027 ± 0.020	-0.001 ± 0.031 ± 0.009	-0.004 ± 0.013 ± 0.002	-0.017 ± 0.019 ± 0.007
dx'	0.006 ± 0.034 ± 0.021	0.001 ± 0.021 ± 0.013	-	0.011 ± 0.017 ± 0.011
	-	-	-0.001 ± 0.004 ± 0.002	-
	0.006 ± 0.013 ± 0.013	0.001 ± 0.015 ± 0.010	-	-
	-	-	-0.002 ± 0.004 ± 0.002	-
dy'	-0.001 ± 0.017 ± 0.014	0.001 ± 0.010 ± 0.009	-0.002 ± 0.004 ± 0.004	0.000 ± 0.018 ± 0.006
	0.009 ± 0.012 ± 0.018	0.004 ± 0.014 ± 0.011	-	-0.005 ± 0.011 ± 0.009
	-	-	-0.05 ± 0.14 ± 0.05	-
	0.71 ± 1.55 ± 0.52	-0.51 ± 1.91 ± 0.35	-	-
dx'	-	-	-0.06 ± 0.15 ± 0.04	-
	-	-	-0.09 ± 0.14 ± 0.06	0.00 ± 0.33 ± 0.08
	0.58 ± 2.67 ± 0.60	0.02 ± 1.40 ± 0.45	-	0.01 ± 0.15 ± 0.12
	1.55 ± 2.10 ± 0.86	-0.21 ± 1.51 ± 0.60	-	-
dy'	-	-	0.02 ± 0.16 ± 0.06	-
	-0.21 ± 1.37 ± 0.39	-0.07 ± 0.57 ± 0.32	-	-
	-	-	-0.01 ± 0.14 ± 0.06	-
	0.19 ± 0.78 ± 0.57	-0.43 ± 0.81 ± 0.44	0.01 ± 0.15 ± 0.06	0.02 ± 0.37 ± 0.07
			-0.03 ± 0.10 ± 0.12	

Table 4.2. Beam parameter summary table for BA30/40 and B95 for years 2019, 2020, 2021, 2023, and 2024. Similar to the beam parameter table in [73], but extended to include BA150. The values are listed as FPU median ± FPU scatter ± measurement uncertainty. The variation of each receiver from year to year is smaller than the measurement uncertainty, except from the beamwidth parameter of the B95 receiver. This coherent shift of $\sim 1.2\%$ is due to the thin window change on December 2022 [75]. All beamwidth parameters have been corrected for the non-negligible size of the chopped source aperture in the far field.

Parameters are presented as FPU median ± FPU scatter ± measurement uncertainty. All detectors that pass cuts and are included in CMB analysis are used to calculate the FPU

median and scatter. The beam widths listed have been corrected for the finite chopper aperture. All differential parameters are calculated as V-H. The uncertainty in each detector’s beam measurement is the standard-deviation in fit parameters across all component maps used to make the composite map. The quantity listed in the table as *measurement uncertainty* is the average of the per-detector uncertainties. The fitted gain parameter is only used to normalize the component maps before coadding, and is not included in the table. Calibration measurements are taken with aluminum TES transition (as opposed to the more sensitive titanium transition) to avoid saturation that could potentially prevent us from fully resolving the beam peak. Hence, gain parameters are not representative of the detector responsivities during CMB observations.

The table shows consistency across multiple years for BA30/40, B95, and BA150 receivers for all parameters within the measurement uncertainty, noting that the BA30/40 detectors are noticeably noisier than B95 and BA150 detectors. There is a 1.2% coherent decrease in beamwidth for BICEP3 between 2019 and 2023, an expected outcome given the December 2022 window replacement. The old slab-type window acted like a meniscus lens on the beam when the receiver was under vacuum. The new window has little impact on BICEP3’s focus resulting in the measured beamwidth decrease [75]. This example illustrates the level of precision we achieve, being able to detect how small changes in optical configuration of the instrument result in a measurable difference in beam parameters. BA150 only had 5 modules deployed in 2023, and 10 modules in 2024 out of 12 total. There is a less noticeable change in BA150’s beam width between 2023 and 2024 as the receiver was slightly out of focus during its first year of observations in 2023. Prior to the 2024 season, the BA150 was re-focused, resulting in a small beam width decrease of less than 1%. However, given the measurement uncertainty this is not a statistically significant change and could be attributed to real characteristics of the extra 5 modules installed in 2024. An upcoming beam mapping campaign in February 2026 after all 12 modules are installed will provide a more complete picture of BA150’s beam parameters. To illustrate how beam parameters can be used to establish automatic cut thresholds we plot focal plane views of the beam ellipticity for BA150 detectors from the 2024 FFBM campaign, as shown in Figure 4.3. Detectors with unusually high

4. Beam Characterization Results

ellipticity will be cut from compositing and those yielding repeatably bad beam parameters are flagged and not included in CMB analysis.




Figure 4.3. BICEP Array 150GHz focal plane view of total ellipticity in its 2024 observing configuration, prior to any automatic cuts 3.5.5. Polarization V detectors are shown in blue and polarization H detectors are shown in red. The total ellipticity is exaggerated by a factor of 20 for visualization purposes. Detectors with unusually high ellipticity are automatically cut from compositing by using an empirical threshold unique to each receiver. Visualizations like this are useful to identify patterns across the focal plane that might reveal trends in beam parameters. The missing detectors from this visualization correspond to either uninstalled modules or detectors that have been turned off due to excessive heating of the focal plane when biased.

All beam parameter results are used for end-to-end instrument performance validation. However, measurements of differential ellipticity are particularly important as they are directly used to subtract differential ellipticity from real CMB timestreams. Since deprojection (just like any filtering process) removes some real sky signal, we choose to subtract differential ellipticity instead of deprojecting it to avoid removing real TE correlations in our measurements.

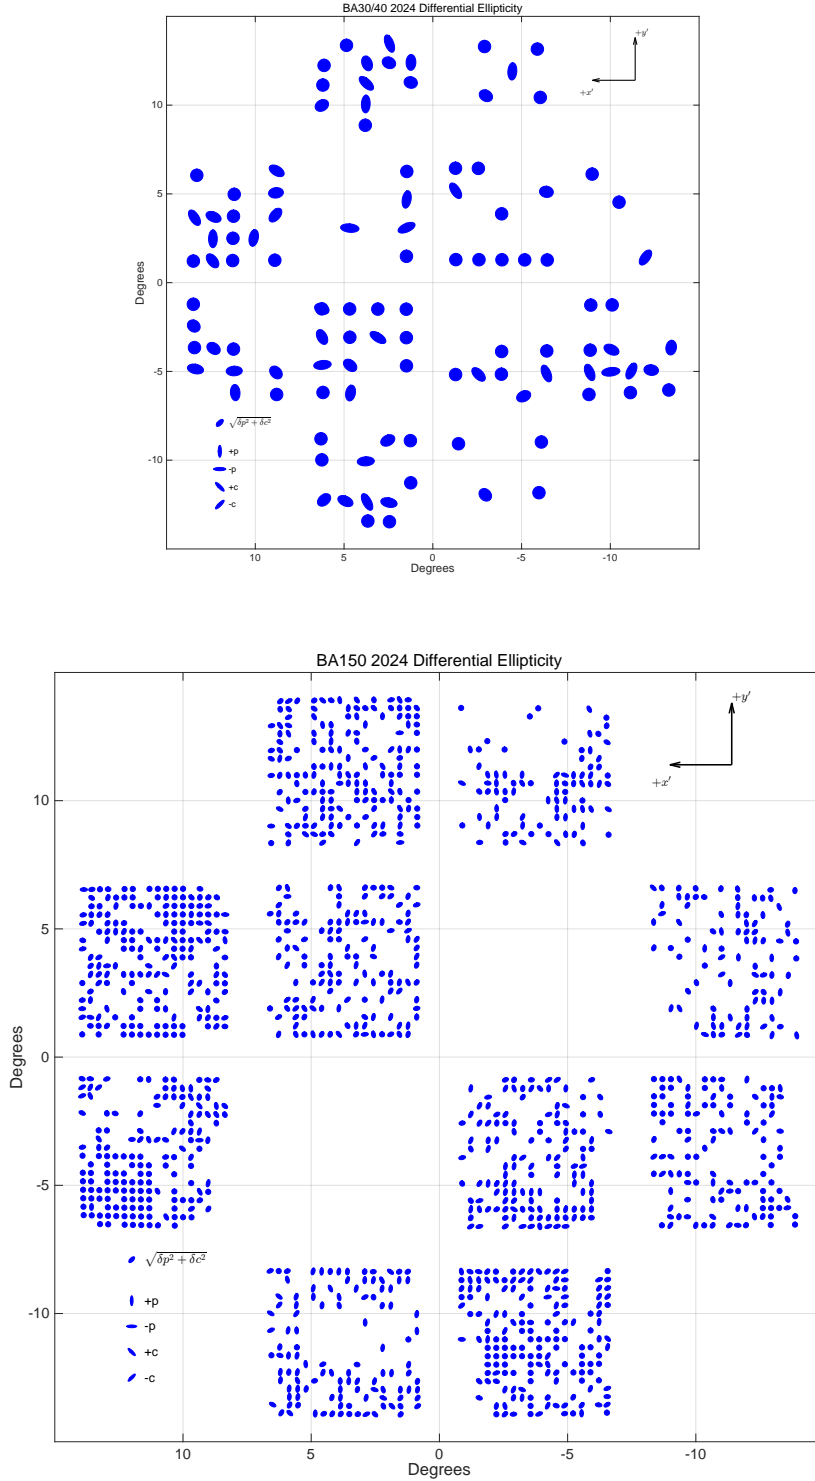


Figure 4.4. Differential ellipticity for BA30/40 (top) and BA150 (bottom) from the 2024 FFBM campaign, exaggerated by a factor of 75 for visualization purposes. The differences are calculated as V-H and are only plotted if both detectors pass beam parameter cuts for the entire season of FFBM measurements. These values are directly used to scale a pre-constructed differential ellipticity template (from the Planck temperature map and its derivatives) before removing it from real CMB timestreams to avoid removing real TE correlations in the sky signal.

Therefore, instead of using the deprojection coefficients (calculated by regressing a pre-constructed differential ellipticity template from the Planck temperature map with real data timestreams), we scale the template by the measured coefficient from unpolarized FFBM measurements before removing it from the data. Inherent beam ellipticity of a detector is likely due to the geometry of the slot antennas themselves. In the reverse time sense, each antenna can be thought of as the sum of many point sources oriented along the antenna direction. An asymmetry in the overall slot antenna geometry would result in residual ellipticity in the V/H axes directions. Within a detector pair, this would contribute to inherent differential ellipticity. We expect and measure it to be small, since the total number of antennas is large and are arranged in square tiles. In addition, a phase or amplitude gradient in the beam as it passes through the optical chain (that is associated with inherent beam ellipticity), would also contribute to a small amount of differential ellipticity in the far field. Referring back to BICEP3’s beam parameters for differential ellipticity plus, (dp in table 3.2), there is a measurable shift from an overall positive to an overall negative value between 2019 and 2023. The number of detectors that contributed to the FPU median is approximately the same between the two years, indicating that this might be a systematic shift in the differential ellipticity. The details are currently being investigated.

4.1.3 Residual Beams and Temperature-to-Polarization Leakage Simulations

As discussed in section 2.7.3, we use deprojection to remove the leading order beam mismatch modes that lead to $T \rightarrow P$ leakage from our timestreams. Namely, we deproject leakage due to differential gain, and pointing, while subtracting differential ellipticity.

Residual Beams

Any higher order structure in the pair difference beams not removed by this process introduces $T \rightarrow P$ leakage that could bias r constraints. We refer to this structure as the *undeprojected*

residuals or simply *residual beams* (the structure left over post deprojection and differential ellipticity subtraction)².

Since the BA30/40 difference beams are mostly dominated by noise (which can be seen in Figure 4.2), the discussion of residual beams is focused on BICEP3 and BA150. Figure 4.5 shows the array averaged $V - H$ residual beams for BICEP3 and BA150 after deprojecting differential gain and pointing. Beam space deprojection has been used here to help illustrate the residual structure.




Figure 4.5. BICEP3 (left column) and BA150 (right column) array averaged pair-difference (top) and residual beams (bottom) after deprojecting differential gain, pointing, and ellipticity. The residual beams show complex structure that is not well described by the leading order beam mismatch modes described in section 2.7.3. If ignored, residual beams directly source $T \rightarrow P$ leakage that could bias constraints on the tensor-to-scalar ratio r . The colorscale shrinks by a factor of ~ 3 between top and bottom, illustrating the effectiveness of deprojection.

The residual beam for both receivers contains complex structure that is not well described by the leading order beam mismatch modes. Notably, the BICEP3 receiver coadded residual

²I always disliked the term *undeprojected residuals*. It will not be used in this thesis. We use the term ‘residual beams’ instead, even though most BK publications use the former term.

beam shows little change from previous measurements (Figure 5.13 in [51]), indicating stability in the optical performance of the instrument over time. Residual beams are not used to directly quantify $T \rightarrow P$ leakage. Rather, we employ a timestream simulation approach that mimics our signal simulations to directly estimate the impact of the residual beams on the total $T \rightarrow P$ leakage in our final maps. We describe these simulations in the next section.

Temperature-to-Polarization Leakage Simulations

We have developed timestream-level beam beam simulations (similar in principle to those described in section 2.7.5) to directly estimate the level of $T \rightarrow P$ leakage due to residual beams in our data (internally referred to as ‘beam sims’). These simulations differ from the standard CMB sims in that they start with a temperature-only unsmoothed sky from Planck ³ with no input polarization (input $Q/U = 0$). In general, our telescopes scanning across the BK sky patch measure the convolution of the instrument beam with the CMB sky:

$$\begin{bmatrix} T_{\text{obs}} \\ Q_{\text{obs}} \\ U_{\text{obs}} \end{bmatrix} = \begin{bmatrix} b_{TT} & b_{TQ} & b_{TU} \\ b_{QT} & b_{QQ} & b_{QU} \\ b_{UT} & b_{UQ} & b_{UU} \end{bmatrix} * \begin{bmatrix} T_{\text{sky}} \\ Q_{\text{sky}} \\ U_{\text{sky}} \end{bmatrix} \quad (4.2)$$

where b_{XY} are the convolution kernels that describe how the instrument beam couples the sky signal in Stokes parameter Y into the observed Stokes parameter X . For an ideal instrument, the only non-zero kernels would be along the diagonal, i.e. b_{TT} , b_{QQ} , and b_{UU} . However, the presence of beam mismatch leads to non-zero off-diagonal kernels, such as b_{QT} and b_{UT} , which describe $T \rightarrow P$ leakage. Setting $Q_{\text{sky}} = U_{\text{sky}} = 0$, the resulting measured signal is given by:

$$\begin{bmatrix} T_{\text{obs}} \\ Q_{\text{obs}} \\ U_{\text{obs}} \end{bmatrix} = \begin{bmatrix} b_{TT} * T_{\text{sky}} \\ b_{QT} * T_{\text{sky}} \\ b_{UT} * T_{\text{sky}} \end{bmatrix} \quad (4.3)$$

³By unsmoothed we mean that the Planck beam has been deconvolved from it.

Therefore, any non-zero power in the observed Q/U maps is solely due to $T \rightarrow P$ leakage from the temperature sky. The residual beams directly measure the convolution kernels b_{QT} and b_{UT} .

For each detector, we convolve a $2^\circ \times 2^\circ$ patch of its composite beam map (the ‘beam thumbnail’) with an unsmoothed Planck T -only sky (i.e. $Q_{\text{sky}} = U_{\text{sky}} = 0$). For that, an intermediate flat sky map is constructed around the detector’s trajectory (since our unpolarized FFBM maps are in local Cartesian coordinates x', y') by interpolating the HEALPix input map and its derivatives to carry out the convolution in the timestream domain. We then produce real detector timestreams by sampling the convolved maps along the pointing trajectory of each detector the same way as during CMB observations. The resulting timestreams undergo the same deprojection and mapmaking process as real data (sections 2.7.2- 2.7.5). This way, the pair-differencing and estimation of the kernels b_{QT} and b_{UT} is done organically during the mapmaking process.

The choice to convolve the per-detector composite beam maps using an intermediate flat sky map gets around having to take a spherical harmonic transform of the beam maps⁴ and convolving in spherical harmonic space. However, this process inherently injects a small amount $T \rightarrow P$ leakage during our filtering and deprojection processes that we refer to as ‘algorithmic floor’, which is easily quantifiable⁵ and is subtracted from the final leakage estimates. The beam thumbnail is small, therefore the flattened algorithmic floor is expected to be very small. For the case of BICEP3 the algorithmic floor is measured to be of $\mathcal{O}(10^{-4})$ ⁶ in BB , roughly a factor of 10 below the $T \rightarrow P$ leakage spectrum at every science bin. However, using larger beam thumbnails (e.g. $4^\circ \times 4^\circ$) would amplify this effect to a point that it could become relevant for our leakage estimates. If we want to directly simulate the impact of extended beam response from near and far sidelobes on $T \rightarrow P$ leakage, we would need a new algorithm that is not prone to this contamination. Such an algorithm is described in detail in section 5.1.

⁴This process is explored in chapter 5 in the context of a more subdominant systematic.

⁵By setting the pair-difference beam to zero and repeating the analysis.

⁶In principle, this value contains other small algorithmic effects such as imperfect matrix purification. Although that is shown to be negligible for $\ell < 150$.

Noise estimates of the leakage are also produced by forming ‘split beam maps’ described in detail in Chapter 4 of [51]. Finally, with the power and noise spectra of the $T \rightarrow P$ leakage maps in hand, we can cross-correlate them with real CMB maps to validate that our leakage estimates are truly present in the real data. To this end, we define the ρ minimum variance estimator originally introduced in [56]:

$$\hat{\mathbf{C}}^{BB} = \rho \langle \mathbf{C}^{BB} \rangle + n. \quad (4.4)$$

where $\langle \mathbf{C}^{BB} \rangle$ is the expected value from a theory BB spectrum for a given r (e.g. $r = 1$), and \mathbf{C}^{BB} is the measured BB spectrum from the leakage maps. n is the noise contribution. We estimate ρ using χ^2 minimization, resulting in:

$$\rho = \frac{\langle \mathbf{C}^{BB} \rangle^T \mathbf{N}^{-1} \mathbf{C}^{BB}}{\langle \mathbf{C}^{BB} \rangle^T \mathbf{N}^{-1} \langle \mathbf{C}^{BB} \rangle} \quad (4.5)$$

\mathbf{N} is the bandpower covariance matrix of the lensed Λ CDM + dust + noise simulations (from section 2.7.5). Since ρ is proportional to a BB spectrum of $r = 1$, it can be interpreted as ‘the amount of r ’ present in the leakage BB spectrum. Careful comparison of the cross spectra between leakage maps and real CMB maps is presented in Chapter 6 of [51] for BICEP3 with updated results for both BICEP3 and BICEP Array in progress. The BK18-era BB autospectrum and cross with real spectrum is shown in Figure 4.6.

We propagate the BB leakage spectrum through our multicomponent likelihood analysis described in section 2.8. We opt to inject the leakage BB spectrum in our 499 simulations of lensed Λ CDM + dust + noise at the power spectrum level. Then we find the shift in the recovered r value using ML search. The recovered bias on r was found to be $\Delta(r) = 1.5 \pm 1.1 \times 10^{-3}$, which is subdominant compared to our statistical uncertainty $\sigma(r) = 0.009$.




Figure 4.6. BICEP3 leakage BB autospectrum (black) and cross with real (blue dots) from [27]. The noise from split beam map simulations crossed with the 499 lensed Λ CDM + dust + noise simulations is used for the error bars. Although the blue dots mostly follow the trend of the BB autospectrum, $T \rightarrow P$ leakage predicted by beam simulations is not confidently detected in the real data.

4.1.4 Crosstalk Beams

All FFBM measurements are used for instrument performance validation. Any features in the beam maps that are not part of the main beam lobe can contribute to a variety of systematics if not properly accounted for in analysis. *Crosstalk Beams* are such features and are understood to impact the cross-correlation of the BB leakage spectrum with the real data ⁷ (Figure 4.6). They are compact secondary beams that originate from the readout system. As discussed in section 2.2, both BICEP3 and BICEP Array detectors are read out using SQUID-based time-domain multiplexing systems. The signal from the TES bolometers is inductively coupled to individual SQ1s, and then summed using a summing coil. Due to non-zero mutual inductance between the SQUIDS in the same SQ1 column, a fraction of the signal from one SQ1 gets picked up by the neighboring SQ1s in the same column. This results in small secondary beams that are frequently seen in beam maps and are validated by independent analyses of cosmic rays. This source of inductive crosstalk can only occur

⁷See chapter 6 of [51].

4. Beam Characterization Results

between detectors ‘upstream’ or ‘downstream’ the readout chain, which gives us a straightforward way to identify them in beam maps. Crosstalk beams are especially pronounced in BICEP3 FFBM measurements but less so in BICEP Array for which extra mitigation steps were taken during fabrication to mitigate their impact [76]. Figure 4.7 shows binned histograms of the BICEP3 upstream and downstream crosstalk beam locations from the main beam in distances of millimeters, for all detectors.




Figure 4.7. Binned histograms of upstream and downstream crosstalk partner distances on the focal plane. The plot illustrates the potential location of the crosstalk beams relative to the main beam. The BICEP3 plate scale that converts angles in the far field to distances on the focal plane is approximately 15 mm/degree⁹. The pattern seems complicated with some detectors having crosstalk partners nearly 3° away from the main beam. However, most crosstalk partners are at around 1° away from the main beam in both the X and Y orientation directions (see text for more details).

The BICEP3 plate scale is approximately 15 mm/degree. The pattern is complicated with crosstalk partners at 1° and 3° away from the main beam in both the X and Y orientations. This plot does not tell us that all detectors have crosstalk beams at these locations, just that these are the locations from the main beam we should expect them at. Notice that the pattern is symmetric about the main beam, under a 90° rotation which matches the module orientation of MCE0 that is oriented 90° with respect to the other MCEs. This is further corroborated by the number of hits in the histogram bins (1/4 of the hits along the $Y = 0$ axis vs the $X = 0$ axis). Detectors at the end/start of the MCE columns have an upstream/downstream partner that is in a different column, resulting in crosstalk beams that are further away from the main beam (at around 3°). From the phenomenology of the crosstalk beams, it is difficult to precisely predict the amplitude of the crosstalk beams, let alone if they cleanly subtract during the pair differencing process. For this reason, we worry

about their impact on the residual beams. Figure 4.8 shows the pair difference and residual beams of a single BICEP3 pair as measured in 2023 FFBM campaign, along with the various stages of beam-space deprojection. The pair is located at the center of the focal plane and belongs to module MCE1 that is oriented at 0° with respect to the focal plane orientation axis.




Figure 4.8. BICEP3 single detector pair difference and residual beams after various stages of beam-space deprojection. The crosstalk beam is clearly visible in the pair difference beam (top left) and remains present after deprojecting all leading order modes (bottom right). The colorscale is kept fixed across all plots to illustrate how the crosstalk beam remains a prominent feature in the residual beam. For this particular detector, the crosstalk beam is located at approximately $y' = \pm 1.05^\circ$ and $x' = 0^\circ$. This illustrates the complicated nature of crosstalk in the residual beams.

The crosstalk beams are clearly visible in pre and post deprojection for both the upstream and downstream locations, although the sign with which crosstalk contamination appears is

defined by the amplitude in the individual V and H beams ¹⁰.

We include an additional example of crosstalk beams from a BICEP3 detector pair belonging to module MCE0 that is oriented at 90° with respect to the focal plane orientation axis. As suggested by Figure 4.7, the crosstalk beams are rotated by 90° compared to the previous example.




Figure 4.9. BICEP3 single detector pair difference and residual beams after various stages of beam-space deprojection. The crosstalk beam is clearly visible in the pair difference beam (top left) and remains present after deprojecting all leading order modes (bottom right). The colorscale is kept fixed across all plots to illustrate how the crosstalk beam remains a prominent feature in the residual beam. For this particular detector, the crosstalk beam is located at approximately $y' = 0^\circ$ and $x' = \pm 1.05^\circ$.

If not accounted for, crosstalk beams directly source $T \rightarrow P$ leakage. For this reason, we

¹⁰If the amplitude of V crosstalk beam is higher than H, the pair difference will show a positive crosstalk beam, and vice versa.

extend our timestream-level deprojection scheme to include crosstalk deprojection. This is done the same way as in section 2.7.3. However, we copy the leakage template corresponding to differential gain and offset it on the sky by the known upstream and downstream crosstalk partners, resulting in two extra templates per detector. We extend the existing deprojection nomenclature to include crosstalk and be more explicit about the modes being deprojected. For example, deprojecting differential gain, pointing along the focal plane x-axis, pointing along the focal plane y-axis, differential ellipticity in the + and \times orientations, upstream crosstalk gain, and downstream crosstalk gain, but no differential beamwidth is denoted by *dp11101111*. Each number corresponds to a certain deprojection template and a 1/0 indicates whether it is deprojected or not ¹¹.

The amplitude of the crosstalk beams in the residual beams varies from pair to pair, depending on the mutual inductance between the SQUIDs in the readout chain. To characterize their amplitude we individually fit an elliptical Gaussian model to each crosstalk beam and location as:

$$A_{\text{crosstalk}} = \frac{\sum_{i,j}^n [V_{i,j}(x', y') - H_{i,j}(x', y')] G_{i,j}(x', y')}{\sum_{i,j}^n G_{i,j}^2(x', y')} \quad (4.6)$$

where $V_{i,j}$ and $H_{i,j}$ are the binned beam map values at pixel (i, j) , $G_{i,j}$ is the elliptical Gaussian model evaluated at pixel (i, j) ¹². The total crosstalk amplitude for each pair is then calculated as the sum of the absolute value of the upstream and downstream amplitudes. Figure 4.10 shows a focal plane view of the total crosstalk amplitude for BICEP3 detectors from the 2023 FFBM campaign.

The average crosstalk amplitude across all BICEP3 detectors from the 2023 FFBM campaign is found to be $\sim 0.37\%$. Discovering the presence of crosstalk in our residual beams and quantifying its amplitude has motivated the extension of our deprojection scheme to include

¹¹A 2 indicates subtracting by a fixed deprojection coefficient just like in section 2.7.3.

¹²The elliptical Gaussian's parameters are fixed to the band-averaged values. An approximation that would have little impact on the amplitude fit.

