

HARVARD UNIVERSITY
Graduate School of Arts and Sciences




DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the
Department of Physics
have examined a dissertation entitled


Beam Systematics and Primordial Gravitational Wave Constraints from the BICEP/Keck Array CMB
Experiments

presented by Michael Tyler St. Germaine

candidate for the degree of Doctor of Philosophy and hereby
certify that it is worthy of acceptance.

Signature 


Typed name: Professor John Kovac, Chair

Signature 

Typed name: Professor Christopher Stubbs

Signature 

Typed name: Professor Ron Walsworth

Signature 

Typed name: Professor Daniel Eisenstein

Date: May 3, 2021

Beam Systematics and Primordial Gravitational Wave Constraints from the BICEP/Keck Array CMB Experiments

A DISSERTATION PRESENTED
BY
MICHAEL TYLER ST. GERMAINE
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
MAY 2021

©2021 – MICHAEL TYLER ST. GERMAINE
ALL RIGHTS RESERVED.

Beam Systematics and Primordial Gravitational Wave Constraints from the BICEP/Keck Array CMB Experiments

ABSTRACT

Inflation, an extension to the standard Λ CDM model which posits a brief, accelerated expansion early in the Universe, naturally solves the horizon and flatness problems with the standard model and provides a source of the initial perturbations that seed large-scale structure. Most inflationary models predict a stochastic background of gravitational waves which would imprint a unique B -mode pattern in the polarization of the Cosmic Microwave Background (CMB) that peaks at degree angular scales. The strength of this signature is parametrized by the tensor-to-scalar ratio r . A nonzero measurement of r would provide direct evidence for inflation.

The BICEP/Keck Array series of CMB experiments has been observing the polarized CMB from the Amundsen-Scott South Pole Station since 2006, using small-aperture, on-axis refractors. Continuous integration of a low-foreground, $\sim 1 - 2\%$ patch of the sky has produced maps over multiple frequencies that lead to world-leading constraints on r from B -mode measurements: $r_{0.05} < 0.034$ and $\sigma_r = 0.009$ using data through the 2018 observing season.

The BICEP/*Keck Array* telescopes measure polarization by taking the difference between two co-located, orthogonally polarized detectors. A prominent systematic that must be controlled is the leakage from the bright CMB temperature sky into the polarization measurement due to mismatch in the angular response patterns (beams) of the two detectors within a pair. In this dissertation, we use high-fidelity *in-situ* measurements of the per-detector beam response in conjunction with specialized simulations to quantify the level of temperature-to-polarization ($T \rightarrow P$) leakage expected in the BK18 data set, and the associated impact on the r constraint. We also use Fourier transform spectrometer measurements of the spectral response to constrain the bias on r from band center uncertainties. We introduce simple metrics evaluated only from the beam maps (without simulations) that estimate the $T \rightarrow P$ leakage and other optical properties, and discuss progress of optical modeling of small-aperture telescope beams, including their ability to predict $T \rightarrow P$ leakage. Finally, we assess the utility of these simulations, metrics, and optical models as we move forward to the next generation of CMB experiments with hundreds of thousands of detectors.

Contents

FRONT MATTER	i
Title	i
Copyright	ii
Abstract	iii
Contents	v
List of Figures	vii
List of Tables	x
Dedication	xi
Acknowledgments	xii
1 INTRODUCTION	1
1.1 The Standard Cosmological Model	3
1.2 The Cosmic Microwave Background	7
1.3 Inflation	17
1.4 Thesis Outline	29
2 THE BICEP/Keck Array EXPERIMENT	33
2.1 Detectors	35
2.2 Detector Readout	39
2.3 Optics	42
2.4 Cryostat	46
2.5 Telescope Mount and Observatory Site	49
2.6 Observation Strategy	52
3 BANDPASS MEASUREMENTS AND SYSTEMATICS	55
3.1 Fourier Transform Spectrometer	57
3.2 Analysis & Results	63
3.3 Uncertainties and Impact on Likelihood Analysis	69