4. Beam Characterization Results

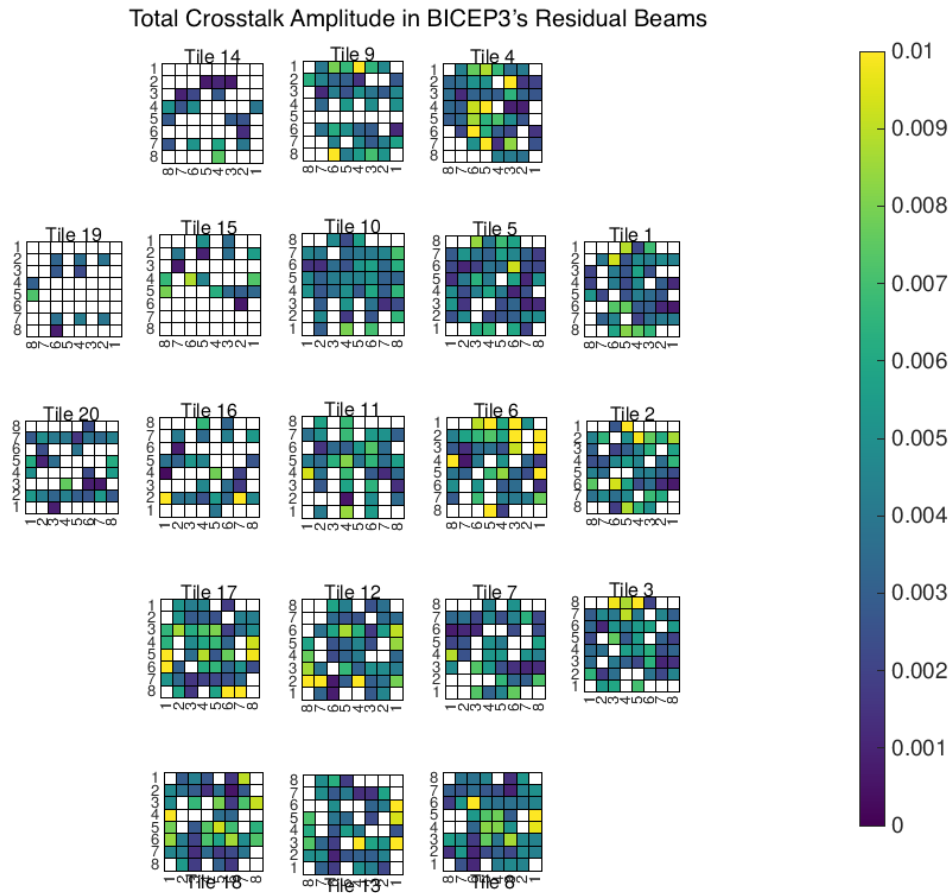


Figure 4.10. BICEP3 focal plane view of total crosstalk amplitude from the 2023 FFBM campaign. As expected the amplitude varies a lot across the focal plane as the crosstalk phenomenology suggests. The average total crosstalk amplitude across all detectors is approximately found to be $\sim 0.37\%$.

crosstalk deprojection. At the current level of sensitivity, crosstalk deprojection is necessary to pass our data consistency and null tests ¹³ and will be incorporated in BK analysis in future publications. We opt to deproject crosstalk as opposed to subtracting it using a fixed coefficient from unpolarized FFBM, since it is known that its amplitude can vary significantly depending on the amplitude of the calibration source being observed. In addition, although on average the crosstalk amplitude prediction from FFBM matches the amplitude measured by independent analyses of cosmic rays, there are systematic differences between the pair-to-pair crosstalk amplitudes between the two datasets.

4.2 Polarized Beam Mapping

In addition to unpolarized beam mapping measurements described in the previous section, we occasionally carry out *polarized* beam mapping measurements using the same FFBM setup. Although data products from these measurements are not directly used in BK analysis, they are used to reconstruct the per-detector polarized beam response, offering unique insights into the polarization performance of our instruments. Polarized FFBM data exist only for BICEP3 from the 2018 BSNS campaign and the 2022 RPS campaign.

The RPS campaign was carried out to (among other things) measure BICEP3’s per-detector polarization angles (ϕ_d in equation 3.22) with the results presented in Chapter 4 of [40] and [72] along with constraints on cosmic birefringence. This dataset consists of 10 equally spaced Dk angles for each of which the RPS stage is rotated to 13 equally spaced angles, resulting in 130 unique $Dk + \phi_s$ angle combinations. For a given Dk angle and source polarization angle ϕ_s , we scan across the RPS in the far field in $9^\circ Az \times 2^\circ El$ raster scans (stepping in El by 0.1° after each $9^\circ Az$ scan) to map out the beam response. The raster is repeated at the same Dk angle but at 13 different RPS angles ϕ_s to form a *rasterset*. We perform 10 total rastersets at 10 different Dk angles to reconstruct the polarized beam response. Special care was taken to ensure that the source’s power setting is stable throughout

¹³See chapter 5 of [38].

4. Beam Characterization Results

the entire campaign (thermal control via PID) and attenuated by $\sim 30dB$ to roughly match the power levels of the thermal chopper. This minimizes any concerns of saturation near the main beam.

The BSNS campaign, consists of 9 unique Dk angles (13 total schedules with some repeats), working around BICEP3's limited boresight rotation capability of 280° instead of 360° due to the cable wrap and pulsetube cryocooler mounting scheme [35]. The key difference with the RPS campaign being that the source's polarization axis is fixed to be aligned with gravity during the entire campaign (no rotation of the RPS stage). The calibration measurement setup and schedules are the same as those used for unpolarized beam mapping. Unlike the RPS campaign, the source is not attenuated and although it is in the far field, we would expect certain level of saturation or non-linearity of the detectors near the main beam.

The analysis of both these datasets follows the same steps as unpolarized beam mapping (described in section 3.5) up to the binning of the component maps¹⁴. Figure 4.11 shows the binned component maps at different Dk angles for a single BICEP3 detector from the 2018 BSNS campaign.




Figure 4.11. BICEP3 component maps for a single pair at different Dk angles from the 2018 polarized FFBM campaign, plotted in 2×2 degree thumbnails. The source's polarization angle is fixed to be aligned with gravity during the entire campaign. The detector's response is minimum when orthogonal to the polarization axis of the source and maximum when aligned with it. The individual maps are normalized to the peak of the brightest component map for V and H separately for visualization purposes. Figure adapted from [73]. Note the gap in Dk coverage between boresight angles of 135° and 293° due to the limited boresight rotation capability of BICEP3.

We expect the detector response to be minimum when the detector's co-polar axis is orthogonal to the polarization axis of the source, and maximum when parallel. Since the source's

¹⁴In [40] a slightly different analysis of the RPS timestreams is described. For this thesis, the RPS timestreams have been re-analysed from scratch under the FFBM analysis pipeline. The difference in the resulting maps is negligible.

polarization axis is fixed, the amplitude of the measured beam solely depends on the Dk angle of observation. To reconstruct the polarized beam response, we need precise knowledge of the angle between the detector's co-polar axis and the polarization axis of the source (ϕ in equation 3.22).

In this section, we describe the polarized beam mapping process that applies to both BSNS and RPS campaigns and show the corresponding beam characterization results. The main results of the RPS campaign including modulation curves and detector polarization angle systematics are presented in [40]. They will be used in Chapter 5 in conjunction with the BSNS beam maps to constrain the impact of polarization-to-polarization leakage due to cross-polar response in our CMB data.

4.2.1 Per-detector Mapmaking for Polarized Beam Mapping

In general, an ideal detector's response to a polarization field dominated by that of the BSNS source in the far field is given by:

$$D = \int [\mathbf{B}_T(\mathbf{x})T_{\text{source}} + \mathbf{B}_p^T \mathbf{R} \mathbf{P}] \delta(\mathbf{x} - \mathbf{x}') d\mathbf{x}' \quad (4.7)$$

where:

$$\mathbf{B}_p = \begin{bmatrix} B_Q(\mathbf{x}) \\ B_U(\mathbf{x}) \end{bmatrix}, \quad \mathbf{R} = \begin{bmatrix} \cos 2\phi & -\sin 2\phi \\ \sin 2\phi & \cos 2\phi \end{bmatrix}, \quad \text{and} \quad \mathbf{P}_{\text{source}} = \begin{bmatrix} Q_{\text{source}} \\ U_{\text{source}} \end{bmatrix} \quad (4.8)$$

where B_T is the detector's unpolarized beam response, B_Q and B_U are the polarized beam responses, T_{source} is the temperature and $Q_{\text{source}}, U_{\text{source}}$ are the polarization fields of the source expressed in the Stokes basis. Unlike the thermal chopper source, the size of the BSNS aperture is a few centimeters making it effectively a point source in the far field. We define the Stokes Q axis to be aligned with the source's polarization axis, and since it is linearly polarized, we have $U_{\text{source}} = 0$. The key assumption here being that the source's cross-polar

response can be ignored, which is appropriate since the polarizer grid in front of the source yields a cross-polar response of $< 0.01\%$. An alternative sensible choice would be to define the integral over $\hat{B}_U = 0$ (i.e. zero monopole response along the axis perpendicular to the source's polarization axis). The two definitions result in similar outcomes in practice [77]. With this setup we expect B_Q to represent the detector's polarized beam response and B_U to represent the cross-polar response. Cross-polar response is defined as detector response to electric fields that are orthogonal to its co-polar axis.

In an environment dominated by the field of the BSNS, a single detector measures the time-averaged power along its co-polar axis. Therefore, equation 4.7 simplifies to ¹⁵:

$$d = B_T + B_Q \cos 2\phi + B_U \sin 2\phi + n \quad \text{with} \quad n \sim \mathcal{N}(0, \sigma^2) \quad (4.9)$$

where n is the noise contribution to the measurement assumed to be normally distributed with zero mean and variance σ^2 . The angle ϕ that is used for mapmaking (equivalent to its definition in equation 3.22) is:

$$\phi = Dk + 90^\circ - DA + \phi_d - \phi_s \quad (4.10)$$

where Dk is the angle of observation, DA is the drum angle that the receiver sits on the Dk stage and ϕ_s is the polarization axis of the source projected onto the focal plane. For the BSNS measurements, ϕ_s is fixed to be aligned with gravity and is not rotated during the entire campaign. For the RPS measurements, ϕ_s takes on 13 different values as the RPS stage is rotated during a raster set.

To reconstruct the polarized beam response we employ a mapmaking method similar to that in section 2.7.2. However, our systematics studies require per-detector, timestream-level simulated TODs. For this reason, we need per-detector polarized beam maps as opposed to pair-sum/difference maps that are used in CMB analysis. For i observations at angles ϕ_i

¹⁵Since $U_{\text{source}} = 0$, T_{source} and Q_{source} match and are absorbed into the definition of B_T and B_Q for simplicity.

that fall into a given pixel in the beam map, we project the per-detector timestreams onto the Stokes Q, U axes by calculating the following mapmaking sums:

$$\langle z \rangle = \sum_i w_i d_i \quad (4.11)$$

$$\langle x \rangle = \sum_i w_i d_i \cos 2(\phi_i) \quad (4.12)$$

$$\langle y \rangle = \sum_i w_i d_i \sin 2(\phi_i) \quad (4.13)$$

where w_i is the weight of the measurement at pixel i . The datasets are taken at very stable conditions and the source's enclosure is temperature controlled, resulting in negligible gain variation over the course of both RPS and BSNS campaigns. In addition, we expect very little noise near the main beam where Dk coverage is good.

Expanding the expressions for $\langle z \rangle$, $\langle x \rangle$, and $\langle y \rangle$ in equation 4.13, they can be written in matrix form as:

$$\begin{bmatrix} \langle z \rangle \\ \langle x \rangle \\ \langle y \rangle \end{bmatrix} = \begin{bmatrix} \sum w_i & \sum w_i \cos 2\phi_i & \sum w_i \sin 2\phi_i \\ \sum w_i \cos 2\phi_i & \sum w_i \cos^2 2\phi_i & \sum w_i \sin 2\phi_i \cos 2\phi_i \\ \sum w_i \sin 2\phi_i & \sum w_i \sin 2\phi_i \cos 2\phi_i & \sum w_i \sin^2 2\phi_i \end{bmatrix} \begin{bmatrix} \hat{B}_T \\ \hat{B}_Q \\ \hat{B}_U \end{bmatrix} \quad (4.14)$$

where N_{hits} is the number of measurements in pixel i and \hat{B}_T , \hat{B}_Q , and \hat{B}_U denote estimated quantities. We invert the matrix to solve for the estimated beam responses as:

$$\begin{bmatrix} \hat{B}_T \\ \hat{B}_Q \\ \hat{B}_U \end{bmatrix} = \mathbf{M}^{-1} \begin{bmatrix} \langle z \rangle \\ \langle x \rangle \\ \langle y \rangle \end{bmatrix} \quad (4.15)$$

where \mathbf{M}^{-1} is the inverse of the matrix in equation 4.14, the 3-dimensional version of the QU decorrelation matrix in equation 2.24. Just like in the case of CMB mapmaking, the matrix \mathbf{M} is not invertible in the case of a single measurement at angle ϕ . Beam map pixels that receive

4. Beam Characterization Results

one or two measurements at adjacent Dk angles will result in \mathbf{M}^{-1} being poorly conditioned and potentially non-invertible. This is especially relevant for the BSNS measurements since the source orientation is fixed and BICEP3's boresight rotation is limited, resulting in a gap in Dk coverage. For RPS measurements, since we reconstruct the polarized beam response using both Dk and ϕ_s angles, this is not a concern. The determinant of \mathbf{M} is the perfect metric to quantify the conditioning of the matrix at every pixel in the map. By defining:

$$\begin{aligned}
 C &= \sum_i \cos 2\phi_i, & S &= \sum_i \sin 2\phi_i, \\
 C_2 &= \sum_i \cos^2 2\phi_i, & cs &= \sum_i \cos 2\phi_i \sin 2\phi_i, \\
 S_2 &= \sum_i \sin^2 2\phi_i, & N &\equiv \# \text{ of measurements}
 \end{aligned} \tag{4.16}$$

we can write the determinant as:

$$\Delta \equiv \det(\mathbf{M}_V) = N(C_2 S_2 - cs^2) - C^2 S_2 - S^2 C_2 + 2CS cs \tag{4.17}$$

$$\approx \frac{N}{2} \sum_{i,j} \sin^2 (2\phi_i - 2\phi_j). \tag{4.18}$$

where in the last step we expand and assume that the sums over $\cos 2\phi_i$ and $\sin 2\phi_i$ are approximately zero¹⁶. To further illustrate the intuition behind the determinant, we define the 3-dimensional analog of vector \mathbf{s}_i from equation 2.23 as $\mathbf{S}_i = [1, \cos 2\phi_i, \sin 2\phi_i]^T$. These vectors now live in \mathbb{R}^3 . Therefore, $\mathbf{M} = \sum_i^N \mathbf{S}_i \mathbf{S}_i^T$ is a sum of outer products of the measurement vectors which makes \mathbf{M} positive semi-definite. Its determinant Δ is either positive, or is zero if the measurement vectors are co-planar/co-linear. In other words, the determinant is a measure of how well the measurement vectors span the 3-dimensional space. In the case of the BSNS measurements, the measurement vectors are defined by the ϕ angles at which we take data. Therefore, the determinant is a measure of how well we estimate \hat{B}_T , \hat{B}_Q , and \hat{B}_U at each pixel in the beam map. Finally, we calculate the covariance matrix of the estimated

¹⁶A good approximation when the measurement angles ϕ_i are equally spaced.

beam responses which is equal to $\sigma^2 \mathbf{M}^{-1}$ and gives us the per-pixel variance estimates for \hat{B}_T , \hat{B}_Q , and \hat{B}_U :

$$\text{Var}(T) = \sigma^2 \cdot \frac{C_2 S_2 - c s^2}{\Delta} \quad (4.19)$$

$$\text{Var}(Q) = \sigma^2 \cdot \frac{N S_2 - S^2}{\Delta} \quad (4.20)$$

$$\text{Var}(U) = \sigma^2 \cdot \frac{N C_2 - C^2}{\Delta} \quad (4.21)$$

In practice, we use both the determinant and the inverse variance maps for weights at each pixel in the mapmaking process to produce per-detector polarized (\hat{B}_Q , \hat{B}_U) and unpolarized (\hat{B}_T) beam maps. The resulting array averaged maps for BICEP3 from the BSNS FFBM campaign are shown in Figure 4.12. The \hat{B}_Q and \hat{B}_U maps are normalized to the peak of the BSNS \hat{B}_T map. The resulting array averaged maps from the RPS campaign are shown in Figure 4.13, normalized to the peak of the RPS \hat{B}_T map¹⁷.

Since the BSNS is linearly polarized, and by definition the Stokes Q axis is aligned with the source's polarization axis, we expect \hat{B}_Q to be similar to the unpolarized beam \hat{B}_T . The four-fold features at radius of $\approx 4^\circ$ and $\approx 13^\circ$ are real optical features present in *all* BICEP3 calibration measurements with good enough SNR. They arise from diffraction off the metal mesh Cardiff filters [37]. The \hat{B}_U map, which represents the detector's cross-polar response to the calibration source, is consistent with zero to first order (i.e. the integral over \hat{B}_U is zero). It is dominated by an overall quadrupole pattern somewhat similar to the residual unpolarized beam seen in Figure 4.5, although the two are not related. In the main beam region, higher order cross-polar response modes are present at sub-percent levels. These modes are expected to systematically mix E-mode with B-mode power ($E \rightarrow B$ leakage) and are explored in detail in Chapter 5.

The mapmaking procedure hinges on precise knowledge of the angle ϕ between the detector's co-polar axis and the source's polarization axis. Any overall angle miscalibration would rotate

¹⁷In both datasets the peak is determined by the 7-parameter Gaussian fit to the \hat{B}_T map. See section 3.5.2.




Figure 4.12. BICEP3 array averaged \hat{B}_T , \hat{B}_Q , \hat{B}_U beam maps from the BSNS FFBM campaign. The maps are inverse variance and determinant weighted to remove poorly conditioned regions. Figure adapted from [73]. The color scale is in normalized dB to the peak of the \hat{B}_T response (\hat{B}_U uses a negative logarithmic scale normalized the same way). Looking at the \hat{B}_Q response, the $(x', y') > 0$ quadrant is noticeably noisier than the rest of the map, and corresponds to the gap in Dk coverage between boresight angles of 135° and 293° . As expected with a linearly polarized source, \hat{B}_Q is very similar to the \hat{B}_T , while the \hat{B}_U map shows no monopole response.




Figure 4.13. BICEP3 array averaged \hat{B}_T , \hat{B}_Q , \hat{B}_U beam maps from the RPS campaign. The color scale is in normalized dB to the peak of the \hat{B}_T response (\hat{B}_U uses a negative logarithmic scale normalized the same way). Just like the BSNS coadded maps \hat{B}_Q is very similar to \hat{B}_T , while the \hat{B}_U map shows no monopole response. The maps become noisier beyond 2° from the main beam in part due to the faster raster scans. The Az , El range of the raster scans is also smaller than the BSNS campaign which results in less coverage of the near sidelobes, hence the tighter choice of x and y -axis limits compared to the BSNS coadds.

some beam response from \hat{B}_Q into a monopole response in the \hat{B}_U map (see equation 5.6). For this reason, it is important to draw a distinction between monopole and higher order cross-polar response modes. The monopole response is defined as the rotationally invariant response in the \hat{B}_U map and can be ‘rotated out’ by the overall angle correction: $\psi = \frac{1}{2} \arctan\left(\frac{A_U}{A_Q}\right)$ where A_Q and A_U are the fitted amplitudes of the \hat{B}_Q and \hat{B}_U maps respectively. The higher order modes cannot be ‘rotated out’ and represent real cross-polar response that integrates to zero over the beam map.

Both the per-detector and array-averaged \hat{B}_U maps show no monopole response, which validates the precision of the per-detector polarization angles measured in [40] and used in cosmic birefringence constraints in [72]. By comparing the peak polarized response with the peak thermal response, we get an indirect measurement of polarization efficiency. This comes out to $99.1\% \pm 0.9\%$ where the uncertainty is listed as two times the standard deviation of the distribution of polarization efficiency across the focal plane. This falls well within the expected range of polarization efficiency for BICEP3 detectors and is consistent with results in [71], further establishing the overall consistency of these measurements. A caveat to the polarization efficiency estimate from BSNS data is that subtle signs of saturation have been observed in the main beam region of the \hat{B}_T maps for some detectors. It is also worth noting that the BSNS measurement was not designed to measure polarization efficiency, but this indirect estimate serves as a metric to understand the impact of saturation/non-linearity effects in the BSNS dataset. Given how close this estimate is to the nominal value of $\approx 99.7\%$, the effect must be present only for a small subset of detectors but definitely something that needs to be accounted for in systematics studies with these maps.

The RPS dataset offers a more precise measurement of the main beam since care was taken to avoid any signs of saturation or non-linearity in the detectors near the main beam (verified by the resulting modulation curves in [40]). To best understand these effects in the BSNS dataset, we take azimuthal averages of the \hat{B}_T and \hat{B}_Q maps for both datasets and compare the resulting radial profiles. The normalization of all the profiles is set to be the peak of the \hat{B}_T profile from the RPS dataset which is expected to have no saturation or non-linearity

effects¹⁸. To do this, the two profiles are regressed between 1° and 3° since both datasets have good SNR and no pathologies in this region. The regression coefficient represents the relative gain between the two datasets and is used to set the normalization of the BSNS profiles to match that of the RPS profiles. The Pearson correlation coefficient R of the regression is used as a metric to establish the level of saturation/non-linearity in the BSNS dataset. A threshold of $R > 0.99$ is used to flag outliers, resulting in $\sim 1.5\%$ of detectors being excluded from downstream analysis. The resulting gain matched profiles of an outlier detector exhibiting signs of unusually high saturation in the BSNS dataset, are shown in Figure 4.14.




Figure 4.14. Azimuthally averaged radial profiles of \hat{B}_T (blue) and \hat{B}_Q (orange) for a single BICEP3 detector for the 2018 BSNS (dashed) and 2022 RPS (solid) polarized FFBM campaigns. The full beam profile is on the left and a zoomed-in view of the main beam is on the right. This detector was specifically picked to illustrate an outlier case of very high saturation in the BSNS dataset. Most detectors show much better agreement between the two datasets. The RPS dataset resolves the main beam very well, while the BSNS dataset measures the near sidelobes with much better signal-to-noise.

The final data products from the polarized FFBM analysis are per-detector maps that combine the high-fidelity main beam measurement from the RPS dataset with the superior SNR in the near sidelobes from the BSNS dataset. Since the two datasets are gain matched for each detector, we stitch them together by taking the RPS measurement within 1° from the main beam and the BSNS measurement beyond 1° . The resulting maps are used in Chapter 5

¹⁸Both \hat{B}_T and \hat{B}_Q profiles are normalized to the peak of \hat{B}_T . Therefore, any difference between the \hat{B}_Q and \hat{B}_T profiles hints at decreased polarization efficiency.

to constrain the impact of polarization-to-polarization leakage due to cross-polar response in our CMB data.

4.3 Far Sidelobe Mapping

To fully characterize the optical performance of our instruments, we perform measurements of the regions beyond the forebaffle cutoff, by far sidelobe (FSL) mapping. The goal is to identify potential sources of contamination and to validate the performance of our baffling scheme. As discussed in section 2.2 we employ comoving, absorptive forebaffles mounted on top of each receiver’s window to minimize the coupling of our detectors to the surrounding environment. These forebaffles substantially suppress sidelobe power at angles larger than $\approx 20^\circ$ away from the main beam. However, sidelobe power beyond the forebaffle increases optical loading if not absorbed by the forebaffle or redirected to the sky by the ground shield. This could lead to excess contamination that would be difficult to constrain during CMB observations. During FSL mapping, we place the BSNS source atop the mast located in the same building as the receiver being mapped (not in the far field) and take data by scanning 360° in azimuth. By repeating this procedure at various source settings (attenuation, polarization angle, and with the forebaffle installed and uninstalled), we map the far sidelobes of our instruments and directly probe the coupling to the forebaffles. The results are then compared with independent measurements of forebaffle loading offering a powerful cross-check of both the understanding of the far sidelobes and the independent loading measurements. The FFBM procedure described in section 3.2 is not suited to map the angular regions beyond the forebaffle cutoff due to the limited coverage of the redirecting mirror. In this section we describe the FSL measurement and results for BICEP3 and BA30/40. Plans for future FSL measurements of the BA150 and BA220/270 receivers are under consideration for the 2026/2027 austral summer season. However, the importance of FSL measurements is subdominant compared to unpolarized FFBM with the thermal chopper. In this section, we loosely follow the FSL description in section 4 of [73].




Figure 4.15. BA30/40 during FSL measurements in March 2021. The comoving forebaffle of BA30/40 is uninstalled during this schedule of sidelobe measurements while those for the Keck Array receivers are installed. Note the 40 GHz calibration source mounted on the mast in the near field. The Dark Sector Laboratory (DSL) hosting BICEP3, and the South Pole Telescope (SPT) can be seen in the distance. Photo from [73].

Measurement Setup

As indicated in Table 4.1, FSL mapping measurements do not utilize the redirecting mirror nor observe a source in the far field. Instead, the BSNS source is placed on top of the mast at the same building as the instrument (MAPO for BA and DSL for BICEP3), which brings it to an elevation of roughly 70° as seen by the instrument. The mast stage is tilted such that it points directly at the receiver. Figure 4.15 illustrates the setup of the FSL mapping measurements for BA30/40 during the austral summer of 2021/2022.

The 40 GHz BSNS is very similar in construction to the 100 GHz source described in section 3.3 but with fewer frequency amplifiers in the RF chain to achieve the desired output frequency. The band pass of the source is also sufficiently wide to cover both the 30 GHz

and 40 GHz detectors.

Scan Strategy

There are a total of 7 parameters we can vary to achieve the optimal scan strategy for FSL mapping summarized in table 4.3. We have the option to install a 90° precision twist before the feedhorn to take measurements at two orthogonal source polarizations. In addition, to maximize the signal-to-noise ratio in the different beam regions, we take measurements at 3 different source attenuation settings. Each attenuation setting is specifically chosen to map the desired beam region with no saturation. The more sensitive titanium transition is used when mapping at no source attenuation, while the rest of the measurements use the aluminum transition. For each source setting, we take measurements at 6 boresight angles (equally spaced only for BA), both with the forebaffle installed and uninstalled.

Parameter	BICEP3	BA30/40
<i>Az</i> speed (deg/s)	1.4	1.4
<i>El</i> step (deg)	0.5	0.5
<i>Az</i> range (deg)	380	410
<i>El</i> range (deg)	34	36.8
Precision Twist	On & Off	On & Off
Forebaffle	On & Off	On & Off
Source Attenuation (dB)	-35, -25, 0	-25, -18, 0
Schedules per source setting	6	6

Table 4.3. Observation setup parameters used for FSL mapping for BICEP3 and BA30/40. Data is taken at 6 equally spaced boresight angles, both with the forebaffle installed and uninstalled. The schedules are repeated after installing a 90° precision twist before the feedhorn to take measurements at two orthogonal source polarizations. The entire set of measurements is repeated at various source power settings to map different beam regions with good signal-to-noise.

The source is electronically modulated at 16 Hz. Raster scans as opposed to out-and-back scans are taken for both receivers, since (as implied by table 4.3) we need > 70 schedules to cover all measurement degrees of freedom. At roughly 9 hr per schedule, this is nearly a full month of FSL observations!