4	BEAM MEASUREMENTS AND ANALYSIS PIPELINE	82
4.1	Far-Field Beam Mapping Procedure	84
4.2	Far-Field Beam Map Analysis Pipeline	90
4.3	Near-Field Beam Maps	110
4.4	Far-Sidelobe Maps	115
5	FAR-FIELD BEAM CHARACTERIZATION	120
5.1	Alternate Polarization Axis Definition	122
5.2	2D Elliptical Gaussian Fit Parameters	123
5.3	Array-Averaged Beams and Beam Window Functions	130
5.4	Validation Against CMB Data	136
5.5	Ghost “Buddy” Beams	141
5.6	Crosstalk Beams	144
5.7	Residual Beams and Boresight Rotation	148
6	BK18 TEMPERATURE-TO-POLARIZATION LEAKAGE SIMULATIONS	154
6.1	T-to-P Leakage Simulation Pipeline	156
6.2	Multicomponent Likelihood Analysis	169
6.3	BK18 Baseline Results	172
6.4	BK18 T-to-P Leakage Simulation Baseline Results	178
6.5	Impact of Crosstalk on T-to-P Leakage	196
7	BEAM METRICS AND MODEL COMPARISONS	203
7.1	Simple Beam Metrics	206
7.2	GRASP Simulations	223
7.3	Huygens-Fresnel Simulations	235
8	CONCLUSIONS	241
	REFERENCES	253

Listing of figures

1.1	Map of CMB Temperature Anisotropies	9
1.2	CMB TT Spectrum	12
1.3	Diagram of E -Mode Polarization	16
1.4	Slow-Roll Inflation	21
1.5	Diagram of B -Mode Polarization	25
1.6	CMB Theory Power Spectra	26
1.7	Current State of BB Measurements	28
1.8	BICEP/ <i>Keck Array</i> Projections	29
2.1	TES Bolometers Used in BICEP/ <i>Keck Array</i>	37
2.2	SQUID Readout Diagram for <i>Keck Array</i>	40
2.3	Optics in BICEP/ <i>Keck Array</i> Receivers	43
2.4	BICEP/ <i>Keck Array</i> Cryostats	47
2.5	Photo of MAPO and DSL at South Pole Station	50
2.6	BICEP3 Standard Observing Schedule	53
3.1	Photo of FTS Interferometer	59
3.2	Photo of FTS on <i>Keck</i>	61
3.3	Sample Interferogram and Spectrum	65
3.4	BICEP/ <i>Keck Array</i> Frequency Spectra	67
3.5	BICEP3 Spectra Before/After LPE Filter Replacement	68
3.6	FTS Coupling Aperture Correction Factor	72
3.7	BICEP3 Band Center Uncertainties	76
3.8	Shift in BK15 Posteriors Due to Band Center Error	80
3.9	BK18 ML Search Results with Band Center Error	81
4.1	Far-Field Beam Mapping Setup	86
4.2	Photogrammetry of Redirecting Mirror	89
4.3	Beam Map Before/After Demodulation	93