Analysis

The raw timestreams are analysed the same way as the FFBM measurements described in section 3.5 but with no masking and a 0.5° pixel size. The elliptical Gaussian fits are performed at the timestream level only to the low source power attenuation setting (e.g. -35dB for BICEP3) since the main beam is completely saturated at other attenuation settings. The pixel window function effects are relevant given that the main beam size of BICEP3 is approximately $\sigma = 0.166^\circ$. Therefore, we do not hope to resolve any main beam features in the beam maps. With this procedure we produce per-detector, per-source attenuation setting, per source orientation state, per-forebaffle state, per- Dk angle beam maps. The next step is to gain match the twist and no-twist maps to produce unpolarized maps. This is done by calculating the source power ratio between the twist and no-twist maps for each detector and applying it before coadding. Let the response of a V detector observing the source with the precision twist installed be $r_{V,twist} \propto g_V s_{twist}$ where s_{twist} is the source power in the twist state and g_V is the gain of the detector¹⁹. Then the source power ratio between the twist and no-twist states is given by:

$$\frac{r_{V,twist} r_{H,twist}}{r_{V,notwist} r_{H,notwist}} = \frac{s_{twist}}{s_{notwist}} \quad (4.22)$$

Using the source power ratio, we gain match the twist and no-twist maps to produce unpolarized maps for each detector, source attenuation setting, forebaffle state, and Dk angle. Having reduced the dimensionality of the dataset by a factor of 2, we proceed with the data reduction process by compositing over the six Dk angles²⁰ to produce a single map for each detector, source attenuation setting, and forebaffle state.

¹⁹There is an angle orientation factor here in terms of the mapmaking angle ϕ defined in 3.22, but not included here, since it cancels out in 4.22.

²⁰The coaddition process is simply median averaging across the six Dk angles.

Stitching power levels

At this stage the only measurement degrees of freedom left are the source power setting and the forebaffle state. The final step of the analysis process is to stitch the per-detector maps across the different source power settings. The resulting data products would be high fidelity beam maps for each detector in the near and far sidelobes (not the main beam). Prior to any stitching, automatic cut thresholds are applied to the maps for each power setting. For BICEP3 cuts are taken from unpolarized FFBM with the thermal chopper from the same year, while for BA30/40 manual cuts are applied to remove non-functioning or poorly performing detectors ²¹. Now, with only good detectors in hand, we proceed with the stitching process. Starting with the lowest power setting, we regress against the map from the next power setting in an overlapping region where both maps have good signal-to-noise and no signs of saturation. To define these regions, we calculate azimuthally averaged beam profiles, which allow us to pick both the stitching and regression regions.

The azimuthally averaged beam profiles at different source power settings can be thought of as vectors. The inner product of these vectors (post normalization) returns the region that the two match or diverge. An inner product close to 1 indicates that the two profiles match well and define the regression region. Conversely, an inner product < 1 indicates that the two profiles start to diverge and define the saturation region of the higher power map. This scheme hinges on picking the correct thresholds for the two regions, which are $[0.99, 0.8]$ for the regression region and < 0.8 for the saturation region. Once the two regions are identified, linear regression is carried out at the map level and the resulting regression coefficient is applied to the lower power map prior to stitching. The data within the stitching region is replaced by the lower power map since it is expected to have better signal-to-noise and less saturation. The process is repeated to bring the resulting map up to the highest power setting (e.g. no source attenuation in the titanium transition). The resulting array coadded maps for BICEP3 are shown in Figure 4.16 for the forebaffle on and off cases.

²¹This is only feasible since BA30/40 has less than 700 optically coupled detectors across both its frequency bands.

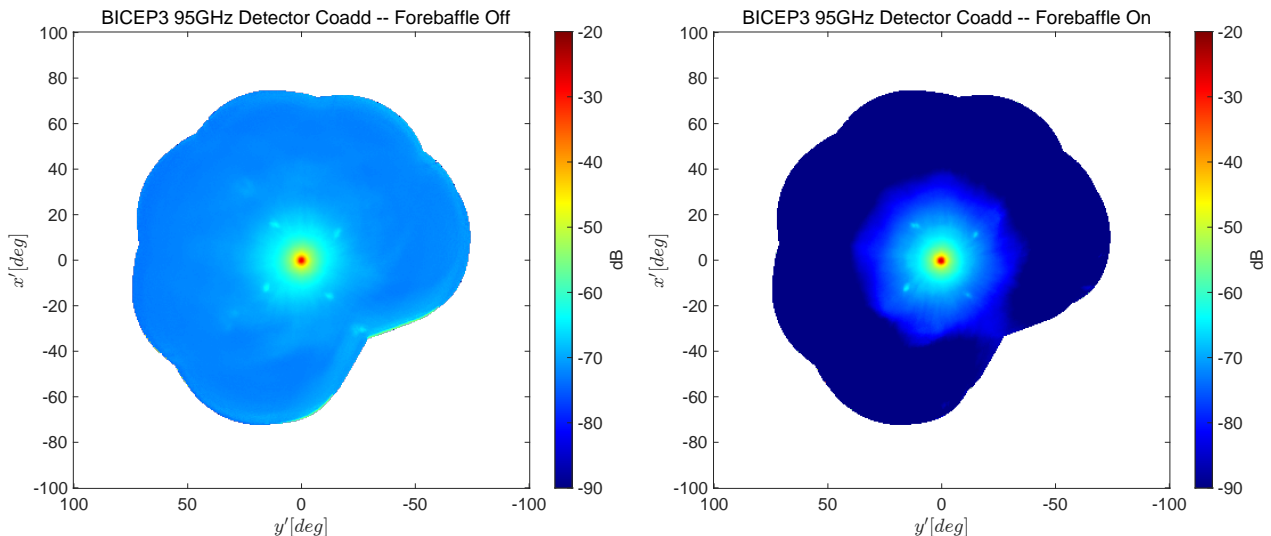


Figure 4.16. BICEP3 far sidelobe array coadded maps with the forebaffle off (left) and on (right). The maps are normalized to the peak of the main beam response of the thermal chopper as measured during FFBM as described in section 4.1 (see text for details). The colorbar is the same for both maps to allow for direct comparison. Notice the range in the x' , y' axes compared to Figure 4.12. The forebaffle substantially suppresses the far sidelobe response at angles larger than $\approx 15^\circ$ away from the main beam. The flower-like structure of the beam map is the result of compositing over the six Dk angles. Figure adapted from [73].

To interpret the FSL maps further, we normalize them to the main beam response of the thermal chopper as measured during FFBM (section 4.1). At the map level, comparison of FSL data products with unpolarized FFBM data is difficult primarily due to the different binning schemes (0.1° bins for unpolarized beam mapping and 0.5° bins for FSL mapping). However, this complication is mitigated when comparing them in azimuthally averaged space. Using a process similar to the one used for power stitching, we regress the two datasets in azimuthally averaged space between 2° and 8° for BICEP3 and 5° to 9° for BA30/40. The resulting regression coefficient is used to normalize the FSL maps to the main beam response of the thermal chopper. On average the regression coefficients are found to be ~ 21.3 dB for BICEP3, ~ 39.4 dB for BA40, and ~ 38.8 dB for BA30, for both forebaffle states. The corresponding coadds for BA40 are shown in Figure 4.17.

The main difference between the FSL maps of BICEP3 and BA30/40 is attributed to the lack of the -35 dB source attenuation setting for BA30/40 which results in the main beam being completely unresolved in the FSL maps due to saturation. Although the forebaffle attenuation is obvious, faint stripes at radii of $\approx 80^\circ$ are present in the BA40 forebaffle-on map. These have a 12-fold symmetry and are consistent with reflections off the dodecagonal

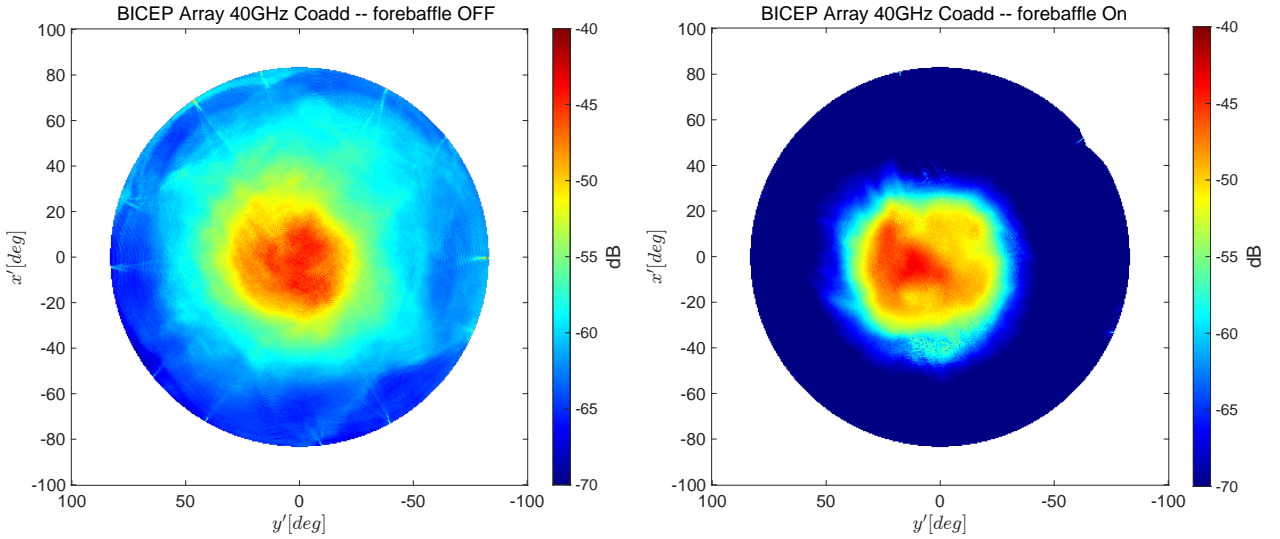


Figure 4.17. BA40 far sidelobe array coadded maps with the forebaffle off (left) and on (right). The maps are normalized to the peak of the main beam response of the thermal chopper as measured during FFBM which peaks roughly 40 dB higher. The colorbar is the same for both maps to allow for direct comparison. Notice the range in the x' , y' axes compared to Figure 4.12. The forebaffle substantially suppresses the far sidelobe response at angles larger than $\approx 20^\circ$ away from the main beam. Contrary to the BICEP3 FSL maps, the main beam is completely unresolved due to saturation. The stripes seen at radii of $\approx 80^\circ$ were originally thought to be due to sun contamination but were later confirmed to be due to reflections off the 12-edged ground shield surrounding BA. These can be faintly seen in the forebaffle-on maps as well although significantly suppressed.

ground shield surrounding BA. Although at heavily suppressed levels, this is the first time we see evidence of reflections from the groundshield in our measurements, hinting that multipoint reflections could optically couple to the receiver even in the presence of the forebaffle. The impact of these reflections on CMB observations is expected to be negligible.

Forebaffle Loading Measurements

The ratio of integrated power of the FSL map with the forebaffle on to the integrated power with the forebaffle off gives an estimate of the fraction of power intercepted by the forebaffle. We refer to this as the *forebaffle sidelobe power ratio*. By knowing the average temperature of the (temperature controlled) forebaffle we can convert this ratio to temperature units. The forebaffle power ratio is found to be $\approx 1.19\% \pm 0.32\%$ for BICEP3, $\approx 2.1\% \pm 0.50\%$ for BA30, and $\approx 1.4\% \pm 0.31\%$ for BA40, where the uncertainties correspond to 68% scatter across detectors. The average temperature of the forebaffle during the observations was measured to be ≈ 270 K for BICEP3 and ≈ 271 K for BA30/40. The resulting forebaffle

loading from FSL measurements is summarized in Table 4.4.

During each summer season at the South Pole, we routinely take measurements to directly probe the forebaffle loading. The loading is derived by comparing the detector saturation power extracted from partial load curves (PLCs) taken during different forebaffle configurations. These include measurements with the forebaffles on and off, as well as with various absorptive or reflective linings installed to establish measurement consistency. Independent analysis of these measurements yields direct forebaffle loading estimates that can be compared with those derived from FSL maps.

Receiver	FB sidelobe power [K_{RJ}]	Predicted FB loading [K_{RJ}]	FB loading measured [K_{RJ}]
BICEP3 [95]	3.21	3.2	3.3
BA30	5.67	5.6	5.6–7.6
BA40	3.78	3.7	3.6–4.0

Table 4.4. Comparison of measured forebaffle (FB) sidelobe power, predicted FB loading, and measured FB loading. The forebaffle loading is derived from FSL maps, the predicted loading from the forebaffle design [35], and the measured loading from direct measurements of the forebaffle loading. The table illustrates the consistency between two independent measurements validating the analysis of the FSL maps and our understanding of the forebaffle performance. BA30 has noticeably higher forebaffle loading than BA40, consistent across measurements.

As illustrated in Table 4.4, the loading estimates from the FSL maps are consistent with the direct forebaffle loading measurements, and with those predicted from the forebaffle design. This consistency validates the analysis of the FSL maps and that the somewhat uncertainty-prone power stitching procedure is working as intended (at least on average across a frequency band).

Receiver Response Characterization

With the FSL maps in hand, having verified their consistency and normalization we proceed to stitch them with the unpolarized FFBM maps in azimuthally averaged space to produce composite beam profiles that cover the main beam, near sidelobes, and far sidelobes. The resulting composite beam profiles for BICEP3 and BA40 are shown in Figure 4.18 and Figure 4.19 respectively.

4. Beam Characterization Results




Figure 4.18. BICEP3 composite beam profiles for the forebaffle on and off. The profiles are normalized to the peak of the main beam response of the thermal chopper as measured during FFBM as described in section 4.1 (see text for details). The forebaffle substantially suppresses the far sidelobe response at angles larger than $\approx 15^\circ$ away from the main beam, although the effect is spread out due to averaging. The bands around the profiles represent the profiles of 99% detectors in the array, giving a sense of FPU scatter. Data quality cuts have been applied on top of the automatic cuts described in section 3.5 to remove poorly behaved detectors. This plot characterizes the angular response of BICEP3 and is meant to be compared with Figure 17 of [48] which shows the composite beam profile of the BICEP1 receiver. Stitching the FSL maps with the unpolarized FFBM maps occurs at $\approx 4.8^\circ$ in a region where both datasets have excellent signal-to-noise and no signs of saturation. Figure updated from [73].




Figure 4.19. BA40 composite beam profiles for the forebaffle on and off. The profiles are normalized to the peak of the main beam response of the thermal chopper as measured during FFBM as described in section 4.1. The forebaffle substantially suppresses the far sidelobe response at angles larger than $\approx 20^\circ$ away from the main beam (further compared with BICEP3). The bands around the profiles represent the profiles of 99% detectors in the array, giving a sense of FPU scatter. This plot characterizes the angular response of BA40 and illustrates the attenuation of the far sidelobes by the forebaffle. Stitching the FSL maps with the unpolarized FFBM maps occurs at $\approx 6.3^\circ$ which is further out compared to BICEP3 due to lack of the -35 dB source attenuation setting for BA30/40.

4. Beam Characterization Results

The forebaffle attenuation (although smoothed by the effect of averaging) starts at $\approx 15^\circ$ for BICEP3 and $\approx 20^\circ$ for BA40, which is consistent with the forebaffle design. However, the attenuation level is roughly 10 dB higher for BICEP3 compared to BA40. This could be due to the bandwidth difference between the two instruments, as well as more ambient power coupling hinting at increased sidelobe loading for BA30/40 reflected in table 4.4.

Chapter 5

BICEP3 Systematics Studies With Linearly Polarized Source

The leading order systematic contamination in CMB polarization measurements aimed at detecting primordial B-modes is the leakage of the much brighter temperature signal into the polarization signal ($T \rightarrow P$ leakage — section 4.1.3). However, given the current sensitivity of our measurements, (BICEP3 has been observing since 2016!) we want to quantify a more subdominant systematic which is the leakage of the brighter E-modes into the faint B-modes ($E \rightarrow B$ leakage) due to higher-order cross-polar response. To avoid mistaking $E \rightarrow B$ leakage due to cross-polar response with other sources that mix E and B modes such as partial sky coverage or unaccounted detector gain variations, we use the term polarization-to-polarization leakage ($P \rightarrow P$ leakage) to refer to this systematic effect ¹. This choice is more consistent with the beam simulations terminology we use internally for our beam simulations and will be used throughout the rest of this thesis.

At $\ell \sim 80$ the E-mode power is ~ 100 times brighter than the theory B-mode power for $r = 0.1$ ² as shown in Figure 1.5. To quantify this leakage we generalize our beam simulations

¹In the past, the terms $E \rightarrow B$ leakage simulations have been used to refer to simulating how gain variations across the focal plane can cause some E-mode signal to leak into B-modes [55], hence the choice of a new term that inevitably will irritate some people.

²Temperature is a factor of ~ 100 above the E-modes, so it is clear why $E \rightarrow B$ leakage is subdominant to $T \rightarrow P$.

framework to include the true polarized detector beams as measured during polarized far-field beam mapping (section 4.2). In this chapter, we describe the framework for $P \rightarrow P$ leakage simulations and introduce a complete description of how to simulate azimuthally asymmetric beam systematics on the curved sky without approximations. In addition, we use beam maps from polarized FFBM to model elnod gain systematics related to Dk rotation in BICEP3 CMB data. From this point on, visually smooth colorscales are used (Python’s viridis or MATLAB’s parula) to display beams and related data products. Traditionally, ‘jet’ is used to display beams which is known to have a visually sharp transition from green to blue ³.

5.1 Polarization-to-Polarization Leakage Simulations

In our temperature-to-polarization leakage simulations (section 4.1.3) we implicitly make the assumption that the polarized beam response of the detectors is the same as the unpolarized beam response (up to some polarization efficiency factor that we measure independently ⁴). This is a good approximation for $T \rightarrow P$ leakage systematics studies. For $P \rightarrow P$ leakage, however, we need to use the full polarized beam response of the detectors. The starting point for developing the machinery for $P \rightarrow P$ leakage simulations is equation 4.7 but adapted for the case that the detector is observing the CMB instead of the BSNS ⁵. Since we are interested in exploring polarization systematics, we set the temperature sky to zero, i.e. $T_{\text{sky}} = 0$, which gives the following equation for the average power measured by a detector:

$$d(\alpha) = \cos(2\alpha) \left[B_Q * Q_{\text{sky}} + B_U * U_{\text{sky}} \right] + \sin(2\alpha) \left[B_Q * U_{\text{sky}} - B_U * Q_{\text{sky}} \right]. \quad (5.1)$$

where α is the angle of observation of the detector on the sky, from equation 2.8, and B_Q and B_U are the beam responses of the detector as defined in equation 4.8. We can think of the mapmaking process as accumulating the observed signal over many sky observation angles α

³Unless you are the beams guru (Kirit S. Karkare) with vision ‘perfectly calibrated’ to the jet colorscale.

⁴Recent results constrain polarization efficiency for BICEP3 to be roughly 99.7% [40].

⁵Dropping the delta function and using $*$ to denote convolution integral.

and projecting it onto the sky Stokes basis:

$$Q_{\text{obs}} \propto \langle d(\alpha) \cos(2\alpha) \rangle \quad (5.2)$$

$$= B_Q * Q_{\text{sky}} + B_U * U_{\text{sky}}, \quad (5.3)$$

$$U_{\text{obs}} \propto \langle d(\alpha) \sin(2\alpha) \rangle \quad (5.4)$$

$$= B_Q * U_{\text{sky}} - B_U * Q_{\text{sky}} \quad (5.5)$$

where the brackets denote the average over all observation angles α . Comparison with equation 4.2 (after setting $T_{\text{sky}} = 0$ here as well) allows us to read off the convolution kernels for polarization in terms of our measured calibration maps. The result is:

$$\begin{bmatrix} Q_{\text{obs}} \\ U_{\text{obs}} \end{bmatrix} = \begin{bmatrix} B_Q & B_U \\ -B_U & B_Q \end{bmatrix} * \begin{bmatrix} Q_{\text{sky}} \\ U_{\text{sky}} \end{bmatrix}. \quad (5.6)$$

meaning that $b_{QQ} = b_{UU} = B_Q$ and $b_{QU} = -b_{UQ} = B_U$. We measure B_Q and B_U for each detector from the polarized FFBM campaigns described in section 4.2. No uniform polarization efficiency factor is applied to equation 5.1 since \hat{B}_Q is normalized to the peak unpolarized response (re-introducing hats to denote measured quantities). Therefore, all polarization efficiency information for a given detector is baked into \hat{B}_Q .

With knowledge of the convolution kernels we can perform a back-of-the-napkin calculation, to get a rough estimate of the expected level of $P \rightarrow P$ leakage by comparing the maximum and minimum values of the per-detector B_U maps which would give us an upper bound on the leakage. Figure 5.1 shows this comparison. The maximum/minimum values of the B_U maps are $\pm 0.8\%$ of the peak beam response, which would naively translate to a $P \rightarrow P$ leakage of $\sim 6.5 \times 10^{-5}$. However, the actual leakage is expected to be smaller due to cancellation effects when coadding over many Dk observations and averaging over multiple pairs of detectors. This upper bound is already a factor of 20 smaller than the estimated $T \rightarrow P$ leakage of

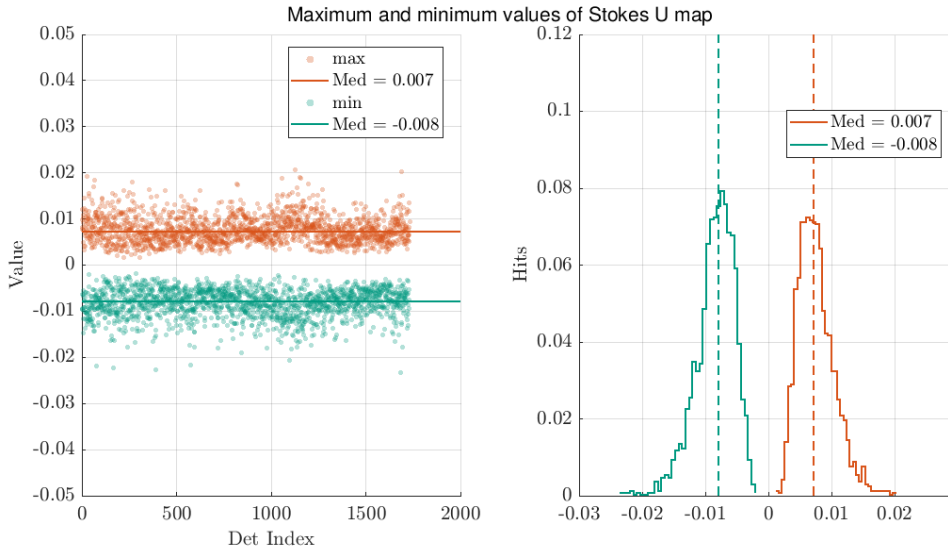


Figure 5.1. Maximum and Minimum values of B_U maps measuring the cross-polar response of BICEP3 detectors. The two distributions are well separated and symmetric around zero, hinting at no net monopole response. The max/min are $\pm 0.8\%$ of the peak beam response, which would naively translate to a $P \rightarrow P$ leakage of $\sim 6.5 \times 10^{-5}$. However, the actual leakage is expected to be smaller due to cancellation effects when coadding over many Dk observations and averaging over multiple pairs of detectors.

BICEP3 in BK18. The remainder of this section is split into two parts. The first gives a high-level overview of the timestream simulation process for $P \rightarrow P$ leakage. The second part describes the technical details on the practical implementation and challenges of folding the subdominant cross-polar response into the simulated TODs, which is general enough to handle any azimuthally asymmetric beam.

Simulation Overview

With *a priori* intuition of the amplitude of the effect we can use equation 5.1 to make simulated timestreams for each BICEP3 detector for the 2018 observing season after we specify an appropriate input sky. We generate constrained, unsmoothed, and lensed Λ CDM realizations in the HEALPix pixelization of $N_{\text{side}} = 512$ using lensing information from SPT3G [78]. The map information is contained in the spherical harmonic coefficients $a_{\ell m}^T$, $a_{\ell m}^E$, and $a_{\ell m}^B$, that we massage to explicitly set $a_{\ell m}^B = 0$. Therefore, the resulting input Q_{sky} and U_{sky} maps are lensed, but contain no B-mode signal. At this stage the T_{sky} map is set to zero. The simulated Q_{sky} and U_{sky} maps are shown in Figure 5.2.

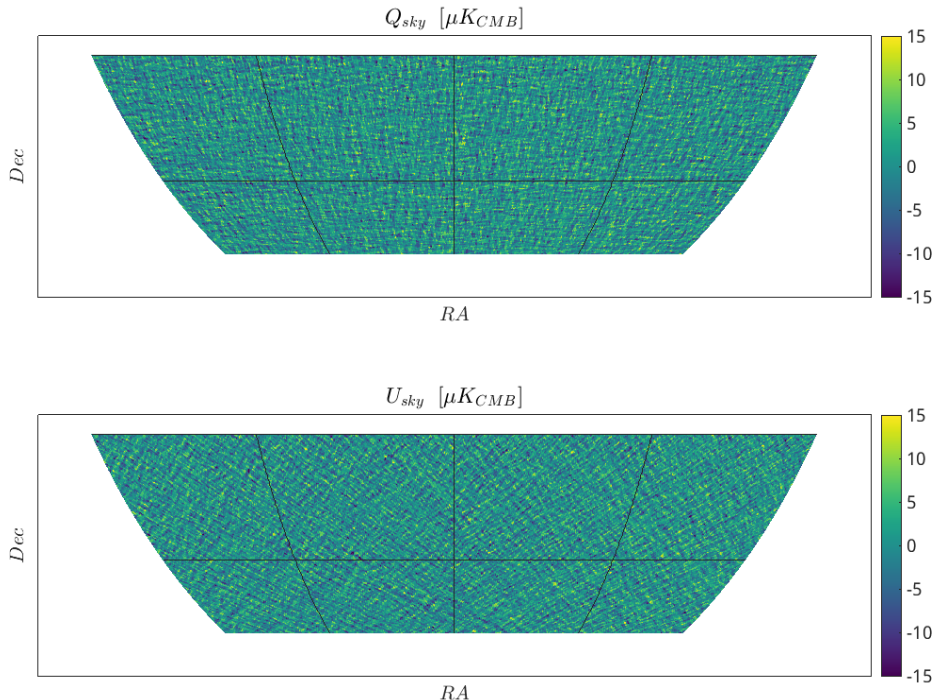


Figure 5.2. Single realization of constrained, no B-mode, lensed, and unsmoothed CMB maps in HEALPix pixelization of $N_{\text{side}} = 512$ for the BK field. The maps are in units of μK_{CMB} and are the input to the $P \rightarrow P$ leakage simulations. The T_{sky} map is set to zero. Q_{sky} (top) and U_{sky} (bottom).

The per-detector simulated timestreams are then fed into the simulations pipeline described in section 2.7.5 and undergo the same filtering, deprojection, and mapmaking process as the real CMB data. Contrary to the $T \rightarrow P$ leakage simulations where any polarization power in the output $Q_{\text{sky}}/U_{\text{sky}}$ would be due to leakage, here we are looking for excess B-mode power in the output sky that is due to cross-polar response described by the \hat{B}_U convolution kernel. For this reason, we make a second set of simulated timestreams where we explicitly set $\hat{B}_U = 0$, which also operates as a test case for removing imperfect purification residuals⁶. The difference of the resulting output Q/U skies (referenced as the *cross-polar leakage map*) would yield the leakage B modes due to higher-order cross-polar response. The cross-polar leakage map is then used in power spectrum estimation that applies the BK18-era, single year, bandpower window functions as described in section 2.7.5.

⁶Ambiguous modes ending up in the BB power spectrum due to imperfect separation of the E and B subspaces.

Practical Implementation and Challenges

In practice, this process has a few complications. Our beams are in local, Cartesian coordinates (x', y') which is the main driver for producing an intermediate flat sky map around the detector's trajectory to carry out the beam convolution in $T \rightarrow P$ leakage simulations. As discussed in section 4.1.3, this process introduces the algorithmic floor (nominally removed from the final $T \rightarrow P$ leakage spectra) of $\mathcal{O}(10^{-4})$. This amount of injected B-mode leakage is already a factor of 2 larger than the upper limit of the expected $P \rightarrow P$ leakage. As a test, the $P \rightarrow P$ leakage simulations were carried out using the intermediate flat sky for the beam convolution step. Unsurprisingly, the algorithmic floor completely drowned the leakage we are trying to estimate, and resulted in nonphysical spectra. The solution is to dodge the use of the intermediate flat map entirely and perform the convolution on the part of the sphere around the detector's trajectory.

For a general beam shape, the convolution on the sphere of the sky map and a beam map (rotated to the detector's trajectory and to match the Dk angle of observation) at each point t along the scanning trajectory takes the form:

$$d(t) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \sum_{m'=-\ell}^{\ell} a_{\ell m} D_{mm'}^{\ell}(R) b_{\ell m'} \quad (5.7)$$

where $a_{\ell m}$ are the spherical harmonic coefficients of the sky, $D_{mm'}^{\ell}(R)$ is the Wigner D-matrix that encodes how the spherical harmonic coefficients of the beam $b_{\ell m'}$ are rotated by parallel transport R (defined in equation 2.2) to the corresponding sky coordinates at each point along the scan trajectory.

It is clear from equation 5.7 that for azimuthally asymmetric beams ⁷, convolution on the sphere is not as simple as multiplication in spherical harmonic space due to the complicated coupling of the m and m' spherical harmonic coefficients. In fact, the cross polar response kernel \hat{B}_U is by definition azimuthally asymmetric. To maintain generality and introduce a

⁷The Wigner D-matrices simplify dramatically for the case that $b_{\ell m'} = 0$ unless $m' = 0$ and convolution is simply multiplication in spherical harmonic space.

simulation scheme that can be used in the future to precisely model systematics that are not azimuthally symmetric (i.e. crosstalk, reflections off detector tile frames – see appendix B, polarized ghost beams – see section 5.2) we explore the full beam convolution on a sphere. The problem quickly becomes computationally expensive in both wall-clock time and computing resources when expanding the beam onto the full sky and attempting to perform operations (like rotating the beam onto the detector’s trajectory). Although simulated timestreams have been produced with this method for a handful of detector pairs, to get simulated timestreams for the ~ 1200 pairs of BICEP3 we carry out the convolution in real space⁸.

We define a point cloud of HEALPix pixels initialized at the north celestial pole with $N_{\text{side}} = 512$ to resample our beams onto the sphere out to radius equal to the beam thumbnail which is fixed at 2° ⁹. The convolution kernels will now be denoted as $\hat{B}_Q(\hat{n})$ and $\hat{B}_U(\hat{n})$ to keep track of the sky direction. An example of the point cloud overlaid onto the rectangular beam map thumbnail is shown in Figure 5.3.




Figure 5.3. Rectangular beam thumbnail (left) and point cloud locations (right) of a representative BICEP3 detector. These locations are used to resample the flat beam onto the curved sky, initialized at the North celestial pole. The beam shown has been integral normalized to unity.

To avoid aliasing we average pixels of the rectangular beam map that fall within the same HEALPix pixel. This results in a slight smearing in the radial direction along the diamond shaped pixels. However, since the beam thumbnail is small this effect is not dominant.

⁸Real space and pixel space are used equivalently.

⁹Larger thumbnails could be used, but compute time scales with number of points.

Once we calculate each detector’s trajectory on the curved sky (using the pointing model – section 2.6) during a certain CMB observation (a tag) we rotate the point cloud corresponding to the convolution kernels $\hat{B}_Q(\hat{n})/\hat{B}_U(\hat{n})$ to the corresponding sky coordinates while matching the Dk angle of observation. This is done using parallel transport, which in the language of equation 2.2 takes the form: $R(RA, \pi/2 - Dec, Dk)$. With the detector’s trajectory in hand, and the point cloud of the beam in the same pixelization, the final ingredient is to interpolate the input map ($Q_{\text{sky}}/U_{\text{sky}}$) along that trajectory. The interpolated sky maps will be denoted as $Q'_{\text{sky}}/U'_{\text{sky}}$. For our CMB simulations this is done using Taylor series interpolation off the input map and its derivatives up to second order. This is possible because these sims are pre-smoothed with the array averaged B_ℓ and the map derivatives are pre-computed. In this case, in order to use the same interpolation scheme we must calculate the covariant derivatives of the point cloud and how it couples to the derivatives of the input map for each detector. On a curved manifold like \mathbb{S}^2 affine connection information becomes important for the second order derivative terms (first derivative is still the gradient). Since the second order derivatives are cumbersome, and we have to calculate two of them for each detector (one for $\hat{B}_Q(\hat{n})$ and one for $\hat{B}_U(\hat{n})$), for each point along the scanning trajectory, we opt to use nearest neighbor interpolation. Comparison between nearest neighbor and Taylor interpolation up to first order yields near identical results, so we stick to nearest neighbor as it reduces compute time by a factor of ~ 3.5 ¹⁰. The resulting point cloud corresponding to the $\hat{B}_Q(\hat{n})$ convolution kernel of a single detector is plotted in Figure 5.4.

At this point, we have all the ingredients to carry out real space convolution on the sphere using the following expression:

$$d_t = \sum_p \left[R^{-1} \hat{B}_Q(\hat{n}_{t,p}) Q'_p + R^{-1} \hat{B}_U(\hat{n}_{t,p}) U'_p \right] \Delta\Omega_p \quad (5.8)$$

where $\hat{n}_{t,p}$ denotes the sky direction at time t and pixel p along the detector’s trajectory and $\Delta\Omega_p$ is the solid angle of the HEALPix pixel. We generate two sets of these simulated

¹⁰The computation is vectorized and parallelized to be executed by 20 CPU cores. This factor would be much higher otherwise.




Figure 5.4. Representative point cloud from a single BICEP3 detector on the HEALPix grid, plotted in gnomonic view. As plotted the beam has not been rotated to the detector trajectory yet, but has been rotated about the North celestial pole direction to match the DK orientation of the particular CMB observation. It is also normalized so that the integral over the entire point cloud equals unity.

timestreams, one using the measured $\hat{B}_U(\hat{n})$ and one for the test case of $\hat{B}_U(\hat{n}) = 0$. Post mapmaking, and the pairmap coaddition process, we take the difference of the resulting Q/U skies to get the cross-polar leakage map. Power spectrum estimation on the cross-polar leakage map yields the $P \rightarrow P$ leakage spectra plotted in Figure 5.5. The ρ estimator is also calculated for the resulting BB spectrum and is found to be $\rho = 2.61 \times 10^{-5}$, which as expected, is a factor of ~ 2.5 smaller than the upper bound estimated from the maximum value of the \hat{B}_U maps.




Figure 5.5. Final $P \rightarrow P$ leakage spectra and ρ estimate. Overplotted are the theory D_ℓ^{BB} for $r = 0.1$ and $r = 0.1 +$ lensing for reference. As expected, the $P \rightarrow P$ leakage due to cross-polar response is roughly a factor of 10 below the BK18 $T \rightarrow P$ leakage and a factor of ~ 2.5 smaller than the upper bound estimated from the maximum value of the \hat{B}_U maps.

Due to how small this effect is, we do not propagate the leakage spectra through the multi-component likelihood analysis to get a bias on r since it is expected to be negligible. However, the introduction of this framework opens the door for other systematics studies that involve azimuthally asymmetric beam features seen in beam maps, that have potentially larger impact on our measurements. Precise modeling of such features would inform analysis of real CMB data and motivate us to extend our deprojection scheme to include more exotic forms of beam features (e.g. localized ghost beam deprojection).

5.2 Gain Systematics Due to Boresight Rotation

Perhaps the most underrated use of calibration beam maps (both unpolarized and polarized) is to carry out exploratory studies in order to explain real instrument behavior. As discussed in section 2.7.1 there are several data quality metrics related to the elnods that are performed at the beginning and end of each ~ 1 hour CMB observation (e.g. elnod goodness of fit, difference between leading and trailing elnod responses, etc.). Dating back to BICEP3's deployment in 2016, a dependence of the elnod response as a function of Dk angle was observed, visualized in Figure 2.15. Although the elnod gain dependence was first thought to vary in Dk , when plotted as a function of elevation (or Declination) it becomes clear that it is primarily a dependence with elevation. This can be seen in Figure 5.6 where the elnod derived detector gains extracted from fitting to a $csc(El)$ model of airmass are plotted at a given Dk angle divided by the mean elnod gain across all Dk angles for that detector.




Figure 5.6. BICEP3 per-detector mean elnod gain at a given Dk angle normalized across the entire 2018 observing season, plotted by projecting the focal plane onto the curved sky at $El = 55^\circ$. Δel denotes the pointing offset of an individual detector with respect to the boresight in the declination direction. At the South Pole $El \propto -Dec$, thus the negative y -axis corresponds to higher elevations. A $\pm 5\%$ gradient in gain is seen as a function of Dec across multiple Dk observations. The observation is that the elnod gains of detectors pointing at lower elevations than the boresight have systematically lower elnod gains while those pointing at higher elevations have systematically higher elnod gain values.

The observation can be summarized in the following 3 points. For a given observation at boresight angle Dk :

- Detectors pointing at lower elevations than the boresight (higher Dec) have systematically lower elnod gains.

- Detectors pointing at higher elevations than the boresight (lower Dec) have systematically higher elnod gains.
- Detectors pointing at the same elevation with the boresight have elnod gains roughly consistent with those at all other boresight measurements.

This effect, first discovered in BICEP3, is present in both the BICEP2 and Keck Array experiments, but is less pronounced. If not accounted for, gain variations across detectors in the focal plane scanning across different parts of the CMB sky will lead to E-to-B mixing¹¹. Although our relative gain calibration procedure (see section 2.7.1) is constantly evolving in the reanalysis, we mitigate boresight-dependent gain variations by normalizing the elnod gain for each Dk angle and each detector prior to calibrating the timestreams using equation 2.10. The price for this choice is loss of any true boresight-dependent gain variations due to detector pathologies or optical effects, but no evidence of such effects has been observed across any of our calibration datasets up to the writing of this thesis. In addition, our absolute calculation procedure (section 2.7.4) relies on per-tile, per- Dk angle coadded maps over a full season, in part due to the observed Dk dependence of the elnod gains discussed here.

The culprit for this behavior was hypothesized to be due to extended beam response and an excursion from the nominal $csc(El)$ model of the atmosphere that is used to calibrate the elnodes. The $csc(El)$ model assumes that the atmosphere is an optically thin, plane parallel slab, which is a good approximation for the extremely dry and stable atmosphere at the South Pole. However, sidelobe response and bandpass coupling of the detectors to an opaque part of the atmospheric spectrum could lead to deviations from this simple model. This hypothesis was explored in Chapter 3 of [79] by postulating a 1.5% level of sidelobe response, with 1% of it spread uniformly at a radius of 15° centered at the boresight and the remaining 0.5% concentrated in a compact ghost beam (see section 5.2.1). Although localized ghost beams have been observed at a 0.3% level, no diffuse sidelobe response has been observed in any calibration datasets up to the writing of [79]. In this section we use unpolarized BSNS

¹¹We further reduce concerns of E-to-B mixing due to gain variations and systematic polarization angle errors via self-calibration [55].

beam maps (pair-sum of the data products in section 4.2) to extend and test this hypothesis with the deepest measurement of the near sidelobes to date.

5.2.1 Ghost Beams

Enough beam power in the sidelobes would couple to different parts of the atmosphere and could result in the elnod gain pattern with Dk we observe. A known feature in optical systems is *ghost beams*. Generally, they are secondary beams, that could be compact or diffuse, that arise from internal reflections off an optical element that could eventually couple to a different part of the sky. Depending on where along the optics chain the reflection occurs will determine if ghost beams will be refocused (and therefore compact) or diffuse, or both. Colloquially, we refer to localized ghost beams as ‘little buddies’ or ‘buddy beams’. We use anti-reflection (AR) coatings on both surfaces of each optical element to minimize this effect, but small levels of ghosting are inevitable. BICEP3 has a known reflection off the 50K metal mesh filter that results in buddy beams that can be seen in beam maps. In the reverse time sense, each detector’s pointing ray couples to a certain sky direction. A reflection off the 50K filter then travels back to the focal plane, reflects again, and travels through the optics, exiting the receiver at a position on the sky that is exactly opposite to the boresight from where it originated. A sketch of how this reflection occurs is shown in Figure 5.7.

Buddy beams are a known feature seen in far-field beam maps of instruments dating back to BICEP2 and Keck Array. Since they are a real optical feature, they can be seen in all FFBM datasets wide enough to capture the near sidelobes. This includes polarized and unpolarized FFBM datasets that use the thermal chopper and the BSNS. Composite beam maps of a handful of BICEP3 detectors from polarized FFBM that uses the BSNS source are plotted in Figure 5.8 to show an example of buddy beams. These beam maps are located in strategic locations across the focal plane to illustrate the buddy beam phenomenology.

A wealth of beam features can be seen in these maps including the square features at $\sim -45\text{dB}$ that trace the outline of BICEP3’s tile frames and are localized at their physical locations in




Figure 5.7. Sketch of how localized ghost beams arise from internal reflections off BICEP3’s 50K metal mesh filter located above the alumina lenses in the skyward direction. A detector’s pointing ray (purple) couples to a certain sky direction. The reflection off the 50K filter (red) then travels back to the focal plane, reflects again, and travels through the optics, exiting the receiver at a position on the sky that is exactly opposite to the boresight from where it originated. Since the reflection off the 50K filter happens after traveling through both lenses, the resulting ghost beam on sky is compact. In terms of focal plane coordinates (introduced in section 3.4), we expect the localized ghost beam to be located at $(-X, -Y)$ if the main beam is located at (X, Y) . The ghost beam adds power to the sidelobes in the direction opposite of the detector’s pointing, which could couple to different parts of the atmosphere and lead to the elnod gain dependence we observe. We refer to the compact ghost beam as a “buddy beam”.




Figure 5.8. Example buddy beams of various BICEP3 detectors from various tiles on the FPU plotted in x', y' coordinates. The box in the top left corner shows the relative location of the main beam of each detector in the focal plane. All maps are plotted in dB using the same color range. The individual beam maps are from the polarized FFBM dataset that uses the BSNS and correspond to detectors from tiles 2 (right), 9 (top), 11 (center), 13 (bottom), 20 (left) ¹³. In this coordinate system, the boresight is located halfway between the main beam and the buddy beam (circled in magenta). The buddy beam for most center detectors coincides with the main beam, hence a detector towards the edge of Tile 11 is chosen to showcase how buddy beams appear on center detectors. These beam maps have been composited over all calibration measurements during the 2018 calibration season (13 measurements total). A wealth of beam features can be seen in these maps, which motivates the use of the BSNS dataset to characterize the ghost beams and their contribution to the elnod gain dependence. The ‘fishtail’ features seen around -35 dB are only present for detectors near the edge of the focal plane and are due to part of the beam power missing the field lens (see Chapter 7 of [51] for more details).

the focal plane¹⁴. A brighter ‘fishtail’ feature can be seen at $\sim -35\text{dB}$ for detectors near the edge of the focal plane, which is due to part of the beam power missing the field lens verified by simulations using the GRASP optical modeling software [51].

We expect ghost beams to have little impact in our $T \rightarrow P$ leakage systematics since they arise in the same way for both A and B detectors. However, characterizing them for each detector is important to verify this intuition. The same 7-parameter elliptical Gaussian fit used to extract beam parameters from beam maps is applied to the ghost beams. The fitted amplitude of the buddy beams is found to be $\sim 0.3\%$ of the main beam amplitude as shown in Figure 5.9, which is much lower than the $\sim 1.5\%$ level of sidelobe response postulated in [79]. However, the total ghost beam power in these beam maps is not fully localized in the buddy beam. There is a diffuse component around the buddy beam that is not captured by the Gaussian fit.

To quantify the contribution of the localized + diffuse ghost beam component we define the *buddy region* as the region centered on the buddy beam with a radius of 8° . If the buddy beam region overlaps with the inner 6° of the beam map (where most of the beam power is located), we exclude it from the buddy beam region.

The buddy region varies in size across detectors depending on how close to the main beam the buddy beam is located. Therefore, it gets divided by the area of the pixels in the integration region. It is found to be $\sim 1.1\%$ on average of the main beam power as can be seen in the tile plot in Figure 5.10, except Tile 1 detectors which are $\sim 2\%$.

If we wanted to decouple the contribution of the localized buddy beam and the diffuse component we could do so by median-coadding over all detectors in a tile. The localized buddy beam will average away in these coadds. However, the diffuse ghost beam component (which is seen by all detectors in the tile) will constructively add. Integrating these coadds over the buddy region for each tile, yields the contribution of the diffuse ghost beam of $\sim 0.8\%$. Therefore, the total ghost beam power of $\sim 1.1\%$ of BICEP3 has a 0.3% contribution from localized ghost beams and a 0.8% contribution from diffuse ghost beams. Interestingly, Tile 1

¹⁴These features get averaged out when coadding over BICEP3’s detectors and are not visible in Figure 4.12.




Figure 5.9. Fitted amplitudes of buddy beams for all BICEP3 detectors. Note the systematically higher buddy beam amplitudes for Tile 1 detectors. It is associated with the uninstalled tile directly across from it, resulting in the re-reflection off the focal plane being higher than for other tiles. Therefore, the measured ghost beam amplitudes are higher for this Tile. In addition, although buddy beams in general do not show up in the pair-differenced beam maps, buddy beams from Tile 1 detectors are systematically brighter for V detectors than H detectors. This deviation compared to other tiles has led to excluding Tile 1 detectors from CMB analysis. Tiles 14, 15, 16, 19, and 20 have some erroneous buddy beam fits that correspond to non-functioning detectors that sneak past the data quality cuts.

does not stand out as having a higher diffuse ghost beam contribution compared to other tiles, which suggests that the mechanism of the diffuse ghost beam component is different from the localized buddy beam. This investigation has updated our understanding of sidelobe power due to ghost beams for the BICEP3 instrument and serves as the benchmark for BICEP Array ghost beam characterization.

Going back to the original question of whether these ghost beams can explain the elnod gain dependence we observe, the measured sidelobe response of $\sim 1.1\%$ is still below the $\sim 1.5\%$ level of sidelobe response postulated in [79]. However, the morphology of the measured ghost beam is different from the simple model of a uniform ring of sidelobe response at a radius of 15° centered at the boresight. Specifically, it is located exactly opposite to the main beam with respect to the boresight which has the following implications:

During an elnod at a given boresight angle Dk :

- A detector located higher in elevation than the boresight (whose ghost beam would therefore lie below the boresight—position A in Figure 5.8) would see a slightly in-

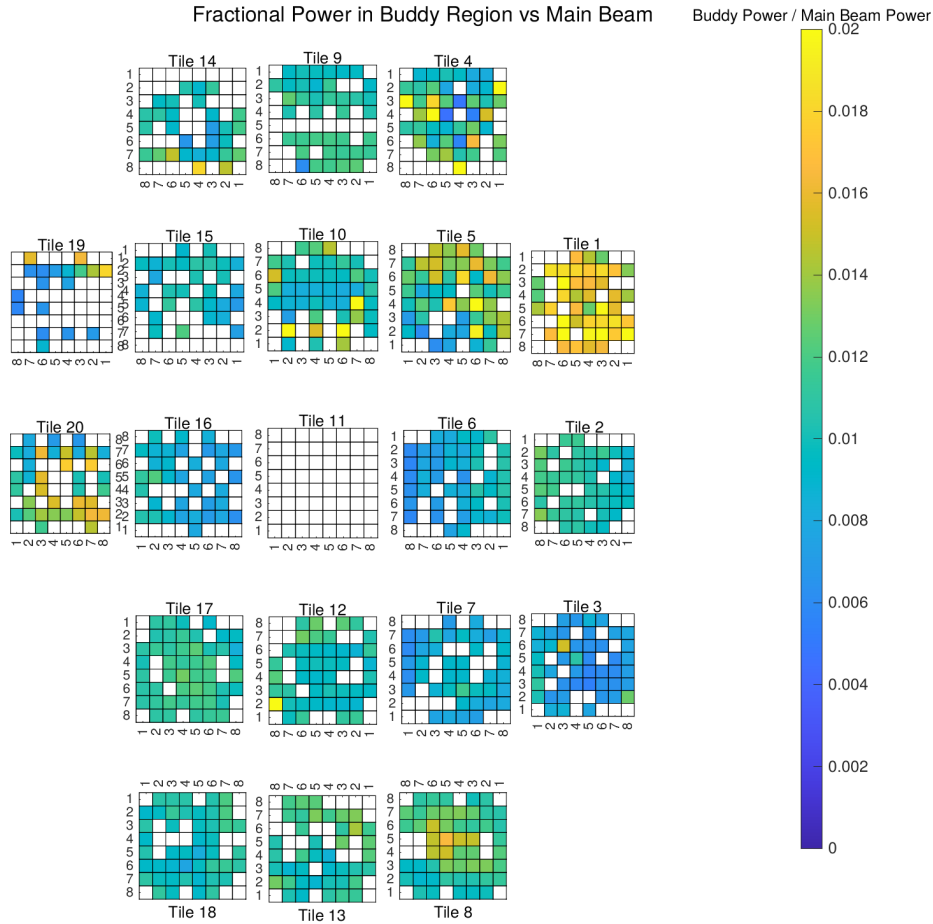


Figure 5.10. Integral of the buddy region as a fraction of the main beam power (see text for definition) quantifying the diffuse ghost beam amplitude in BICEP3 detectors. The resulting integral of the buddy region is $\sim 1.1\%$ of the main beam power on average, except Tile 1 detectors which are $\sim 2\%$. The integral is not calculated for Tile 11 detectors since the buddy region overlaps with main beam.

creased atmospheric response, due to the ghost beam coupling to a brighter part of the atmosphere. **The resulting elnod gain would be overestimated.**

- Conversely, a detector located lower in elevation than the boresight (whose ghost beam would therefore lie above the boresight—position E in Figure 5.8) would see a slightly decreased atmospheric response, due to the ghost beam coupling to a dimmer part of the atmosphere. **The resulting elnod gain would be underestimated.**
- A detector located at the same elevation as the boresight (whose ghost beam would therefore lie at the same elevation as the boresight—position B, C, and D in Figure 5.8) would see an atmospheric response that is biased neither higher nor lower, since the

ghost beam is coupling to the same part of the atmosphere as the main beam but at a different azimuth. The resulting elnod gain would be roughly consistent across different elnod observations at the same Dk angle.

Consider the case of a single detector near the edge of the focal plane (for example, a detector in Tile 20): even though the boresight is pointed at elevation of 55 degrees, the ghost beam can be pointing at elevations as low as 30 degrees. Since the atmosphere is brighter at lower elevations, the ghost beam would bias the timestreams and the fitted elnod gain would be overestimated. Conversely, when the boresight is rotated by 180° , the detector is now pointing at a higher elevation resulting in a fitted elnod gain that is underestimated. If the boresight is rotated by 90° or 270° the ghost beam is pointing at the same elevation as the detector and the elnod gain is mostly unbiased.

This phenomenology qualitatively matches the elnod gain pattern observed in Figure 5.6, and also explains why this effect is less pronounced for receivers with a smaller footprint on sky (BICEP2/Keck Array). However, to quantitatively explain this observed behavior we need a model that takes into account the ghost beam power, its shape and location on the sky as well as a more accurate model of the atmosphere than the $csc(El)$ approximation. This model would allow us to simulate elnods and extract elnod gains for different boresight angles and compare to the observed pattern, which is what we turn to next.

5.2.2 Atmospheric Modeling

We model atmospheric emission using the AM atmospheric radiative transfer code [34]. AM performs layer-by-layer calculations of molecular absorption and continuum processes across radio to far-infrared frequencies, given vertical profiles of temperature, pressure, and constituent abundances. It outputs frequency-dependent opacity, sky brightness, and transmission, and is widely used for site characterization and calibration estimates of ground and airborne experiments. For the atmospheric configuration at the South Pole, we used data from the Modern-Era Retrospective Analysis for Research and Applications, Version 2

(MERRA-2) reanalysis [80]. These meteorological profiles of temperature, pressure, and humidity were used as inputs to the **AM** atmospheric radiative transfer code, specifically during the months of deep Austral winter (June, July, August). The Planck formula for blackbody radiation (equation 1.31) at a specific frequency and temperature in the Rayleigh–Jeans limit ($h\nu \ll kT$) is given by:

$$B_\nu(T) \approx \frac{2kT\nu^2}{c^2}. \quad (5.9)$$

We can then define an equivalent *Rayleigh–Jeans brightness temperature* for a given intensity I_ν as:

$$T_{\text{RJ}} \equiv \frac{c^2}{2k\nu^2} I_\nu. \quad (5.10)$$

Figure 5.11 shows the output temperature spectrum $A(El, \nu)$ as a function of elevation, in Kelvin Rayleigh-Jeans units [K_{RJ}]. The elevation from sea-level is set to 2.835 km. In the same figure the overall BICEP3 bandpass $b(\nu)$ is shown for reference.




Figure 5.11. **Left:** Atmospheric brightness temperature spectra at the South Pole for the average PWV during the deep winter months (June, July, August). **Right:** BICEP3 95 GHz overall bandpass.

The bandpass coupling to the atmosphere beyond 110 GHz is $\sim 2.5\%$. However, this part of

the atmospheric spectrum is bright and could have a non-negligible contribution to the elnod gain. With a more precise model of the atmosphere than $csc(El)$, and BICEP3's bandpass in hand, we move to developing a model that predicts the elnod gains using our calibration beam maps as inputs.

5.2.3 Elnod Gain Model

The biggest advantage of using AM is that it can accurately model the atmosphere at all elevations and account for the elevations that the atmosphere is more opaque ¹⁵.