4.4	Improvement in Demodulation Routine	95
4.5	Beam-Centered Coordinate System	96
4.6	BICEP3 and <i>Keck</i> FPU Projected On Sky	97
4.7	FFBM Mirror Mask	103
4.8	Sample BICEP3 Composite Beam Map	106
4.9	Beam Map Pipeline Transfer Function	111
4.10	Near-Field Beam Mapper Apparatus	113
4.11	Sample BICEP3 Near-Field Beam Maps	115
4.12	Far-Sidelobe Measurement Setup	117
4.13	Sample BICEP3 Far-Sidelobe Map	118
5.1	<i>Keck</i> 2016-2018 Differential Pointing	128
5.2	BICEP3 2016-2018 Differential Pointing	129
5.3	<i>Keck</i> 2016-2018 Differential Ellipticity	131
5.4	BICEP3 2016-2018 Differential Ellipticity	132
5.5	Array-Averaged Beam Maps	134
5.6	BICEP/ <i>Keck Array</i> Beam Window Functions	135
5.7	Deprojection Coefficients vs Beam Parameters	138
5.8	<i>Keck</i> 220 GHz Beam Window Function Derived from CMB	140
5.9	Ghost "Buddy" Beam Diagram	142
5.10	Integrated Buddy Beam Power Per Detector	143
5.11	Integrated Crosstalk Beam Power Per Detector	145
5.12	Phasing of Crosstalk Beams	146
5.13	Band-Averaged Residual Beams	149
5.14	Band-Averaged Residual Beams After Boresight Rotation	152
6.1	Standard Deprojection Templates	161
6.2	Suppression Factor and Bandpower Window Functions for BICEP2	168
6.3	BK18 Q Maps	173
6.4	BK18 E/B Maps	175
6.5	BK18 Power Spectra	176
6.6	BK18 Marginalized Posteriors	179
6.7	BK18 Leakage Simulation Q Maps	181
6.8	BK18 Leakage Q Maps with Varying Deprojection	183
6.9	BK18 Leakage EE/BB Auto Spectra	185

6.10	BK18 Leakage Spectra Cross with Real	190
6.11	BK18 ML Search Results with Injected $T \rightarrow P$ Leakage	194
6.12	<i>Keck</i> 2016 Real+Leakage Simulation Spectra with Crosstalk Deprojection	198
6.13	<i>Keck</i> 2016 Leakage Spectra Cross with Real with Crosstalk Deprojection	200
7.1	BICEP3 Beam Steer from NFBM	208
7.2	BICEP3 Edge Taper from NFBM	210
7.3	BICEP3 NFBM Metric Comparison Plots	213
7.4	$T \rightarrow P$ Leakage in <i>Keck</i> 220 vs 230 GHz	216
7.5	BICEP3 $T \rightarrow P$ Leakage Compared Against NFBM Metrics	218
7.6	BICEP3 and <i>Keck</i> 95 GHz $T \rightarrow P$ Leakage Inner vs Outer Tile	219
7.7	Direct-Island Coupling Deprojection Template	221
7.8	<i>Keck</i> 220 GHz Leakage with Direct-Island Coupling Deprojection	222
7.9	GRASP Near-Field Simulations of BICEP3	227
7.10	GRASP Simulation Truncation at Field Lens	228
7.11	GRASP Far-Field Simulations of BICEP3	230
7.12	GRASP Simulations of BICEP3 Residual Beams	232
7.13	GRASP Residual Beam Compared to Real BICEP3 Residual Beam	233
7.14	Huygens-Fresnel Simulation of BICEP3	237
7.15	Huygens-Fresnel Simulation of BICEP Array High-Frequency Beams	238

Listing of tables

1.1	Nomenclature of 220 GHz Bands	32
2.1	Deployment History of BICEP/ <i>Keck Array</i> Telescopes	35
2.2	BICEP/ <i>Keck Array</i> Focal Plane Unit History	38
2.3	Optical Parameters for BICEP/ <i>Keck Array</i> Telescopes	43
3.1	Average FTS Parameters for BICEP/ <i>Keck Array</i> Frequency Bands	70
3.2	List of FTS Aperture Correction Factors	73
3.3	Band Center Uncertainties	78
4.1	Far-Field Criteria for BICEP/ <i>Keck Array</i> Telescopes	85
4.2	Far-Field Beam Map Schedule Parameters	90
5.1	<i>Keck</i> / BICEP3 2016 2D Gaussian Fit Parameters	124
5.2	<i>Keck</i> / BICEP3 2017 2D Gaussian Fit Parameters	125
5.3	<i>Keck</i> / BICEP3 2018 2D Gaussian Fit Parameters	125
5.4	BICEP3 2016-2018 $V-H$ 2D Gaussian Fit Parameters	126
6.1	BK18 Leakage Simulation ρ and PTE Statistics	191
6.2	<i>Keck</i> 2016 Crosstalk Leakage Simulation ρ Statistics	201

FOR MY MOM, MY DAD, AND MY SISTER.