To model how the ghost beam couples to the atmosphere for each pair, we convolve its beam map with the atmospheric emission sky model from AM and integrate over the bandpass. To perform the convolution we project the beam maps onto the curved sky at the nominal $El \approx 55^\circ$ of CMB observations, such that $\hat{B}_T(x', y') \rightarrow \hat{B}_T(Az, El)$ ¹⁶. From this point forward we refer to the beam's direction on the sky in terms of Az, El as opposed to RA, Dec with the understanding that $El = 90 - Dec$ and $Az = -RA$. We start by projecting the FPU onto the sky via the CMB pointing model (section 2.6). Since we know the Az, El encoder positions during beam mapping, we can project the entire beam map onto the sky the same way, as can be seen in Figure 5.12. Since localized ghost beams are very relevant to this analysis, the main beam locations of two detectors (Tile 1 in purple and Tile 9 in orange) are highlighted along with the expected location of their ghost beams. The beam map projection at a given Dk angle is also shown to illustrate how a detector's ghost beam could be coupling to very different parts of the atmosphere during elnods.

For each point (i) along the elnod trajectory, our model for the elnod gain is given by the following expression:

¹⁵Also known as atmospheric 'self-shielding'

¹⁶The observed detector pointings are used in section 2.6.




Figure 5.12. **Left:** Detector pointings on sky with expected ghost beam locations overlaid. **Right:** Beam map overlotted on sky for a detector at the orange location.

$$g_{i,\text{elnod}} = g_0 \sum_{\nu} \left[A(El, \nu) * \hat{B}_T(Az, El) \right] b(\nu) \quad (5.11)$$

$$= g_0 \sum_{\nu} \sum_n \left[A(n, \nu) \hat{B}_T(Az, El-n) \right] b(\nu) \Delta\Omega_n \quad (5.12)$$

where n is the index over pixels in the beam map, and $\Delta\Omega_n$ is the solid angle element of each pixel. g_0 is a normalization factor that is set to unity since we ensure $\sum_n \hat{B}_T(n) \Delta\Omega_n = 1$ prior to convolution. In practice, we extend $A(El, \nu)$ to be the same across all azimuths since the atmosphere is mostly azimuthally symmetric at the South Pole. There are a few approximations that go into this model. First, we do not measure the beam maps using different frequency sources, so the frequency dependence of the beams is assumed to be negligible (i.e. same beam response at all frequencies within the bandpass). Second, we assume that the bandpass is the same for all detectors and equal to the overall BICEP3 bandpass, which makes the model computationally tractable but introduces some uncertainty in the final elnod gain model. Therefore, this model can only hope to model the overall trend of the elnod gain dependence with Dk and not the exact gain values for each pair. Equation 5.12 is computed along the elnod trajectory that is defined as a 0.8° excursion in

elevation centered at 55° in 10 evenly spaced steps. The Dk angles we consider are: 23° , 68° , 203° , and 248° . The result is a single curve of the elnod response for each detector pair at each Dk angle that takes into account the extended beam response and uses a more accurate model of the atmosphere than $csc(El)$.

To extract the elnod gains from this model and compare the results with Figure 5.6, we repeat the simulated elnod procedure but using delta functions (‘pencil beams’) instead of the measured beam maps. This would eliminate the sum in equation 5.12 effectively assuming that all the power is at the main beam, at the detector location. We then use the $csc(El)$ model of air mass to extract the elnod response of our detectors. This process is meant to mimic the elnod gain extraction procedure used in the real data. Regressing the modeled elnod response against the nominal $csc(El)$ model responses gives us the elnod gain in units of $K_{RJ}/\text{airmass}$ for each pair and each Dk angle. The regression coefficient effectively quantifies how much more (or less) power is picked up by an individual pair due to its extended beam shape and more accurate atmosphere compared to the nominal model of a pencil beam and $csc(El)$ atmosphere. These are plotted as the 1D counterparts to Figure 5.6. The top panel shows the observed elnod gains while the bottom panel shows the modeled elnod gains. Tile 1 is highlighted in red to distinguish it from the other tiles since it is expected to have systematically higher ghost beam power and therefore a stronger dependence of the elnod gain with Dk .

Figure 5.13 shows that the model matches the overall trend of the observed elnod gains fairly well, with Tile 1 consistently showing higher (lower) elnod gains compared to the other tiles (as expected based on its location on the sky as shown in Figure 5.12), due to higher total ghost beam power compared to other tiles.

The correlation plot between observed and modeled elnod gains is plotted in Figure 5.14. As expected, the overall trend of the elnod gains is well captured in the model. However, the individual points do not lie exactly along the $x = y$ line indicating that the elnod gains are not modeled perfectly on the per-pair scale. This is due to the simplifying assumptions discussed in the previous paragraph. Specifically, the assumption of a single bandpass for all

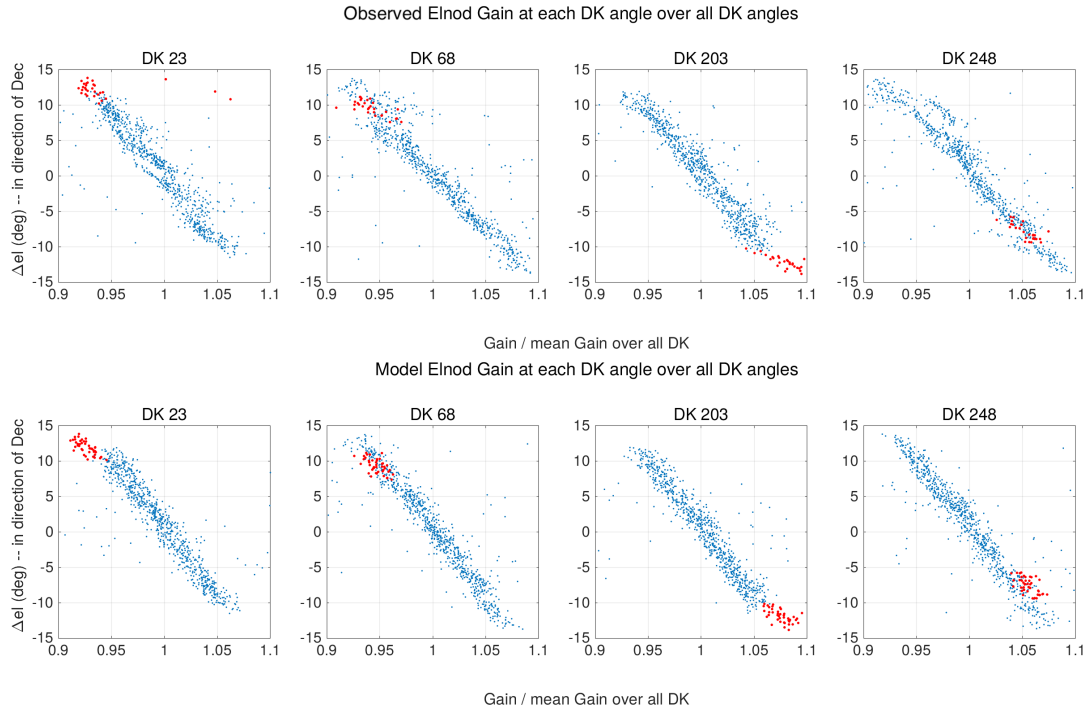


Figure 5.13. **Top:** Observed elnod gains plotted as gain/(mean of all gains). Tile 1 detectors are plotted in red to distinguish them from the other tiles since they are expected to have systematically higher ghost beam power and therefore a stronger dependence of the elnod gain with Dk . **Bottom:** Modeled elnod gains

detectors is expected to be the dominant source of the jitter seen in the correlation plot since some detectors might couple more (or less) to the opaque part of the atmospheric spectrum beyond 110 GHz or below 70 GHz.

For the first time since the discovery of the elnod gain dependence on Dk in BICEP3, we have been able to use deep measurements of the near sidelobes to explain the observed behavior with a physical model that takes into account the total ghost beam power and its coupling to the atmosphere. There is still room for improvement in the model which will allow us to better match the observed elnod gains on a per-pair level. Given all the extra analysis work needed to correct elnod gain systematics with Dk (i.e. per- Dk per-detector relative gain calibration, per-tile per- Dk abscales, self-calibration) it might be worth implementing a version of this model to reduce some of these analysis steps in the CMB data analysis pipeline. However, no such efforts are underway at the time of the writing of this thesis.




Figure 5.14. Scatter plot of the correlation between observed and modeled elnod gains. The $x = y$ line (orange) is plotted for reference. The overall trend of the elnod gains is well captured in the model, but the individual points do not lie exactly along the $x = y$ line indicating that the elnod gains are not modeled perfectly on the per-pair scale. This is due to the simplifying assumptions (see text for details), specifically the assumption of a single bandpass for all detectors which is expected to be the dominant source of the jitter seen in this plot.

Chapter 6

Ground Mapping

Detection of excess B-mode power in the CMB at degree angular scales requires exquisite control of systematic effects. In addition to the well-known instrumental systematics ($T \rightarrow P$ leakage, cross-polar leakage, detector gain systematics), ground-based CMB experiments could also be affected by systematics related to contamination from terrestrial sources. To mitigate the impact of these systematics, we design our instruments with baffles and shielding (blackened, comoving forebaffles, and stationary groundshields) such that each ray from surrounding microwave emitters must reflect at least twice prior to coupling to the detectors. In section 2.7, we described the application of the azimuth-synchronous signal subtraction technique used to remove any signal that is fixed in azimuth (i.e. synchronous with the ground) at the timescale of a single tag (~ 1 hour). However, any coupling to azimuth-synchronous signals that are fluctuating on timescales shorter than a single tag, would escape this filtering process and could potentially contaminate the CMB maps.

In this chapter, we make 360° direct measurements of the ground for both BICEP3 and BICEP Array experiments using the mirror to redirect the beams towards the ground. We focus more on BICEP Array measurements in its 2023 observing season configuration (BA30/40, BA150, K220, and K270 installed), due to the limited ground coverage of the BICEP3 having a taller ground shield¹. To build some intuition and orient the reader to the South Pole

¹Even though we tilt the BICEP3 groundshield for all ground measurements in the direction of MAPO, the coverage of the ground is still poor. Combined with SPT obscuring most of the field of view, the coverage

environment, Figure 6.1 shows a panoramic photo of the ground taken from the roof of the MAPO building that houses the BICEP Array receiver.




Figure 6.1. Panoramic photo of the South Pole environment taken from the roof of the MAPO building during late February 2023. BICEP Array, located at MAPO, is located to the left of the photo (within the groundshield). The MAPO mast with the 24' thermal chopper mounted on top is raised and is visible next to BICEP Array. In the distance, on the right of the photo, we can see the Dark Sector Laboratory (DSL) where BICEP3 and the South Pole Telescope (SPT) are located. Further, in the distance, beyond the BICEP Array groundshield, the elevated South Pole station is faintly visible. In the foreground, we can see two utility stations. This photo offers a good visual representation of what the ground maps with BICEP Array are going to look like in section 6.2.

Analysis of these measurements allows us to characterize the ground signal and its temporal fluctuations at various timescales (including sub-hourly variations) by placing a limit on the expected level of contamination in CMB power spectra as part of our overall systematics budget. We start by describing the measurement strategy and the analysis pipeline used to make *ground maps* and *ground movies*. The former is used to characterize the spatial distribution of the ground signal, as well as calibrating the data to physical units (T_{CMB}), while the latter is used to characterize the temporal fluctuations of the ground signal. Power spectrum analysis of the ground movies gives a lower bound on the *ground suppression factor* (GSF) that describes the total attenuation of the ground signal in the CMB maps (post azimuth-synchronous filtering) from our baffling and groundshield design. Finally, we characterize the total power in the ground signal at each timescale of interest and offer a

is limited.

description on how to simulate ground signal systematics in the CMB maps.

6.1 Measurement Overview

In previous years we have made indirect measurements of the ground signal by simply skipping the demodulation step in the FFBM analysis pipeline and binning in apparent azimuth and elevation. These analysis efforts did not yield conclusive results, but they inspired curiosity that ultimately led to the development of ground mapping calibration measurements. Only the redirecting mirror is needed to set up for the ground measurements. During the 2023 calibration season, after the FFBM campaign with the thermal chopper ended, we simply unmounted the thermal chopper from the DSL and MAPO masts and were ready for ground measurements with BICEP3 and BICEP Array. New schedules were written for these measurements, composed of 365° ² scans in azimuth at $2.8^\circ/\text{s}$ with an elevation step of 0.1° at the end of each azimuth scan. A schedule is composed of all the scans needed to cover the full 365° in azimuth, and the full elevation range of 30° . The setup and schedule parameters for both BICEP3 and BICEP Array are summarized in Table 6.1. The BICEP3 forebaffle is not installed for these measurements as it does not fit under the mirror, but the forebaffles of all receivers in BICEP Array were installed for these measurements.

Ground Mapping Parameters	BICEP3	BICEP Array
Groundshield Configuration	Groundshield tilted in the direction of MAPO	No groundshield tilt
Mirror	Full Coverage	Limited FPU coverage
Type of scan	Out-back scans, el step: 0.1°	Snake rasters, el step: 0.1°
<i>Dk</i> coverage	3 unique DK orientations	Single orientation for each receiver
Schedule duration	~ 22 hours	~ 12 hours
Elevation range	30°	30°

Table 6.1. Ground mapping setup and schedule parameters for BICEP3 and BICEP Array.

The key ingredient when taking these measurements is to ensure that at some point during

²Slight overlap at the beginning and end of the schedule due to the necessary acceleration and deceleration of the telescope at the beginning and end of the schedule, ensuring that the scan speed is constant throughout the azimuth range of interest.

the schedule, every detector of the receiver being mapped is coupled to both the ground and the sky. This is important for the calibration of the data to physical units (T_{RJ} and eventually T_{CMB}). This is achieved by the generous elevation coverage of the schedule, on top of the large footprint of the focal plane onto the horizon ³.

6.2 Analysis Pipeline, Ground Maps, and Ground Movies

The analysis of these measurements is quite different from the FFBM analysis described in section 3.5. For one, there is no need to demodulate the data since we are interested in the ground emission. In this section we discuss the development of the ground mapping analysis pipeline, and the resulting ground maps and ground movies.

Motivated by the enormous computational throughput of the Cannon HPC cluster, the analysis pipeline is designed to process the data per-arcfile as opposed to per-schedule. It is found that running the analysis on a per-arcfile basis and then concatenating the results is a factor of ~ 4 faster than the per-MCE analysis discussed in section 3.5.4 (assuming computational resources are available ⁴). This choice was mainly made by noticing that the Cannon job scheduler would often sneak low-memory, low-CPU, and short duration jobs into queue gaps between the long duration, high-memory, and high-CPU jobs usually requested by the group. The analysis pipeline is composed of the following steps for each arcfile:

- **Calculate pointing timestreams:** Use the FFBM pointing model (section 3.4) to calculate the pointing information and get the apparent azimuth and apparent elevation for each detector at each time sample.
- **Binning into Ground Maps:** Bin data contained in each arcfile into 0.1° square pixels in apparent azimuth and apparent elevation. Post binning, we accumulate the data from all arcfiles into a single map for each detector.

³Stage 3 BK receivers have focal planes roughly 30° across on sky.

⁴Computational resources are allocated based on group fairshare. When fairshare is high job throughput is high and vice-versa.

- **Calibration off the atmosphere:** Use the per-detector accumulated ground maps to calibrate the data to physical units (T_{RJ}) by fitting to a precise atmospheric model by AM [34]. Convert to T_{CMB} using $T_{CMB} = \frac{x^2 e^x}{(e^x - 1)^2} T_{RJ}$ where $x = \frac{h\nu}{k_B T_{CMB}}$ while integrating over the bandpass.

When calculating the pointing timestreams, we use approximate mount/mirror information that is pretty close to reality from the SolidWorks models of BICEP3 and BICEP Array. However, we calibrate the pointing information by double-checking that when the instrument is pointed at zenith, the horizon is located at apparent elevation of 0° . We also pick the apparent azimuth zero-point to be the location of the opposite building (that is DSL for BICEP Array and MAPO for BICEP3).

The choice to bin in 0.1° bins is done to maintain as much resolution as possible in the azimuth direction. The sampling rate is 200 Hz and the scan speed is $2.8^\circ/\text{s}$, so the data is sampled every $\sim 0.014^\circ$ in azimuth, which is much smaller than the 0.1° bin size and makes aliasing negligible.

Calibrating off the AM atmospheric model is essential for comparing the ground signal to the azimuth-synchronous signal that is removed from the CMB data. For this we use a configuration of AM similar to the one described in section 5.2.3, but the meteorological profiles taken from MERRA-2 (precipitable water vapor, temperature, pressure, etc.) were from the months of late Austral summer (March, April, May). In addition, we extend the frequency range of the requested atmospheric spectrum to cover all frequencies of interest (30 GHz, 40 GHz, 100 GHz, 150 GHz, 220 GHz, and 270 GHz). A plot of the atmospheric spectrum $A(El, \nu)$ that spans the full frequency range of interest is shown in Figure 6.2.

For the fit to the atmospheric model, we pick a quiet part of the atmosphere seen in the accumulated per-detector ground maps and average the signal in azimuth to get a 1D profile of the atmospheric signal as a function of elevation D_{El} . The elevation range of the fit is from 3° to 10° to avoid the very low elevations where the surrounding buildings and structures (e.g. SPT) might intrude. The cutoff at 10° is chosen to ensure that detectors located below the boresight (and therefore will only see the atmosphere towards the last 15° of the elevation




Figure 6.2. Atmospheric model for the South Pole at the time of the ground measurements calculated by AM [34] as a function of elevation. The meteorological profiles used as inputs to AM were taken from MERRA-2 reanalysis during the months of late Austral summer (March, April, May). The atmospheric spectrum is plotted across the full frequency range of interest to cover all frequencies of BICEP Array and BICEP3. Used to calibrate the per-detector ground maps to physical units (T_{RJ}).

range) are included in the fit. For an accurate fit, we also need to account for the bandpass of the detectors, which is taken to be the average bandpass of all detectors in a given frequency band from FTS measurements [81]. Then the atmosphere for a given receiver i is modeled as ⁵:

$$D_{EL,i} = g_i \sum_{\nu} b_i(\nu) A(EL, \nu) \Delta\nu + d_0 \quad (6.1)$$

The gain is extracted via χ^2 minimization between the data and the model as:

$$g_i = \frac{\sum \tilde{D}_{EL,i} \tilde{I}_{atm}}{\sum \tilde{I}_{atm}^2} \quad (6.2)$$

where the tilde denotes mean-subtracted quantities ⁶ and $I_{atm} = \sum_{\nu} b_i(\nu) A(EL, \nu) \Delta\nu$ is the band-averaged atmospheric model. The use of the average bandpass for all detectors in a given frequency band is a simplifying assumption that is expected to throw off the calibration of each detector by a small amount. The fitting result for BA150 is shown in Figure 6.3.

⁵The bandpass is normalized to sum to unity.

⁶Just good practice.




Figure 6.3. Example of fitting the atmospheric model to BA150 detectors. The 1D atmospheric profiles from the ground maps (blue) fits the atmospheric model (yellow) well enough for the purposes of this calibration. The red dotted lines denote the fitting region from 3° to 10° in elevation.

The Pearson correlation coefficient between the data and the model is calculated during the regression process and a threshold of 0.98 is set to reject detectors with correlation coefficient below this threshold⁷. This ensures that detectors that are falling off the mirror or see little to no atmosphere are excluded (e.g. those located below the boresight with respect to the increasing elevation direction). We use this threshold as a data quality cut that is applied to all data products from the ground measurements.

Polarization dependent effects play a big role here, since we observe the ground at an oblique angle. At microwave frequencies, the ground emits thermal radiation whose brightness is set by its emissivity. By Kirchhoff's law for an opaque surface, emissivity and reflectivity are related by $\varepsilon = 1 - R$, so a more reflective surface emits less. Compacted snow and ice (walking paths between buildings) are expected to have different emissive properties than fresh snow as the former is layered, and the latter is porous (porosity and layering would tend to reduce the overall resulting polarization). Considering Fresnel reflection at oblique incidence angle θ_i , and transmission angle θ_t (under Snell's law) we can think of the reflected polarization in terms of the parallel 'p', and perpendicular 's' polarization components. s is

⁷Although the threshold might seem aggressive, more than 90% of rgl detectors pass this cut for BICEP3 and more than 85% of those on-mirror for BA150.

the component perpendicular to the plane of incidence and p is the component parallel to the plane of incidence⁸. In terms of the indices of refraction of the two media n_1 for air and n_2 for snow/ice, the reflection amplitudes are given by⁹:

$$R_p = \left| \frac{n_2 \cos \theta_i - n_1 \cos \theta_t}{n_2 \cos \theta_i + n_1 \cos \theta_t} \right|^2 \quad (6.3)$$

$$R_s = \left| \frac{n_1 \cos \theta_i - n_2 \cos \theta_t}{n_1 \cos \theta_i + n_2 \cos \theta_t} \right|^2 \quad (6.4)$$

For $n_2 \sim 1.3$ and $\theta_i > 60^\circ$ we have $R_s > R_p$, so by Kirchoff's law $\varepsilon_p > \varepsilon_s$: the p -polarized emission is brighter than the s -polarized emission. Since the plane of incidence is vertical (it contains the surface normal and the line of sight), the p -polarization has its E-field in this vertical plane, while the s -polarization is horizontal. Therefore, a detector aligned with the zenith direction measures p , and one aligned with the horizon measures s . In the case of BICEP3 where we have observations of the ground at multiple DK angles, we pick a Dk angle that roughly puts the V detectors aligned with the zenith direction (parallel to p). The resulting coadded maps for the V and H detectors of BICEP3 are shown in Figure 6.4.

The horizon, cold groundshield, atmosphere, and the MAPO structure all seem largely unpolarized, and their amplitude does not change from V to H detectors. However, the ground signal is clearly brighter in the V detectors compared to the H detectors, which is consistent with the expectation that $\varepsilon_p > \varepsilon_s$: the V (p -aligned) detectors see higher emissivity and therefore a brighter ground signal. This observation is true for all frequencies of BICEP Array as well. This observation might motivate further investigation of how the polarization of the ground signal might couple to the CMB maps. The difference is relatively small ($\sim 10\%$) and as we will see in section 6.3, the unpolarized DC suppression of the ground signal is already enormous ($\sim 10^9$ on average). This topic will be revisited at the end of the chapter when discussing how to simulate ground contamination in the CMB maps to include it in

⁸The s, p designations correspond to the German *senkrecht* for perpendicular and *parallel* - interesting that only one stuck around in German.

⁹The inhomogeneous light that is reflected cannot be described by a single propagation angle, so this is a heuristic description to build some intuition.

6. Ground Mapping




Figure 6.4. Coadded ground maps for BICEP3 H (top) and V (bottom) detectors, centered in the direction of MAPO. The Dk angle is chosen such that the V detectors are roughly aligned with the zenith direction (parallel to p) and the H detectors are roughly aligned with the horizon (parallel to s).

our systematic error budget.

The resulting coadded maps for all frequency bands of interest are shown in Figures 6.5 and 6.6.

A wealth of features can be seen in these maps, especially at the higher frequencies. Although these measurements could be used for groundshield performance validation, a measurement setup error (winch cable left hanging off the groundshield, seen in the BA150 map) has prevented meaningful analysis. For the case of BICEP3, the coverage is limited due to SPT obscuring a large fraction of the field of view and the tilted groundshield blocking the ground in the direction opposite to MAPO. For this reason, the BICEP3 maps are not included in the temporal fluctuation analysis. The BICEP Array maps have excellent coverage of the ground throughout the Az range. The MAPO mast with the thermal chopper mounted on top is visible at $El \sim 20^\circ$ in the K220 and K270 maps.

At this point the units are converted to T_{CMB} using:




Figure 6.5. Coadded ground maps for BA30 (top) BA40 (middle) and BICEP3 (bottom) plotted as close to the real aspect ratio. The color scale is in T_{RJ} and varies from frequency to frequency. The BA30/40 maps can be compared to the optical panorama in Figure 6.1 to identify the various structures seen in the distance and the foreground. The lowest elevations of these maps are covered by the groundshield that is designed to redirect light to the cold sky and therefore has temperature close to the zenith temperature. Focusing on the BICEP3 map, the coverage is much more limited due to the tilted groundshield that is blocking the ground in the direction opposite to MAPO. In addition, the SPT is obscuring a large fraction of the field of view since it is located in the same building as BICEP3 (DSL). Isolating the ground in the BICEP3 maps has been difficult compared to BICEP Array, due to how close in elevation the horizon is to the groundshield cutoff. Therefore, it is not included in the temporal fluctuation analysis in section 6.4.

$$T_{\text{CMB}} = T_{\text{RJ}} \frac{N_{\text{RJ}}}{N_{\text{CMB}}} \quad (6.5)$$

for discrete bandpasses with bins i the conversion factor is given by ¹⁰:

¹⁰Assuming bandpass is normalized to sum to unity and i is evenly spaced

6. Ground Mapping




Figure 6.6. Coadded ground maps for BA150 (top), K220 (middle), and K270 (bottom) plotted as close to the real aspect ratio. The color scale is in T_{RJ} and varies from frequency to frequency, with the horizon being the brightest at higher frequencies. During the ground observations, a winch cable was left hanging off the groundshield and is causing the contamination within the groundshield seen in the BA150 map. In addition, the elevation coverage of the K220 and K270 maps is large enough to pick up the thermal emission of the thermal chopper that is located on the MAPO mast ($El \sim 20^\circ$). Since the wavelengths of these receivers are smaller compared to the previous plot, a lot more detail can be seen in the ground structure.

$$N_{RJ} = \sum_i b_i \nu_i^2, \quad (6.6)$$

$$N_{CMB} = \sum_i b_i \nu_i^2 \frac{x_i^2 e^{x_i}}{(e^{x_i} - 1)^2}, \quad \text{with } x_i = \frac{h\nu_i}{k_B T_{CMB}} \quad (6.7)$$

The coadded maps are used to: visualize the spatial distribution of the ground signal, verify the calibration procedure, determine the time-averaged power in the ground signal. However, they do not contain any information on the temporal fluctuations of the ground signal.

Ground Movies

From this point on, we focus on BICEP Array measurements since the BICEP3 maps have ground coverage limitations. To characterize the temporal fluctuations of the ground signal, we make *ground movies*. The first choice in this analysis is the timescale. The natural choice is to define our time resolution as the duration of a single halfscan (i.e. a single 360° azimuth scan) at the end of which we step in elevation by 0.1° . These take ~ 280 seconds to complete including the time spent on turnarounds. With this definition, we see that a given slice in azimuth will be observed by different detectors at different times during the scan, stressing the importance of accurate calibration procedure ¹¹.

To go about constructing the ground movies (and therefore visualize the time fluctuation of ground signal), we make *field-of-view maps*. These are instantaneous detector coadds that capture the ground temperature fluctuation after each halfscan, binned into a larger pixel size of 0.8° in elevation and 0.4° in azimuth ¹². This process starts at the timestream level. There are roughly 170 total halfscans in a BICEP Array ground mapping schedule, resulting in 170 field-of-view maps. Not all of these maps include the ground signal since at the beginning and the end of the schedule these maps mostly depict the groundshield and the atmosphere respectively. Hence, the field-of-view maps are trimmed to only include the elevation range that the ground is visible (namely $-8^\circ < El < 0^\circ$) resulting in 11 total bins.

The field-of-view maps are accumulated from an entire schedule and the right-going and left-going halfscans are separated. We form *signal maps* by coadding the right-going and left-going field-of-view maps. We also form *noise maps* by taking the difference between the right-going and left-going field-of-view maps ¹³. Due to the snake raster scan strategy, the right-going and left-going halfscans are not exactly the same since they are taken at different elevations. Therefore, we cannot expect the ground signal to subtract cleanly in the noise maps. However, since the raster elevation step is only 0.1° and the field-of-view maps are binned in 0.8° elevation bins, combined with the high SNR of the measurement (as can be

¹¹This is why we use the band averaged AM atmosphere instead of simple $csc(El)$.

¹²There have been quite a few iterations of this choice including up to 1.2° El bins.

¹³The scan-direction jackknife maps introduced in section 2.7.6.

6. Ground Mapping

seen in Figures 6.5 and 6.6), the noise maps subtract away cleanly. This is illustrated in Figure 6.7.

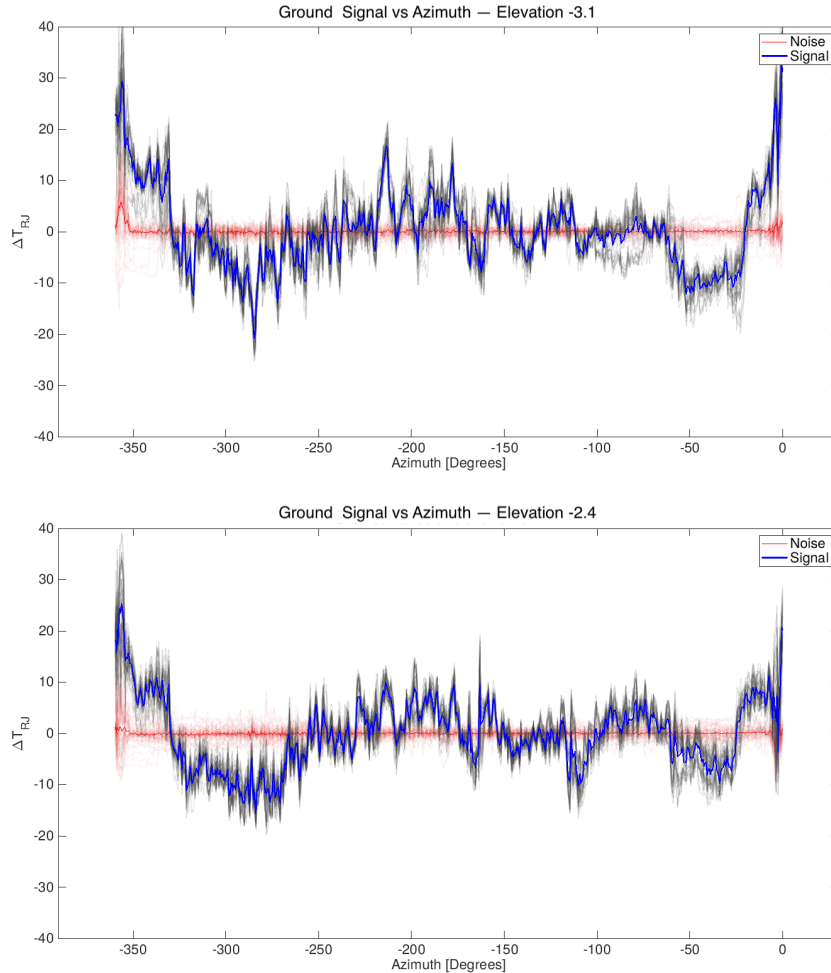


Figure 6.7. BA150 ground signal as a function of azimuth for two El bins centered at -3.1° (top) and at -2.4° (bottom). The faint lines denote the signal (black) and noise (red) for each halfscan, while the mean of ground signal (blue) and mean of noise (red) are plotted in bold. The azimuth range is shifted to put the direction of DSL and MAPO mast at the edges of the plot. These regions correspond to unusually high noise and signal and are excluded from downstream analysis.

The regions where the signal and noise maps are unusually high correspond to the directions of the MAPO mast and DSL, which are excluded from downstream analysis ¹⁴.

¹⁴In practice, only the first 2 bins ($El = 0$ and $El = -0.8$) are affected by this.

6.3 Ground Suppression Factor

For each observing BICEP Array frequency ν , the ground signal $g_\nu(Az, El, t)$ and the corresponding noise $n_\nu(Az, El, t)$ are Fourier transformed along the temporal axis (halfscan index) to get $\tilde{g}_\nu(Az, El, f)$, and $\tilde{n}_\nu(Az, El, f)$. A tapered cosine window is applied to the halfscan axis to mitigate spectral leakage, and zero padding up to the closest power of 2. The Nyquist frequency is ~ 1.8 mHz given the halfscan duration of ~ 280 seconds. The power spectra of the signal and noise are calculated as:

$$\left. \begin{aligned} G_\nu(Az, El, f) &= |\tilde{g}_\nu(Az, El, f)|^2 \\ N_\nu(Az, El, f) &= |\tilde{n}_\nu(Az, El, f)|^2 \end{aligned} \right\} \nu \in \{30, 40, 150, 220, 270\} \text{ GHz} \quad (6.8)$$

where I have absorbed the normalization factor from Parseval's theorem, and windowing into the definitions of G_ν and N_ν for simplicity. The signal and noise power spectra are averaged over the El bins that measure the ground signal the best ($-5.5^\circ < El < -0.8^\circ$). The resulting data products are $\tilde{G}_\nu(Az, f)$ and $\tilde{N}_\nu(Az, f)$.

With the signal and noise power spectra in hand, we turn to the azimuth-synchronous signal subtracted from the CMB data. In general, we do not expect that the entirety of this contamination is due to coupling to an azimuth fixed structure on the ground. In fact, we have evidence that part of the azimuth-synchronous signal is due to magnetic pickup [52] in the readout system¹⁵. Another effect expected to get caught in the azimuth-synchronous signal are things like atmospheric scan-synchronous fluctuations.

Therefore, comparison of $\tilde{G}_\nu(Az, f)|_{f=0}$ to the power spectrum of the azimuth-synchronous signal yields a lower bound on the ground suppression factor (GSF), meaning that the ground signal is suppressed by at least a factor of GSF in the CMB maps. In other words, we define the term ‘scan-synchronous’ to include all effects that are synchronous with the motion of the telescope in azimuth. With this definition, the term ‘azimuth-synchronous’ is used to

¹⁵This is based on the fact that the azimuth-synchronous signal is present in the dark detectors, and it is synchronous with the motion of the telescope.

refer to a subset of scan-synchronous effects that are due to coupling to an azimuth fixed structure on the ground.

The power spectrum analysis of the azimuth-synchronous signal is similar to that described at the beginning of this section with the main difference being the lack of relevant temporal information (tags are repeated under the same observing conditions multiple days apart). The pair-sum of the azimuth-synchronous signal $\xi_\nu(Az, Dk)$ from a single tag at boresight angle Dk that is removed from the CMB timestreams is used to compare to $\tilde{G}_\nu(Az, f)|_{f=0}$ for a given observing frequency. $\xi_\nu(Az, Dk)$ is subject to the same relative and absolute gain calibrations as the real data to be converted to T_{CMB} units (μK_{CMB}). We take the Fourier transform of $\xi_\nu(Az, Dk)$ for right-going $\tilde{\xi}_{\nu,+}(Az, Dk)$ and left-going $\tilde{\xi}_{\nu,-}(Az, Dk)$ scans separately and take the cross spectrum as:

$$\tilde{\Xi}_\nu(m, Dk) = \tilde{\xi}_{\nu,+}(m, Dk) \tilde{\xi}_{\nu,-}^*(m, Dk), \quad (6.9)$$

where m is the azimuthal mode number. The main assumption here is that the noise between the two is uncorrelated, and therefore the cross spectrum is a good estimator of the power spectrum of the azimuth-synchronous signal ¹⁶.

Due to the wide range of atmospheric conditions during the CMB observations that are known to affect the azimuth-synchronous signal, the concept of ‘good weather’ needs to be defined. Our automatic round2 cuts are perfect for this purpose (section 2.7.1), hence we adjust their thresholds to empirically define it. Our most powerful per-channel weather related cut comes from the leading and trailing elnods during a single tag observation, and represents the χ^2 goodness-of-fit test to the $csc(El)$ model of the atmosphere. During cloudy conditions, the elnods deviate substantially from the $csc(El)$ model and the χ^2 is biased high. These tags are either cut or down weighted in the analysis. We use an aggressive threshold of $\chi^2 < 5$ which is a factor of 5 more aggressive than the threshold used in the CMB analysis. The second weather related cut is a per-pair cut on a χ^2 test on the pair-difference elnod signal set to be

¹⁶The pair-difference signal was also used as a noise estimate of the azimuth-synchronous signal yielding similar results at the power spectrum level.

$\chi^2 < 2$ also a factor of 5 more aggressive than the threshold used in the CMB analysis. The final custom cut threshold we apply to define ‘good weather’ is also a per-channel cut that allows us to probe the stability of the atmosphere during the duration of the tag. It is a cut on the fractional difference in elnod gains between the leading and trailing elnods to be less than 10% (a factor of 2 more strict than the threshold used in the CMB analysis)¹⁷. Azimuth fixed signals from all BICEP Array detectors that pass round 2 and good weather cuts are accumulated over deep winter months (May, June, July, August 2023) and the corresponding $\tilde{\Xi}_\nu(m, Dk)$ are calculated (~ 200 observations). Figure 6.8 shows the resulting $\tilde{\Xi}_\nu(m, Dk)$ and $\tilde{G}_\nu(m, f)|_{f=0}$ for BA150 as an example.

Although a single Dk and a single BICEP Array receiver are shown in Figure 6.8, the noise is significantly subdominant to the time-averaged ground signal across all frequencies and all Dk angles.

When pointing at $El = 55^\circ$ which is the average elevation during CMB observations, each Dk angle puts different receivers at different heights with respect to the groundshield height. For example, BA150 at DK angle of 127° is located below BICEP Array’s boresight. Naively, we might expect that a receiver located above the boresight might couple to the ground more strongly than a receiver located below the boresight. We indirectly probe for such dependence by producing a GSF for each CMB Dk . We model the suppression factor of the ground signal as:

$$\text{GSF}_{\nu, Dk}(m) = \frac{\tilde{G}_\nu(m, f)|_{f=0} - \tilde{N}_\nu(m, f)|_{f=0}}{\tilde{\Xi}_\nu(m, Dk)} \quad (6.10)$$

The GSF is then averaged in the azimuth range where the azimuth-synchronous signal is believed to originate (i.e. angular scale of surrounding buildings), which is roughly $5 - 7^\circ$. This scale is determined by eye from examining the coadded ground maps in Figures 6.5 and 6.6. Figure 6.9 shows the resulting GSF for all BICEP Array frequencies and all CMB Dk angles.

¹⁷There is a true temperature gradient measured in pair-sum as the atmosphere changes during the tag duration, and we want to allow for that.




Figure 6.8. For BA150 as an example, plotted are the power $\tilde{G}_{150}(m, f)|_{f=0}$ (brown) and noise $\tilde{N}_{150}(m, f)|_{f=0}$ (orange) spectra of the ground signal as measured during ground mapping, with the power spectrum of the azimuth-synchronous signal $\tilde{\Xi}_{150}(m, Dk = 23^\circ)$ (purple) overplotted for comparison. The purple shaded region denotes 95% of the azimuth-synchronous signal power spectra across all ‘good weather’ tags included in this analysis (May, June, July, August 2023). A noise model of white + $1/f$ noise is fit to the noise spectrum that is also overplotted (yellow) (see appendix C). These spectra are used to calculate the ground suppression factor to get a lower bound on the coupling of the ground signal to the CMB maps. The azimuth-synchronous signal comes from our CMB observations which cover a much smaller azimuth range ($\sim 60^\circ$) compared to the ground maps ($\sim 360^\circ$).




Figure 6.9. Ground suppression factor (GSF) for all BICEP Array frequencies and all CMB Dk angles. The error bars correspond to 95% of the GSF values across all ‘good weather’ tags included in this analysis. The Dk angles in blue correspond to the drum locations that the given receiver is located at the bottom of the drum and the Dk angles in red correspond to the drum locations that the given receiver is located at the top of the drum. The GSF is at least 10^8 for all frequencies and all Dk angles, with BA30 showing systematically lower GSF compared to other BA frequencies. Keck receivers (K220, K270) are consistent with each other but lower than the values of BA40 and BA150. No dependence on Dk angle is observed.

No dependence of the GSF on Dk angle is observed, which might have been expected if the ground signal was the dominant component of the scan-synchronous signal, or if the groundshield was shorter which is reassuring.

6.4 Temporal Fluctuation

With $\tilde{G}_\nu(m, f)$ and $\tilde{N}_\nu(m, f)$ in hand, we turn to characterizing the temporal fluctuation of the ground signal. Although the power spectra of the azimuth-synchronous signal are not relevant to this part of the analysis, Figure 6.10 shows a visualization equivalent to Figure 6.8 but evaluated at $f = 0.83 \text{ mHz}$ which corresponds to a timescale of approximately 20 minutes. The ground signal is still dominant over the noise at this timescale.

The set of timescales we are interested in is: $t \in \{16, 20, 32, 40, 64, 80, 128\}$ minutes. Technically, since we observe the ground for almost 5.5 hours during a single ground mapping schedule, we could include the 160-minute timescale as well, but given that it is way above the timescale of a single tag, it is not as relevant to this analysis. For each timescale of




Figure 6.10. BA150 ground signal power spectrum $\tilde{G}_{150}(m, 0.83 \text{ MHz})$ (brown) and noise power spectrum $\tilde{N}_{150}(m, 0.83 \text{ MHz})$ (orange) at the timescale of 20 minutes. The DC azimuth fixed signal power spectrum $\tilde{\Xi}_{150}(m, Dk = 23^\circ)$ (purple) is overplotted for easy comparison to Figure 6.8, but not used in the analysis of temporal fluctuation. Even for short timescales, the ground spectra are still signal dominated across all azimuthal modes.

interest, we define a simple metric as the average of the temporal power spectrum over the azimuth range of the CMB observations ($5^\circ < Az < 55^\circ$) for each t bin. We expect the temporal fluctuation of the ground signal to look like a power law in time, with more power at longer timescales, and less power at shorter timescales. This is shown in Figure 6.11 for all BICEP Array observing frequencies.




Figure 6.11. Temporal fluctuation of the ground signal for all BICEP Array frequencies. Plotted is the average of the power spectrum $\tilde{G}_\nu(Az, f)$ over the azimuth range of the CMB observations ($5^\circ < Az < 55^\circ$) for each timescale bin.

6.5 Future Work

Having characterized both the time-averaged and time-varying components of the ground signal, the goal is to fold them into our simulations framework to get an estimate of the level of ground contamination in the CMB maps. There are two ways of going about this:

- **Directly model the measured ground contamination:** Convolve the far sidelobe

responses of our instruments with the per-detector ground maps. This approach allows us to make simulated timestreams that contain the ground signal (similar in principle to the process described in section 5.1). However, this approach is computationally expensive and data products like FSL maps do not exist for BICEP Array receivers yet. The convolution process on the curved sky for such a project would definitely require the machinery developed for the $P \rightarrow P$ leakage simulations.

- Empirically model the ground contamination that escapes the scan-synchronous subtraction process:** As we have shown, the ground temperature fluctuates at timescales below the duration of a single tag. We can use $\tilde{G}_\nu(Az, f)$ to estimate how much less power is contained in timescales of < 1 hour compared to the DC component of the ground signal. Then, we would use these factors to scale down the azimuth-synchronous signal that is removed from the CMB timestreams, and inject it into the simulated timestreams. Since the azimuth-synchronous signal contains more than just the ground signal, this approach would likely overestimate the level of ground contamination in the final CMB spectra, resulting in an upper limit on the level of ground contamination. What is motivating about this approach is that the machinery to complete this investigation already exists within our simulations framework, and that it uses data products that already exist. The resulting upper limit however, would be more representative of the scan-synchronous contamination rather than the ground contamination alone.

Going back to the polarization of the ground emission, it is unclear how it should be included in the simulations framework. The ground emission is measured to be partially polarized, but the potential coupling of the polarized component to the detectors during CMB observations would require multi-path reflections that would be difficult to model. However, using the maps shown in Figure 6.4, we can estimate the index of refraction of the ground. The process of doing this starts with the Fresnel reflection coefficients from equation 6.4, re-written to match the problem statement in terms of V and H detectors. For this particular case, V detectors are aligned with the zenith direction and H detectors are aligned with the horizon direction, to avoid carrying an overall angle that modulates the polarization dependence of

the reflection coefficients. The resulting reflection coefficients (after using Snell's law and setting $n_1 = 1$ for air) are:

$$R_H(Az, El) = \left| \frac{\cos \theta - \sqrt{n^2 - \sin^2 \theta}}{\cos \theta + \sqrt{n^2 - \sin^2 \theta}} \right|^2 \quad (6.11)$$

$$R_V(Az, El) = \left| \frac{n^2 \cos \theta - \sqrt{n^2 - \sin^2 \theta}}{n^2 \cos \theta + \sqrt{n^2 - \sin^2 \theta}} \right|^2 \quad (6.12)$$

Our ground maps represent brightness temperature maps coadded separately for V and H detectors, and are denoted by $T_V(Az, El)$, $T_H(Az, El)$. The quantity we measure when ground mapping is therefore:

$$T_{(p)}(Az, El) = \epsilon_{(p)}(Az, El) T_{\text{phys}}, \quad (6.13)$$

where $p \in \{V, H\}$ denotes polarization orientation, $\epsilon_{(p)}(Az, El)$ is the emissivity which is a function of the elevation and azimuth that we observe the ground at, and T_{phys} is the physical temperature of the snow/ice surface. The emissivity is related to the Fresnel reflection by $\epsilon_{(p)}(Az, El) = 1 - R_{(p)}(Az, El)$. Based on Figure 6.4, we conclude that $R_V(Az, El) \neq R_H(Az, El)$. It is more convenient to express the reflection coefficients in terms of the zenith angle $\theta = 90^\circ - El$ rather than El . In terms of θ , our measurements are at oblique angles of roughly $\theta \sim 80^\circ$ where the polarization dependence of the reflection coefficients is expected to be large. We express the ratio of the brightness temperatures measured by V and H detectors as:

$$\frac{T_V(Az, \theta)}{T_H(Az, \theta)} = \frac{\epsilon_V(Az, \theta)}{\epsilon_H(Az, \theta)} = \frac{1 - R_V(Az, \theta)}{1 - R_H(Az, \theta)}. \quad (6.14)$$

where the only unknown parameter is the index of refraction n , and we have assumed that the

underlying physical temperature of the ground is the same for both polarizations. To get an overall estimate of n , we average over the azimuth range such that the observed emission only depends on θ $\langle T_{(p)}(Az, \theta) \rangle_{Az} \rightarrow \langle T_{(p)}(\theta) \rangle$. For a given zenith angle θ_0 corresponding to each elevation bin, we can numerically solve for n . We can extend this process for different regions of the ground that are dominated by different media (e.g. porous snow vs compacted ice). Surface roughness can be modeled as reducing polarization contrast by mixing polarizations as:

$$R^{(p)} \rightarrow R^{(p)} \exp(-k^2 \sigma^2 \cos^2 \theta), \quad (6.15)$$

where k is the wavenumber and σ is the RMS surface height. Neglecting roughness would therefore lead to underestimating n . By a quick back-of-the-napkin calculation, ignoring roughness we get $n \approx 1.4$ which is lower than that of solid ice (roughly $n_{ice} \approx 1.78$). In addition, since ground emission at the South Pole site is layered and originates from a finite depth rather than a sharp boundary, we expect the effective index of refraction to be lower than that of solid ice, and closer to that of porous snow. Table 6.5 summarizes the expected refractive indices of the different media so that we can compare them with our extracted value of $n \approx 1.4$.

Material	n
Air	1.0
Snow (low density)	1.2–1.35
Compacted Ice (firn)	1.3–1.5
Solid ice	~ 1.78

If we wanted to get a more accurate estimate of the index of refraction, we would need to improve our model of the ground emission to include more realistic features of the snow/ice, such as:

- Include surface roughness as a free parameter in the model

- Subsurface emission should be modeled with radiative transfer equations rather than Fresnel reflection coefficients

Although these measurements provide an estimate of the index of refraction of the ground, it is not of particular interest to our analysis of the ground contamination in the CMB maps, and by extension to the systematics analyses for BK receivers.

Chapter 7

Conclusion

The emphasis of this work has been on beam characterization, beam systematics, and exploratory data analysis that deepens our understanding of our instruments. It is through a wide range of measurements (unpolarized/polarized FFBM, FSL, time-constant measurements, ground mapping etc.), that we have been able to build physically motivated models that explain instrumental systematics and their impact on our measurements. Ultimately, control of these systematics will lead to improved constraints on r and potentially a non-zero detection of primordial B-modes. As such, we emphasize that a wide variety of calibration measurements should be prioritized for future calibration campaigns, specifically for BICEP Array receivers, in order to achieve the best possible instrument performance.

As the number of detectors increases and the map depth improves, the requirements on the control of systematics become more stringent. Temperature-to-polarization leakage will continue to be the dominant systematic for the current and future level of sensitivity. Physically motivated deprojection schemes will still be the most effective way to mitigate it. Subdominant systematics such as polarization-to-polarization leakage and sidelobe contamination will start to become more relevant as well, emphasizing the need for a comprehensive understanding of the instrument's beam response. With the simulation machinery developed in this thesis, we have all the tools to directly simulate the impact of beam systematics without the flat-sky approximation. This method is the most reliable way to quantify how

systematics enter our measurements, since it uses the true polarized and unpolarized beam maps, capturing the full complexity of the beam response. Once we have the autospectrum of a given systematic, we can directly compare it to real data at the power spectrum level (internally referred to as the ‘cross-with-real’) to determine how much of the systematic is present in our data. Depending on the level of contamination, we could opt to then remove it from the spectra of the real data (+ simulations) prior to maximum likelihood estimation ¹.

In our most recent publication—BK18 [27]—we note that the bias on r is $\Delta r \sim 1.5e - 3$ and has scaled down nearly proportionally with our level of sensitivity $\sigma(r)$ from previous results like BK15 [82]. Looking forward to BK24, we expect an improvement in $\sigma(r)$ of roughly a factor of 5 (including delensing with SPT-3G). If the reader is to take anything from this work, it would be that all the results presented in previous chapters are used to inform our systematics studies so that we can keep pushing down $\sigma(r)$ without being limited by systematics. We do not guarantee that the bias on r from systematics like $T \rightarrow P$ leakage will continue to scale down proportionally with $\sigma(r)$, but we aim to keep it subdominant by at least a factor of 2-4.

¹We did not do this for BK18 since the level of contamination is a factor of 6 lower than $\sigma(r)$.

Bibliography

- [1] Bennett, C., Bogges, N., Cheng, E., Hauser, M., Kelsall, T., Mather, J., Moseley Jr, S., Murdock, T., Shafer, R., Silverberg, R., et al., “Scientific results from coBE,” *Advances in Space Research* **13**(12), 409–423 (1993).
- [2] Collaboration, P., Ade, P. A., Aghanim, N., Alves, M., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., Aussel, H., et al., “Planck 2013 results. i. overview of products and scientific results,” (2013).
- [3] Aghanim, N. et al., “Planck 2018 results. vi. cosmological parameters,” *Astron. Astrophys* **641**, A6 (2020).
- [4] Baumann, D., [*Cosmology*], Cambridge University Press (2022).
- [5] Einstein, A., “Die feldgleichungen der gravitation,” *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften* **1915**, 844–847 (1915).
- [6] Hubble, E., “A relation between distance and radial velocity among extra-galactic nebulae,” *Proceedings of the national academy of sciences* **15**(3), 168–173 (1929).
- [7] Riess, A. G. and Breuval, L., “The local value of h_0 ,” *Proceedings of the International Astronomical Union* **18**(S376), 15–29 (2022).
- [8] Guth, A. H., “The inflationary universe,” in [*Cosmology*], 411–446, CRC Press (2023).
- [9] Mukhanov, S., “Cosmological perturbations from inflation,” *Journal of Physics A: Mathematical and Theoretical* **40**(25), 6561–6572 (2007).

-
- [10] Starobinsky, A. A., “Stochastic de sitter (inflationary) stage in the early universe,” in [*Field Theory, Quantum Gravity and Strings: Proceedings of a Seminar Series Held at DAPHE, Observatoire de Meudon, and LPTHE, Université Pierre et Marie Curie, Paris, Between October 1984 and October 1985*], 107–126, Springer (2005).
- [11] Kofman, L., Linde, A., and Starobinsky, A. A., “Towards the theory of reheating after inflation,” *Physical Review D* **56**(6), 3258 (1997).
- [12] Akrami, Y., Arroja, F., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A., Barreiro, R., Bartolo, N., Basak, S., et al., “Planck 2018 results. i. overview and the cosmological legacy of planck,” *arXiv preprint arXiv:1807.06205* (2018).
- [13] Lewis, A. and Challinor, A., “Camb: Code for anisotropies in the microwave background,” *Astrophysics source code library* , ascl-1102 (2011).
- [14] Story, K., Reichardt, C., Hou, Z., Keisler, R., Aird, K., Benson, B., Bleem, L., Carlstrom, J., Chang, C., Cho, H.-M., et al., “A measurement of the cosmic microwave background damping tail from the 2500-square-degree spt-sz survey,” *The Astrophysical Journal* **779**(1), 86 (2013).
- [15] Lesgourgues, J., “The cosmic linear anisotropy solving system (class) i: Overview,” *arXiv preprint arXiv:1104.2932* (2011).
- [16] Penzias, A. A. and Wilson, R. G., “A measurement of excess antenna temperature at 4080-mc/s,” *The Astrophysical Journal* **142**, 419–421 (1965). <https://api.semanticscholar.org/CorpusID:120759902>.
- [17] Weiland, J., Odegard, N., Hill, R., Wollack, E., Hinshaw, G., Greason, M., Jarosik, N., Page, L., Bennett, C., Dunkley, J., et al., “Seven-year wilkinson microwave anisotropy probe (wmap) observations: planets and celestial calibration sources,” *The Astrophysical Journal Supplement Series* **192**(2), 19 (2011).

- [18] Novaes, C., Bernui, A., Ferreira, I., and Wuensche, C., “Searching for primordial non-gaussianity in planck cmb maps using a combined estimator,” *Journal of Cosmology and Astroparticle Physics* **2014**(01), 018–018 (2014).
- [19] Hu, W. and Sugiyama, N., “Small-scale cosmological perturbations: an analytic approach,” *The Astrophysical Journal* **471**(2), 542–570 (1996).
- [20] Silk, J., “Cosmic black-body radiation and galaxy formation,” *Astrophysical Journal*, vol. 151, p. 459 **151**, 459 (1968).
- [21] Eisenstein, D. J. and Hu, W., “Baryonic features in the matter transfer function,” *The Astrophysical Journal* **496**(2), 605–614 (1998).
- [22] Jackson, J. D., [*Classical electrodynamics*], John Wiley & Sons (1999).
- [23] Dodelson, S. and Schmidt, F., [*Modern cosmology*], Academic press (2020).
- [24] Kovac, J. M., Leitch, E., Pryke, C., Carlstrom, J., Halverson, N., and Holzzapfel, W., “Detection of polarization in the cosmic microwave background using dasi,” *Nature* **420**(6917), 772–787 (2002).
- [25] Zaldarriaga, M. and Seljak, U., “Gravitational lensing effect on cosmic microwave background polarization,” *Physical Review D* **58**(2), 023003 (1998).
- [26] Hanson, D., Hoover, S., Crites, A., Ade, P., Aird, K., Austermann, J., Beall, J., Bender, A., Benson, B., Bleem, L., 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 (2013).
- [27] Ade, P. A., Ahmed, Z., Amiri, M., Barkats, D., Thakur, R. B., Bischoff, C., Beck, D., Bock, J., Boenish, H., Bullock, E., et al., “Improved constraints on primordial gravitational waves using planck, wmap, and bicep/keck observations through the 2018 observing season,” *Physical review letters* **127**(15), 151301 (2021).

-
- [28] Hand, E., *ILC Analysis of BICEP2 and Keck Array CMB Polarization Data through the 2015 Observing Season*, PhD thesis, University of Cincinnati (2023).
- [29] Abergel, A., Ade, P. A., Aghanim, N., Alves, M., Aniano, G., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., et al., “Planck 2013 results. xi. all-sky model of thermal dust emission,” *Astronomy & Astrophysics* **571**, A11 (2014).
- [30] Adam, R., Ade, P. A., Aghanim, N., Alves, M., Arnaud, M., Ashdown, M., Aumont, J., Baccigalupi, C., Banday, A., Barreiro, R., et al., “Planck 2015 results-x. diffuse component separation: Foreground maps,” *Astronomy & Astrophysics* **594**, A10 (2016).
- [31] Chiang, H. C., Ade, P. A., Barkats, D., Battle, J. O., Bierman, E. M., Bock, J. J., Dowell, C. D., Duband, L., Hivon, E. F., Holzapfel, W. L., et al., “Measurement of cosmic microwave background polarization power spectra from two years of bicep data,” *The Astrophysical Journal* **711**(2), 1123–1140 (2010).
- [32] Irwin, K. D. and Hilton, G. C., “Transition-edge sensors,” in [*Cryogenic particle detection*], 63–150, Springer (2005).
- [33] Ade, P. A., Aikin, R. W., Barkats, D., Benton, S., Bischoff, C. A., Bock, J., Brevik, J., Buder, I., Bullock, E., Dowell, C., et al., “Detection of b-mode polarization at degree angular scales by bicep2,” *Physical Review Letters* **112**(24), 241101 (2014).
- [34] Paine, S., “The am atmospheric model,” *arXiv preprint arXiv:1905.05694* (2019).
- [35] Collaboration, B., Keck, A., PAR, A., et al., “Bicep/keck xv: The bicep3 cmb polarimeter and the first three year data set,” (*No Title*) (2021).
- [36] Ade, P., Ahmed, Z., Amiri, M., Barkats, D., Thakur, R. B., Bischoff, C., Beck, D., Bock, J., Boenish, H., Buza, V., et al., “Bicep/keck xix: Extremely thin composite polymer vacuum windows for bicep and other high throughput millimeter wave telescopes,” *arXiv preprint arXiv:2411.10428* (2024).

- [37] Ade, P. A., Pisano, G., Tucker, C., and Weaver, S., “A review of metal mesh filters,” *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III* **6275**, 248–262 (2006).
- [38] Cheshire, J., *Generating Deepest-Yet Cosmic Microwave Background Polarization Maps with the BICEP/Keck Experiments*, PhD thesis, University of Minnesota (2024).
- [39] Hui, H., Ade, P., Ahmed, Z., Aikin, R., Alexander, K. D., Barkats, D., Benton, S. J., Bischoff, C. A., Bock, J. J., Bowens-Rubin, R., et al., “Bicep array: a multi-frequency degree-scale cmb polarimeter,” in [*Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*], **10708**, 75–89, SPIE (2018).
- [40] Cornelison, J., *Constraints on Electromagnetic Parity Violation in the Cosmic Microwave Background using BICEP3*, PhD thesis, Harvard University Graduate School of Arts and Sciences (2023).
- [41] Battistelli, E. S., Amiri, M., Burger, B., Halpern, M., Knotek, S., Ellis, M., Gao, X., Kelly, D., Macintosh, M., Irwin, K., 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**(3), 908–914 (2008).
- [42] De Korte, P. A., Beyer, J., Deiker, S., Hilton, G. C., Irwin, K. D., MacIntosh, M., Nam, S. W., Reintsema, C. D., Vale, L. R., and Huber, M. E., “Time-division superconducting quantum interference device multiplexer for transition-edge sensors,” *Review of Scientific Instruments* **74**(8), 3807–3815 (2003).
- [43] Irwin, K. D., Cho, H.-M., Doriese, W. B., Fowler, J. W., Hilton, G. C., Niemack, M. D., Reintsema, C. D., Schmidt, D. R., Ullom, J. N., and Vale, L. R., “Advanced code-division multiplexers for superconducting detector arrays,” *Journal of Low Temperature Physics* **167**(5), 588–594 (2012).
- [44] Grayson, J. A., Ade, P., Ahmed, Z., Alexander, K. D., Amiri, M., Barkats, D., Benton, S., Bischoff, C. A., Bock, J., Boenish, H., et al., “Bicep3 performance overview and

-
- planned keck array upgrade,” in [*Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*], **9914**, 157–173, SPIE (2016).
- [45] Palladino, S. M., *Constraining Primordial Gravitational Waves with BICEP/Keck Array Telescopes and Developing the BICEP Array Housekeeping System*, PhD thesis, University of Cincinnati (2021).
- [46] Ade, P., Ahmed, Z., Amiri, M., Barkats, D., Thakur, R. B., Bischoff, C., Beck, D., Bock, J., Boenish, H., Bullock, E., et al., “Bicep/keck xv: The bicep3 cosmic microwave background polarimeter and the first three-year data set,” *The Astrophysical Journal* **927**(1), 77 (2022).
- [47] Finkbeiner, D. P., Davis, M., and Schlegel, D. J., “Extrapolation of galactic dust emission at 100 microns to cosmic microwave background radiation frequencies using firas,” *The Astrophysical Journal* **524**(2), 867–886 (1999).
- [48] Takahashi, Y. D., Ade, P. A., Barkats, D., Battle, J. O., Bierman, E. M., Bock, J. J., Chiang, H. C., Dowell, C. D., Duband, L., Hivon, E. F., et al., “Characterization of the bicep telescope for high-precision cosmic microwave background polarimetry,” *The Astrophysical Journal* **711**(2), 1141–1156 (2010).
- [49] BOENISH, H., BULLOCK, E., BUZA, V., CHESHIRE IV, J., CONNORS, J., CORNELISON, J., CRUMRINE, M., CUKIERMAN, A., DENISON, E., DIERICKX, M., et al., “Bicep/keck xv: The bicep3 cmb polarimeter and the first three year data set,”
- [50] Yoon, K., *Design and Deployment of BICEP: A Novel Small-Aperture CMB Polarimeter to Test Inflationary Cosmology*, PhD thesis, California Institute of Technology (2007).
- [51] Germaine, M. T. S., *Beam Systematics and Primordial Gravitational Wave Constraints from the BICEP/Keck Array CMB Experiments*, PhD thesis, Harvard University (2021).
- [52] Bullock, E., *Polarization Angle Calibration And B-Mode Characterization With The Bicep And Keck Array CMB Telescopes*, PhD thesis, University of Minnesota (2018).

- [53] Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M., “Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere,” *The Astrophysical Journal* **622**(2), 759 (2005).
- [54] Lau, K., *Constraining Inflation Models with the BICEP/Keck B-mode Experiments*, PhD thesis, University of Minnesota (2023).
- [55] Ade, P., Aikin, R., Barkats, D., Benton, S., Bischoff, C. A., Bock, J., Brevik, J., Buder, I., Bullock, E., Dowell, C., et al., “Bicep2. iii. instrumental systematics,” *The Astrophysical Journal* **814**(2), 110 (2015).
- [56] Aikin, R. W., *Testing Inflationary Cosmology with the Bicep1 and Bicep2 Experiments*, PhD thesis, California Institute of Technology (2013).
- [57] Smith, K. M., “Pseudo-cl estimators which do not mix e and b modes,” *Physical Review D—Particles, Fields, Gravitation, and Cosmology* **74**(8), 083002 (2006).
- [58] Ade, P., Ahmed, Z., Aikin, R., Alexander, K. D., Barkats, D., Benton, S., Bischoff, C. A., Bock, J., Bowens-Rubin, R., Brevik, J., et al., “Bicep2/keck array. vii. matrix based e/b separation applied to bicep2 and the keck array,” *The Astrophysical Journal* **825**(1), 66 (2016).
- [59] Bunn, E. F., Zaldarriaga, M., Tegmark, M., and de Oliveira-Costa, A., “E/b decomposition of finite pixelized cmb maps,” *Physical Review D* **67**(2), 023501 (2003).
- [60] Willmert, J., *Constraining inflationary B-modes with the BICEP/Keck array telescopes*, PhD thesis, University of Minnesota (2019).
- [61] Lewis, A., “Lenspix: Fast mpi full sky transforms for healpix,” *Astrophysics Source Code Library*, ascl-1102 (2011).
- [62] Halal, G., Clark, S. E., Cukierman, A., Beck, D., and Kuo, C.-L., “Filamentary dust polarization and the morphology of neutral hydrogen structures,” *The Astrophysical Journal* **961**(1), 29 (2024).

-
- [63] An, K., “Sulla determinazione empirica di una legge di distribuzione,” *Giorn Dell’inst Ital Degli Att* **4**, 89–91 (1933).
- [64] Collaborations, B., Ade, P., Ahmed, Z., Aikin, R., Alexander, K., Barkats, D., Benton, S., Bischoff, C., Bock, J., Bowens-Rubin, R., et al., “Bicep2/keck array x: Constraints on primordial gravitational waves using planck, wmap, and new bicep2/keck observations through the 2015 season,” *arXiv preprint arXiv:1810.05216* (2018).
- [65] Collaboration, P., Arnaud⁶⁷, M., Ashdown⁶³, M., Atrio-Barandela¹⁷, F., Aumont⁵⁴, J., Baccigalupi⁸², C., Banday⁸⁸, A., Barreiro⁵⁹, R., Battaner⁹⁰, E., Benabed⁵⁵, K., et al., “Planck intermediate results,” *A&A* **586**, A134 (2016).
- [66] Bennett, C. L., Larson, D., Weiland, J. L., Jarosik, N., Hinshaw, G., Odegard, N., Smith, K., Hill, R., Gold, B., Halpern, M., et al., “Nine-year wilkinson microwave anisotropy probe (wmap) observations: final maps and results,” *The Astrophysical Journal Supplement Series* **208**(2), 20 (2013).
- [67] Hamimeche, S. and Lewis, A., “Likelihood analysis of cmb temperature and polarization power spectra,” *Physical Review D—Particles, Fields, Gravitation, and Cosmology* **77**(10), 103013 (2008).
- [68] Lewis, A. and Bridle, S., “Cosmological parameters from cmb and other data: A monte carlo approach,” *Physical Review D* **66**(10), 103511 (2002).
- [69] Hester, J. J., Scowen, P. A., Sankrit, R., Burrows, C. J., Gallagher III, J. S., Holtzman, J. A., Watson, A., Trauger, J. T., Ballester, G. E., Casertano, S., et al., “Wfpc2 studies of the crab nebula. i. hst and rosat imaging of the synchrotron nebula,” *Astrophysical Journal v. 448, p. 240* **448**, 240 (1995).
- [70] Ade, P., Ahmed, Z., Aikin, R., Barkats, D., Benton, S., Bischoff, C., Bock, J., Bowens-Rubin, R., Brevik, J., Buder, I., et al., “Bicep2/keck array xi: Beam characterization and temperature-to-polarization leakage in the bk15 data set,” *The Astrophysical Journal* **884**(2), 114 (2019).

- [71] Cornelison, J., Verges, C., Ade, P., Ahmed, Z., Amiri, M., Barkats, D., Thakur, R. B., Beck, D., Bischoff, C. A., Bock, J. J., et al., “Improved polarization calibration of the bicep3 cmb polarimeter at the south pole,” in [*Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*], **12190**, 829–848, SPIE (2022).
- [72] Collaboration, B., “Bicep/keck xviii: Measurement of bicep3 polarization angles and consequences for constraining cosmic birefringence and inflation,” *arXiv preprint arXiv:2410.12089* (2024). v3, revised 17 Feb 2025.
- [73] Giannakopoulos, C., Vergès, C., Ade, P. R., Ahmed, Z., Amiri, M., Barkats, D., Thakur, R. B., Bischoff, C. A., Beck, D., Bock, J. J., et al., “Calibration measurements of the bicep3 and bicep array cmb polarimeters from 2017 to 2024,” *arXiv preprint arXiv:2409.16440* (2024).
- [74] Karkare, K. S., *Multifrequency beam systematics and measurements of B-mode polarization with the BICEP/Keck Array cosmic microwave background experiments*, PhD thesis, Harvard University (2017).
- [75] Eiben, M., Barkats, D., Balkaski, A., Crystian, S., Dierickx, M. I., Goldfinger, D. C., Grimes, P. K., Kimberk, R., Kovac, J. M., Meiners, G., et al., “Laminate polyethylene window development for large aperture millimeter receivers,” in [*Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*], **12190**, 1083–1096, SPIE (2022).
- [76] Shiu, C., Soliman, A., O’Brien, R., Steinbach, B., Bock, J. J., Frez, C. F., Jones, W. C., Megerian, K. G., Moncelsi, L., Schillaci, A., et al., “Design and performance of a 30/40 ghz diplexed focal plane for the bicep array,” *The Astrophysical Journal Supplement Series* **272**(1), 12 (2024).
- [77] Ade, P. A., Aikin, R., Barkats, D., Benton, S., Bischoff, C. A., Bock, J., Bradford, K., Brevik, J., Buder, I., Bullock, E., et al., “Bicep2/keck array. iv. optical characterization and performance of the bicep2 and keck array experiments,” *The Astrophysical Journal* **806**(2), 206 (2015).

-
- [78] Ge, F., Millea, M., Camphuis, E., Daley, C., Huang, N., Omori, Y., Quan, W., Anderes, E., Anderson, A., Ansarinejad, B., et al., “Cosmology from cmb lensing and delensed ee power spectra using 2019–2020 spt-3g polarization data,” *Physical Review D* **111**(8), 083534 (2025).
- [79] Yang, H.-I. E., *Constraining Primordial Gravitational Waves and Characterizing B-mode with the Keck Array and BICEP3 CMB experiments*, PhD thesis, Stanford University (2021).
- [80] Gelaro, R., McCarty, W., Suárez, M. J., et al., “The modern-era retrospective analysis for research and applications, version 2 (merra-2),” *Journal of Climate* **30**(14), 5419–5454 (2017).
- [81] Eiben, M., *Windows to the Universe: Improving the Sensitivity of High Throughput Millimeter Telescopes*, PhD thesis, Harvard University (2025).
- [82] Ade, P., Ahmed, Z., Aikin, R., Alexander, K., Barkats, D., Benton, S., Bischoff, C., Bock, J., Bowens-Rubin, R., Brevik, J., et al., “Constraints on primordial gravitational waves using p lanck, wmap, and new bicep2/k eck observations through the 2015 season,” *Physical review letters* **121**(22), 221301 (2018).
- [83] Zhang, C., *The Development and Performance of the First BICEP Array Receiver At 30/40 GHz For Measuring The Polarized Synchrotron Foreground*, PhD thesis, California Institute of Technology (2023).
- [84] Soliman, A., Ade, P., Ahmed, Z., Amiri, M., Barkats, D., Thakur, R. B., Bischoff, C., Bock, J., Boenish, H., Bullock, E., et al., “Optical design and characterization of 40-ghz detector and module for the bicep array,” *Journal of Low Temperature Physics* **199**(3), 1118–1126 (2020).
- [85] Soliman, A., Ade, P., Ahmed, Z., Aikin, R., Alexander, K., Barkats, D., Benton, S., Bischoff, C., Bock, J., Bowens-Rubin, R., et al., “Design and performance of wide-band

- corrugated walls for the bicep array detector modules at 30/40 ghz,” in [*Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*], **10708**, 551–562, SPIE (2018).
- [86] Mertler, C. A., Vannatta, R. A., and LaVenia, K. N., [*Advanced and multivariate statistical methods: Practical application and interpretation*], Routledge (2021).

Appendices

A BA30/40 Detector Time Constant Measurements

In this appendix we describe direct calibration measurements of detector time constants for the BICEP Array 30/40 GHz receiver using the linearly polarized BSNS in the near field. These measurements were performed in the 2023-2024 austral summer season only for the BA30/40 receiver. This was the first time that direct measurements with a calibration source were conducted for this receiver. We start by briefly reviewing transition-edge sensor (TES) bolometer time constants and how they arise, and then describe the measurement and results along with a comparison of time constants extracted from galaxy scans. For the TES bolometer review we rely on [32], and [83].

TES bolometers operate by measuring the temperature change of a thermally isolated absorber that is weakly coupled to a cold bath. The absorber and TES reside on a small island with total heat capacity C and are connected to the bath through thermal links with conductance G . In the absence of electrical feedback, the thermal evolution of the island temperature T in response to a perturbation in absorbed power can be described by the energy balance equation

$$C \frac{dT}{dt} = P_{\text{opt}} + P_{\text{J}} - G(T - T_b), \quad (\text{A.1})$$

where P_{opt} is the optical power absorbed by the detector, P_{J} is the electrical Joule power

dissipated in the TES, and T_b is the bath temperature. Linearizing ² this equation about the operating point yields a characteristic thermal response time

$$\tau_0 = \frac{C}{G}, \quad (\text{A.2})$$

which is the intrinsic thermal time constant of the detector. Physically, τ_0 represents the time required for the island to return to equilibrium after a small perturbation in power. Increasing the heat capacity slows the response, while increasing the thermal conductance to the bath speeds it up. TES detectors are typically operated under voltage bias, which produces negative electrothermal feedback (ETF). Under voltage bias the Joule power is

$$P_J = \frac{V^2}{R(T, I)}, \quad (\text{A.3})$$

where V is the bias voltage and $R(T, I)$ is the TES resistance, which depends strongly on both temperature and current in the superconducting transition. If the absorbed optical power increases, the temperature of the TES rises, which increases the resistance. Because the device is voltage biased, this increase in resistance reduces the current through the TES, thereby reducing the Joule heating. The reduction in Joule power partially cancels the original increase in absorbed power, stabilizing the detector and accelerating its thermal response.

The strength of this feedback is quantified by the electrothermal loop gain

$$\mathcal{L} = \frac{P_0 \alpha}{GT_0}, \quad (\text{A.4})$$

where P_0 is the Joule power at the operating point, T_0 is the TES temperature, and:

$$\alpha = \frac{\partial \ln R}{\partial \ln T} \quad (\text{A.5})$$

²Linearizing in the context of this appendix refers to Taylor expanding around the operating point and keeping terms up to first order.

is the logarithmic sensitivity of the TES resistance to temperature. A large α corresponds to a steep superconducting transition, which produces strong electrothermal feedback. For typical TES detectors, α can be on the order of 100, leading to loop gains \mathcal{L} of several tens or more. For BICEP/Keck instruments the loop gain is typically $\mathcal{L} \sim 70 - 200$, depending on the operating point of the detector and the steepness of the transition. Electrothermal feedback modifies the thermal response of the detector by reducing the effective time constant. Linearizing the coupled electrical and thermal equations of motion gives

$$\tau_{\text{eff}} = \frac{\tau_0}{G(1 + \mathcal{L})}. \quad (\text{A.6})$$

Thus, strong electrothermal feedback ($\mathcal{L} \gg 1$) can reduce the effective time constant by more than an order of magnitude compared to the intrinsic thermal timescale. This reduction allows TES detectors to respond rapidly to changes in optical power while maintaining thermal stability crucial for CMB observation at high data rates. In the frequency domain, the detector response to a modulation in optical power is approximately that of a single-pole low-pass filter with transfer function

$$H(f) = \frac{1}{1 + i2\pi f\tau_{\text{eff}}}. \quad (\text{A.7})$$

This transfer function is used in our deconvolution scheme described in section 2.7 using the measured time constants from galaxy observations [83]. It is these time constants that we try to measure via direct measurements with the linearly polarized BSNS. In addition to the thermal response, the bias circuit introduces an electrical time constant determined by the inductance L of the readout circuit and the TES resistance R ,

$$\tau_{\text{el}} = \frac{L}{R}. \quad (\text{A.8})$$

For stable operation the electrical response must be faster than the thermal response, requiring

$$\tau_{\text{el}} \ll \tau_{\text{eff}}. \tag{A.9}$$

Real TES detectors often exhibit more complex behavior than a single thermal time constant because the absorber, TES film, and supporting structures can have distinct heat capacities and weak internal thermal links (reduced G). These internal degrees of freedom introduce additional thermal modes, leading to a detector transfer function that may contain multiple poles in the frequency domain. For these reasons, we do not model the detector response from first principles, but instead use our CMB measurements during the galaxy phase to empirically determine τ_{eff} for each detector. This is done by comparing the observed galaxy signal in the time domain of the right and left-going scans separately and by forming 'scan-direction jackknife' maps. We fit for the time constant that minimizes the residual signal in the scan-direction jackknife maps ³.

With the measurement parameters now defined, we move on to the measurement itself. We design new calibration schedules to perform 2 minute long 'noise stares' (not scans) on the Titanium transition with the 40 GHz BSNS mounted on top of the mast on the same building as the receiver. The chop rate is set to 0.5 Hz (electronically chopped and therefore essentially a perfect square wave, unlike the chopper) to fully explore both the up-going and down-going transitions. From a high power FSL measurement on Titanium with this receiver and these biases (see section 4.3), we empirically determine the region we expect our detectors to saturate (very close to the source), which guides our choice of pointing offset from the source and its power level. Looking at FSL timestreams on the Titanium transition, we empirically define the saturation region to be $\sim 20^\circ$ radius around the source. The source attenuation is set to -15 dB and the instrument is pointed 40° away from the source in azimuth and 5° away in elevation to avoid saturating the detectors. To ensure that both detector polarizations couple to the source, we repeat the noise stares at 8 equally spaced Dk angles. A set of 8 noise stares is referred to as a *stare set*. Each stare set is taken at the nominal CMB bias points, and at 9 evenly spaced bias offsets from the nominal bias points, for a total of 10 bias

³This is a simplification of the actual process we use in the reanalysis. It is ever evolving as we work our way to the BK24 results, so further discussion is deferred.

A. BA30/40 Detector Time Constant Measurements

points per stare set.

Figure A.1 shows a schematic of the measurement setup in (X,Y) boresight-centered coordinates (see section 3.4), where the BA30/40 detectors are projected onto the sky and the source geometry overlotted.

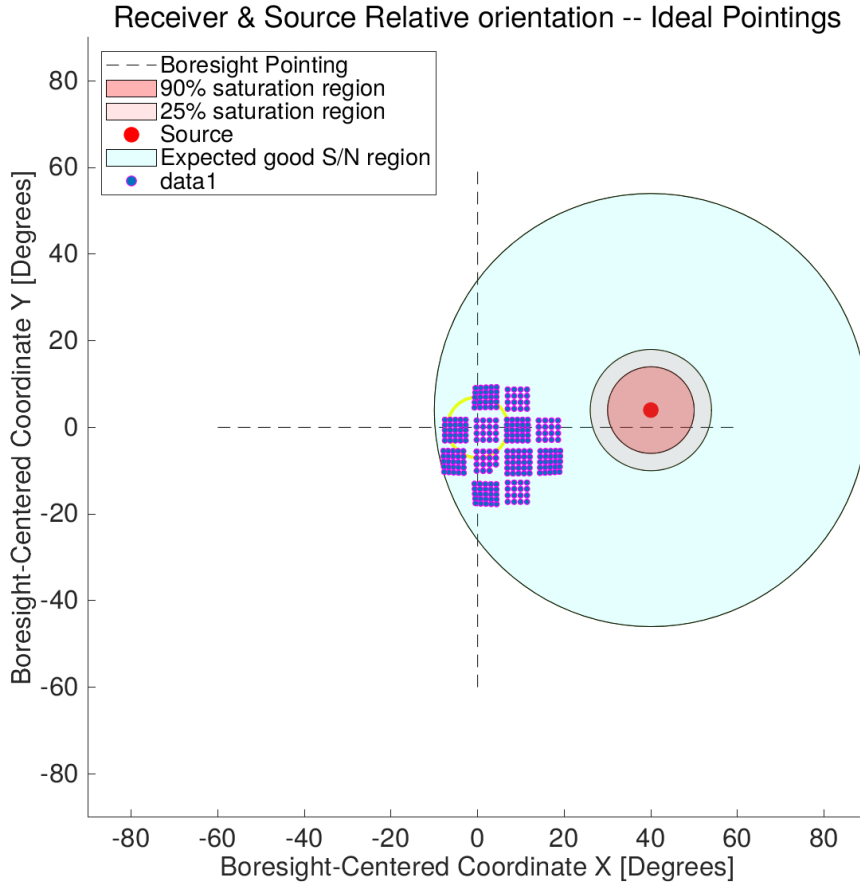


Figure A.1. Schematic of the time constant measurement setup in (X,Y) boresight-centered coordinates with the 40 GHz BSNS electronically chopped at 0.5 Hz. The BA30/40 FPU is projected onto the sky and the source geometry overlotted. The empirically defined saturation region is shown in red. Although this plot shows a single Dk angle, we execute ‘noise stares’ at 8 equally spaced Dk angles to ensure coupling of both detector polarizations with the polarization axis of the source.

Up-going is defined as the transition corresponding to increasing optical power, while down-going is the transition corresponding to decreasing optical power. Although to first order the time constant should be the same for both transitions, real TES detectors can have different time constants for the up-going and down-going transitions. Close examination of equation A.6 reveals that since the loop gain \mathcal{L} depends on the logarithmic sensitivity α (which

changes along the superconducting transition), the time constant can differ for the transition corresponding to increasing optical power (up-going) and the transition corresponding to decreasing optical power (down-going)–(see sections 2.5-2.8 in [32]). We define τ_{\uparrow} and τ_{\downarrow} for the up-going and down-going transitions, respectively and from equation A.6 we see that: $\tau_{\uparrow} < \tau_{\downarrow}$. This is expected to be a strong function of both operating point and the shape of the superconducting transition (e.g. nonlinearity in $R(T, I)$).

Analysis of Noise Stares

A detector’s response to the chopped source signal that looks like a perfect square wave in the time domain indicates short time constants. A rounded step response indicates longer time constants. We fit a simple 3-parameter exponential model to the up-going and down-going transitions to extract τ_{\uparrow} and τ_{\downarrow} for each detector and bias point. The model is given by:

$$s(t) = A(1 - e^{-(t-t_0)/\tau}) + B, \quad (\text{A.10})$$

To be precise, we first stack the transitions across the 2-minute stare and average them together to get a single up-going and down-going transition per stare set, to which the model is fit. Time constants are grouped by fitted amplitude (metric of how well the detector couples to the source), Dk angle, and bias point. There are many degrees of freedom in the measurement. For example, we can look at the dependence of the time constants on the bias point, or on the Dk angle, or on the fitted amplitude, or on distance from the source, etc. Here we show only the more interesting trends. Figure A.2 shows the raw timestreams of the noise stares for a representative 40 GHz detector at the nominal bias point, along with the fitted model to the up-going and down-going transitions (relevant statistics in the plot legend).

As expected, the representative detector shown in Figure A.2 has very short time constants on the order of $\mathcal{O}(10)$ ms for both the up-going and down-going transitions. However, a fraction of BA30/40 detectors ($\sim 10\%$) have substantially longer fitted up-going and down-

A. BA30/40 Detector Time Constant Measurements

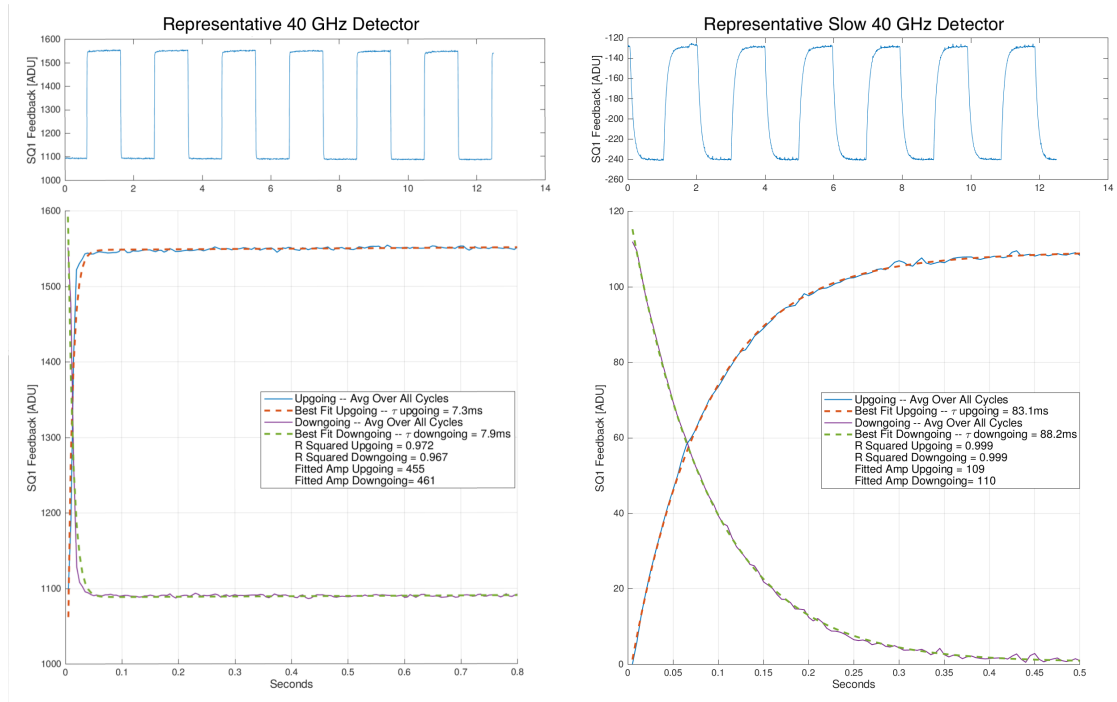


Figure A.2. Raw timestreams of the noise stars for a representative 40 GHz detector at the nominal bias point (left) and a slow 40 GHz detector (right). The population of slow detectors is $\sim 10\%$. The up-going and down-going transitions are fit with an exponential model to extract τ_{\uparrow} and τ_{\downarrow} respectively. Relevant statistics of the fit and extracted time constants are shown in the plot legend.

going time constants on the order of $\mathcal{O}(100)$ ms. In addition, the fitted time constants for the up-going transition are systematically shorter than those for the down-going transition.

The dependence of the time constants on the bias point is visualized in Figure A.3. As expected, the time constants for both the up-going and down-going transitions are a strong function of the bias point. BA30/40 is the first receiver in the BICEP/Keck program that has detectors with seemingly long time constants. To summarize the time constant measurements from this analysis, we pick those extracted from the fit with the largest amplitude (best coupling to the source) for each detector and bias point. For the nominal bias points, Figure A.4 summarizes them.

The time constants for BA30 detectors are longer on-average than those for BA40 detectors, also noted in [83]. Although the bin size of these histograms is too large to resolve systematic differences between up-going and down-going transitions, the up-going time constants are systematically shorter than the down-going ones. For the nominal CMB biases, the median up-going time constant for BA30 is $\sim 22.0 \pm 1.1$ ms, while the median down-going time

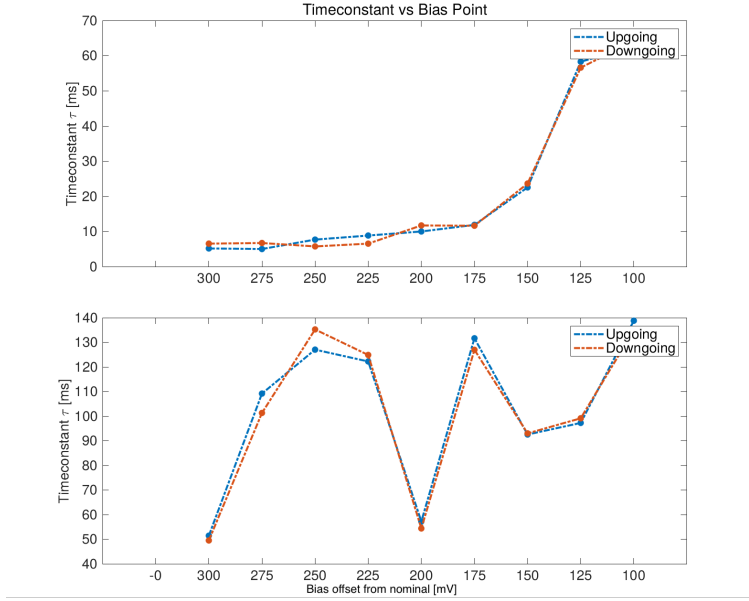


Figure A.3. Dependence of the fitted time constants on the bias point for the same representative 40 GHz detector shown in Figure A.2. As expected, the time constant for both the up-going and down-going transitions are a strong function of the bias point. The nominal bias point is not plotted here.

constant for BA30 is $\sim 23.4 \pm 1.0$ ms. The corresponding median up-going and down-going time constants for BA40 are $\sim 8.3 \pm 0.7$ ms and $\sim 9.68 \pm 0.8$ ms, respectively. We have 2 noise stare sets at the nominal bias points, and the quantity quoted as the uncertainty in the time constant is the standard deviation of the measurements⁴. An argument can be made here that the standard deviation of the measurements is not a good representation of the uncertainty, since we have no reason to believe the underlying distribution of the time

⁴Ensuring good coupling with the source by not including time constants extracted from detector responses of less than 100 ADU.

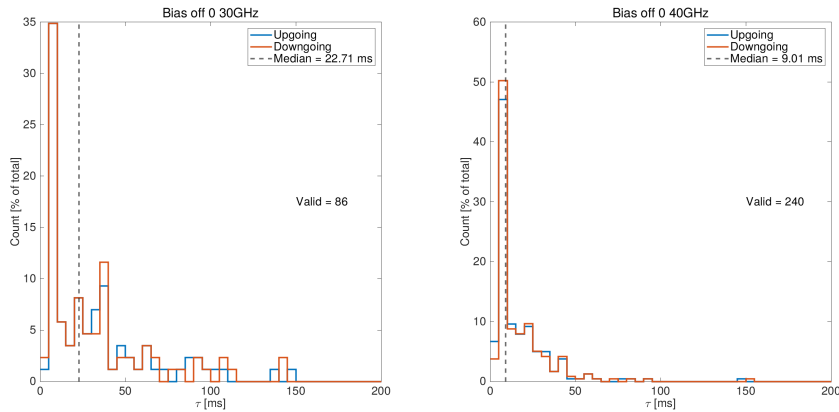


Figure A.4. Time constant measurements for BA30 and BA40 detectors at the nominal bias points. Time constants for BA30 (left) are longer on-average than those for BA40 (right).

A. BA30/40 Detector Time Constant Measurements

constants is Gaussian. For this reason, we cannot confidently say that we have measured a systematic difference between the up-going and down-going time constants, and more noise stares at the nominal bias points would be required to make more definitive statements about this.

Finally, we compare the time constants extracted from the noise stares with those from the galaxy scans. The correlation is plotted in Figure A.5 and illustrates a very poor correlation between the two. This may indicate that the two datasets are measuring different physical quantities, which is not well understood. Given the poor correlation, these time constants are not used in the mainline pipeline as we have more confidence in those extracted from the galaxy scans. In addition, there are no more plans to perform noise stares with the BSNS for BA30/40 or for other receivers.




Figure A.5. Correlation plot between time constants extracted from the noise stares and those extracted from the galaxy scans. The $x = y$ line is shown in orange. The correlation is very poor, and at this point it is not understood why.

B BA150 Corrugation features

One of the main uses of unpolarized FFBM with the thermal chopper is to probe for any anomalous features in the beam maps that might contaminate our undeprojected residuals. In this section, we describe the discovery and impact of beam features in the BA150 receiver associated with reflections off the corrugated tile frames that hold the detector tiles in place. As mentioned in section 2.2, all modern BICEP/Keck receivers have corrugated tile frames, to mitigate reflections that would otherwise couple to the detectors [84]. In fact, prior to fabrication we simulate the expected beam features arising from reflections off the corrugated tile frames, using the commercially available software package CST Microwave Studio ⁵. Our corrugation design, specifically optimized for each observing frequency, uses seven corrugation slots with a depth of $\lambda/4$, height of $\lambda/8$, and frame height of 7.2 mm. A schematic of the corrugated tile frames is shown in Figure B.1.


We can easily verify by converting mm distances on the FPU to degrees on the sky if there are any optical features in the beam maps associated with reflections off the corrugated tile frames. Unpolarized FFBM with the thermal chopper is perfect for this application (calibration data with the 150 GHz BSNS are planned but have not been taken yet). As discussed in [85], the corrugation design parameters (depth, height, width, etc.) are optimized such that there is no coupling of the detectors closest to the frame edges to the corrugated tile frames, and therefore no beam features associated with them. This is verified in the unpolarized FFBM beam maps for all BA150 tiles.

The BA150 tiles are ordered in the general control program (GCP) by their rows and columns ⁶. Each tile has 18 rows and 18 columns of orthogonal and co-located detector pairs (648 total optically coupled detectors per tile ⁷). During the 2024 FFBM campaign, we discovered that the detectors in the rows/columns that are second closest to the tile edges (i.e. rows/columns 2 and 17) show beam features that could be attributed to coupling to the


⁵<https://www.3ds.com/products/simulia/cst-studio-suite>

⁶This is not the same as multiplexing rows/columns of the readout system.

⁷Each tile contains some non-optically coupled detectors (dark detectors, or open detectors) as well, so the total number read out is 656.



(a) Schematic of the corrugated tile frames used in BICEP Array 150 GHz receiver. The corrugation design is optimized using the CST modeling software to minimize reflections for detectors closest to the tile edges.



(b) Geometry of the corrugated tile frames used in BICEP Array 150 GHz receiver. Relevant design parameters are included.

Figure B.1. With knowledge of BA150’s plate scale to convert from mm distances on the FPU to degrees on the sky (curved focal plane), we can verify if there is an optical feature of the detector’s beam associated with reflections off the corrugated tile frames.

corrugated tile frames. Figure B.2 shows the locations on the tiles we see the beam features and the corresponding composite beam maps from a representative sample of detectors in those locations.

These beam features align with the direction of the tile frames (i.e. all detectors in row 2 show features in tile frame direction) and if a detector is close to the corner of the tile (e.g. row 2 and column 2) there are two beam features in the direction of the two tile frames. Although beam maps from a single tile are shown in this figure, all BA150 tiles show the

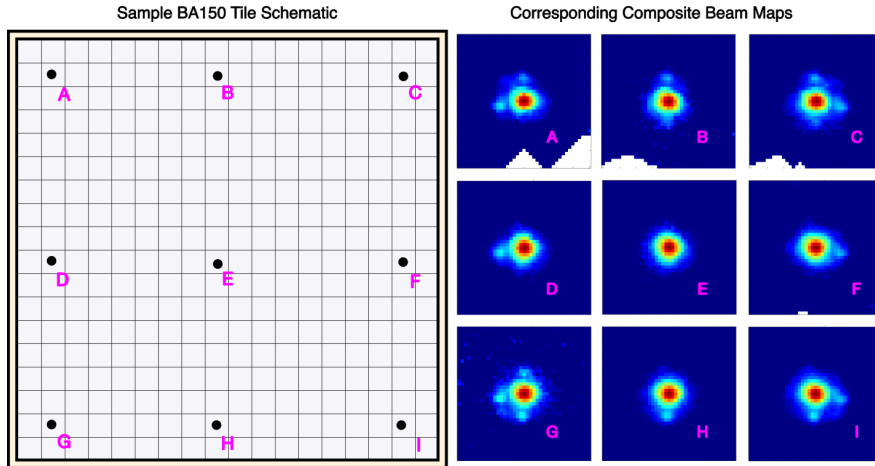


Figure B.2. BA150 tile schematic (left) and corresponding $2^\circ \times 2^\circ$ beam maps from strategically selected detectors (right) at the indicated locations (A,B,C,...). The beam maps were taken during the 2024 FFBM campaign with the thermal chopper, and are shown in the x', y' coordinates defined in section 3.4. The detectors in the rows/columns that are second closest to the tile edges (i.e. rows/columns 2 and 17) show localized beam features in the direction of the tile frames. Detector E (and all detectors far from the tile edges) do not show this behavior. Although beam maps from a single tile are shown here, all BA150 tiles show the same behavior.

same behavior.

Since the $2^\circ \times 2^\circ$ beam thumbnails are plotted in detector-centered coordinates (x', y') defined in section 3.4, we can verify if they align with the location of the tile frame. To do this, we use the plate scale of BA150 (~ 15.9 mm/deg)⁸, to convert the mm distances on the FPU to degrees on the sky. The frame's total width is ~ 22.5 mm. For the detectors second closest to the tile edges, the inner edge of the frame would be located at $\sim 0.4^\circ$, and the outer edge of the frame would be located at $\sim 1.8^\circ$.

Looking at the beam maps in Figure B.2, the beam features visually appear to be located at $\sim 0.7^\circ$, which is right in the middle of the tile frame. However, to extract the exact location we fit a 2D elliptical Gaussian (equation 3.30) to the beam maps after masking the main beam out to 3σ ⁹. We extract the amplitude and locations of the beam features plotted in Figure B.3.

The beam features are located at $\sim 0.7^\circ$, which is right in the middle of the tile frame as expected. More interestingly, the amplitude of the beam features is different between polarization A (above the diagonal) and B (below the diagonal) detectors, which would be

⁸BA150's FPU is curved, so this number varies as a function of distance from the center of the FPU.

⁹This sigma is taken per-detector as described in section 4.1.2.

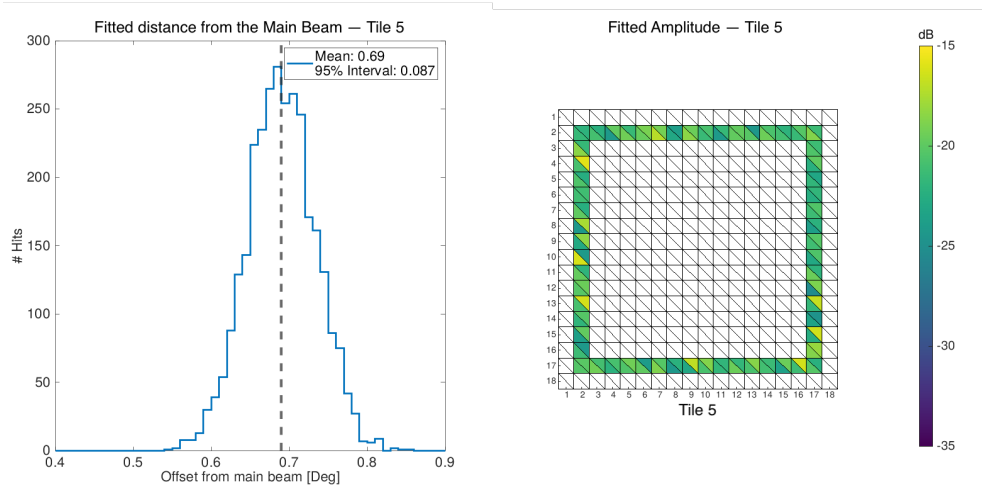


Figure B.3. Amplitude (right) and location (left) of the beam features associated with reflections off the corrugated tile frames for detectors in second-closest rows/columns to the tile edges. The amplitude plot only highlights the detectors which are expected to show the beam features (i.e. those in rows/columns 2 and 17). Each square corresponds to a detector pair and each diagonal separates the A and B amplitudes. The features are located at $\sim 0.7^\circ$, which is right in the middle where the tile frame is expected to be. Since these correspond to reflections, it is unsurprising that their amplitude is different between polarization A and B detectors.

consistent with reflections preferentially coupling to one of the two orthogonal polarizations.

The amplitude of the beam features is on the order of ~ 20 dB relative to the beam peak.

As BA150’s map weight increases with more and more CMB data, we expect to reach the level of sensitivity where these beam features will start to contaminate our timestreams noticeably. For example, they might lead to a failure to pass the tile-inner-vs-outer jackknife test [38] since the outermost rows/columns of the tiles consistently see these features. We might consider deprojecting or subtracting them. In the case of subtraction, this analysis provides all necessary information to fix the template coefficients (see section 2.7.3) and remove them from the timestreams. Prior to doing this, since these features are localized but not azimuthally symmetric, we could use the machinery developed in section 5.1 to simulate their impact on our final BB spectrum. This is another high-impact project (along with the ground contamination simulations described in section 6.5) that requires a minimal amount of work to complete since the codebase and data products are already in place.

C Noise modeling in Ground Spectra

Modeling the noise in the ground spectra $\tilde{N}(m, f)$ (dropping the ν for simplicity)¹⁰ is important for making accurate statements about the power contained at sub-hour timescales. Although the ground mapping presented in chapter 6 is a very high signal-to-noise measurement, we need to verify that the noise contribution is not drowning the signal. In addition, having a good understanding of the noise power spectrum would allow us to generate random realizations of noise timestreams with statistical properties empirically motivated by the data. For this, we need a model for the noise power spectrum. In this appendix, we start with a discrete noise timestream signal $x[n]$ and a likelihood function for the noise power spectrum:

$$d[n], \quad n = 0, 1, \dots, N - 1, \quad (\text{C.1})$$

where each $d[n]$ is an independent and identically distributed (I.I.D.) Gaussian random variable with zero mean and variance σ^2 :

$$d[n] \sim \mathcal{N}(0, \sigma^2). \quad (\text{C.2})$$

When making power spectra we first take the Fourier transform and then take the absolute square of the real and imaginary components of the complex Fourier amplitude. The discrete Fourier transform (DFT) of the sequence $d[n]$ at frequency bin k is defined as:

$$X[k] = \sum_{n=0}^{N-1} d[n] e^{-j\frac{2\pi}{N}kn}. \quad (\text{C.3})$$

Since $d[n]$ are Gaussian and the DFT is a linear transform, $X[k]$ is a complex Gaussian random variable:

¹⁰In this appendix ν does not denote the observing frequency like it did in chapter 6.

$$X[k] = X_r[k] + jX_i[k], \quad (\text{C.4})$$

where,

$$X_r[k] = \sum_{n=0}^{N-1} d[n] \cos\left(\frac{2\pi}{N}kn\right), \quad X_i[k] = -\sum_{n=0}^{N-1} d[n] \sin\left(\frac{2\pi}{N}kn\right). \quad (\text{C.5})$$

Both $X_r[k]$ and $X_i[k]$ are linear combinations of Gaussian variables, which makes them also Gaussian. We are now interested in the statistical properties of the real and imaginary part of the Fourier transformed noise signal.

Because each $d[n]$ has zero mean,

$$\mathbb{E}[X_r[k]] = \mathbb{E}[X_i[k]] = 0. \quad (\text{C.6})$$

The variances are:

$$\text{Var}[X_r[k]] = \sigma^2 \sum_{n=0}^{N-1} \cos^2\left(\frac{2\pi}{N}kn\right), \quad (\text{C.7})$$

$$\text{Var}[X_i[k]] = \sigma^2 \sum_{n=0}^{N-1} \sin^2\left(\frac{2\pi}{N}kn\right). \quad (\text{C.8})$$

The orthogonality of the DFT basis vectors implies that the sums are equal:

$$\sum_{n=0}^{N-1} \cos^2\left(\frac{2\pi}{N}kn\right) = \sum_{n=0}^{N-1} \sin^2\left(\frac{2\pi}{N}kn\right) = \frac{N}{2}. \quad (\text{C.9})$$

Hence:

$$\text{Var}[X_r[k]] = \text{Var}[X_i[k]] = \frac{N}{2}\sigma^2. \quad (\text{C.10})$$

Since $X_r[k]$ and $X_i[k]$ are independent zero-mean Gaussian variables with equal variance, the power at each frequency bin $P[k]$, is the sum of squares of two independent Gaussian

random variables [86]. If Z is a zero-mean Gaussian with variance σ^2 , then Z^2/σ^2 follows a chi-square distribution with 1 degree of freedom:

$$\frac{Z^2}{\sigma^2} \propto \chi_1^2. \quad (\text{C.11})$$

Therefore,

$$\frac{X_r[k]^2}{\text{Var}[X_r[k]]} + \frac{X_i[k]^2}{\text{Var}[X_i[k]]} \sim \chi_2^2, \quad (\text{C.12})$$

which is a chi-square distribution with 2 degrees of freedom.

Thus, the power $P[k]$ is distributed as:

$$P[k] \propto \text{Scaled } \chi_2^2, \quad (\text{C.13})$$

with scale parameter $\text{Var}[X_r[k]] = \frac{N}{2}\sigma^2$.

Probability density function (PDF) of $P[k]$:

A χ^2 distribution with 2 degrees of freedom is equivalent to an exponential distribution with parameter $\frac{1}{2}$:

$$f_Z(z) = \frac{1}{2}e^{-z/2}, \quad z \geq 0. \quad (\text{C.14})$$

Therefore, the PDF of $P[k]$ is exponential with mean $\lambda = \frac{N}{2}\sigma^2$:

$$f_{P[k]}(p) = \frac{1}{\lambda}e^{-p/\lambda}, \quad p \geq 0. \quad (\text{C.15})$$

where p is the value that $P[k]$ can take.

Substituting in and writing in terms of x we get:

$$P_\nu(f) = \frac{1}{N\sigma^2} \exp\left(-\frac{f}{N\sigma^2}\right), \quad f \geq 0. \quad (\text{C.16})$$

where $P_\nu(f)$ is the PDF of the noise power spectrum at frequency ν and f is the value that the noise power spectrum can take at that frequency. A common parameterization of the noise power spectrum is a combination of 'white noise' and '1/f noise':

$$\tilde{N}_\nu = N_0 \left(1 + \left(\frac{\nu}{\nu_*}\right)^\alpha\right), \quad (\text{C.17})$$

where N_0 is the white noise level, ν_* is the knee frequency where the 1/f noise starts to dominate, and α is the slope of the 1/f noise. With this model governing the statistical properties of the noise power spectrum equation C.16 can be written as:

$$P_\nu(f) = \frac{1}{N_\nu} \exp\left(-\frac{f}{N_\nu}\right), \quad f \geq 0. \quad (\text{C.18})$$

where N_ν is normalized to unity and $\langle f \rangle = N_\nu$ (shown via integration by parts). This PDF is equivalent to the likelihood of obtaining a particular value of the power spectrum, f , given model parameters N_0 , ν_* , and α . So we can rewrite it as

$$\mathcal{L}_\nu(N_0, \nu_*, \alpha | f) = \frac{1}{N_\nu(N_0, \nu_*, \alpha)} \exp\left(-\frac{f}{N_\nu(N_0, \nu_*, \alpha)}\right). \quad (\text{C.19})$$

The main difference between equation C.18 and equation C.19 is that for equation C.18 we hold the model fixed and consider the probability for different values of f , while for equation C.19 we fix the data realization f and allow the model parameters to vary. $P(f)$ is normalized so that it integrates to 1 over $f \in [0, \infty)$, while the normalization of the likelihood with respect to the model parameters is unspecified.

We make the assumption that the noise is uncorrelated and stationary, meaning that the statistical properties of the noise do not depend on time. Of course, this is not true during the ground mapping schedules as a whole, but $\tilde{N}(m, f)$ is fitted with this model at each

frequency bin independently, making this assumption reasonable. Effectively, this means that each bin of the power spectrum is independent of the others which is a sufficient assumption for the Fourier modes of the Az direction of the ground spectra m .

Therefore, the joint likelihood for the entire power spectrum is the product of the likelihoods at each frequency:

$$\mathcal{L} = \prod_{\nu} \mathcal{L}_{\nu}. \quad (\text{C.20})$$

It is often more convenient to work with the log-likelihood. Taking the logarithm converts the product into a sum:

$$\log \mathcal{L} = \sum_{\nu} \log \mathcal{L}_{\nu}. \quad (\text{C.21})$$

Substituting the likelihood expression for each frequency gives:

$$\log \mathcal{L} = - \sum_{\nu} \left(\frac{f}{N_{\nu}} + \log N_{\nu} \right). \quad (\text{C.22})$$

The fit shown in Figures 6.8 and 6.10 is obtained by maximizing this log-likelihood with respect to the model parameters N_0 , ν_* , and α , or equivalently by minimizing the negative log-likelihood $-\log \mathcal{L}$. This could be part of the timestream-level simulations discussed at the end of Chapter 6 to simulate noise timestreams in conjunction with the signal timestreams to quantify the amount of azimuth-synchronous signal power in our CMB spectra.