Acknowledgments

My time in graduate school has been a journey – one that would be impossible to do alone. I have relied on the love and support of countless friends and family, and it is often difficult to find the opportunity to give appropriate thanks and recognition to these people who mean so much to me. The work presented here is a testimony to the enduring guidance I’ve been fortunate enough to have these past few years and throughout all my life.

First, I would like to thank my mentor and advisor, John Kovac. The most important aspect to success in graduate school, in my opinion, is the relationship one has with their advisor. By that metric, I could not be happier with the way grad school played out for me. From the frequent one-on-one meetings, teaching AY191, and working together in the lab at Harvard and at Pole, I have learned so much and gained enough confidence to believe that I can become great at anything to which I apply myself. I would also like to thank the members of my committee – Chris Stubbs, Ron Walsworth, and Daniel Eisenstein – for their wisdom and support, especially in the final months as I developed and polished this document.

My coworkers and friends of the Harvard CMB lab, past and present, made every day so much easier. All of this work builds on a foundation that was developed by

previous beam expert Kirit Karkare, who deserves special recognition for his infinite patience and wisdom as I learned the ropes. Other lab members who initially helped me feel warm and welcome in the lab are Colin Bischoff, Denis Barkats, Jake Connors, Victor Buza, Rachel Bowens-Rubin, Maddie Hickman, and Sam Harrison. As the years went on, others joined the lab who also aided me in honing my skills as a CMB scientist and shaped me into a better person overall. I thank Miranda Eiben for her cheerfulness, fellow Arizonan James Cornelison for our common love of Mexican food, Ben Schmitt for his love of American football (and the Hail Murray!), David Goldfinger for his appreciation of board games, and the French contingency of Denis, Ben Racine, Clara Verges, and Marion Dierickx for the great conversations, the laughs, and my head start on learning French. Special thanks to Marion for being a great rock climbing partner and Pole team leader. I am also grateful to Steve Sansone for handling our (often last-minute) machining jobs with expertise and great humor, and to Irene Coyle for her warm heart and for keeping this lab running smoothly.

Other members of the *BICEP/Keck Array* collaboration have provided guidance and fond memories, especially in our time together at Pole. I have significantly overlapped with Howard Hui, Lorenzo Moncelsi, Justin Willmert, Ed Young, Ari Cukierman, Cyndia Yu, Kenny Lau, and Eric Yang, who all helped make my time in Antarctica even more unforgettable, including our 100% participation rate in the 2018 Race Around the World, and that time the SMuRFs dyed their hair “blue.”

Of course, my friends old and new have been there with me every step of the way.

I feel incredibly fortunate at this age to still be in such close contact with friends I made in middle school and high school. To my fellow members of the 501st – Tony, Michael, Chris, Alex, Jason, and Ryan – thank you for the countless long nights of entertainment and laughs. I thank Katie for the adventures in Arizona, Boston, Portland, and British Columbia. Ed has always been a great friend, and was the only friend I had in Boston when I first came here – thanks for the long talks about sports and for crushing me repeatedly in Madden. The new friends I made here in Boston quickly became family, and it has been wonderful going through grad school together with them, from “work hard, party hard” G1s to “we’re too old for this” G6s. Special thanks to lifelong friends Andy and Catherine for simply being there for me (and Maddie) whenever I needed camaraderie.

Lastly, I must thank my family – Mom, Dad, and Danielle – for the continuing, unconditional love and support they have given me my whole life. They had to endure a few holiday seasons with me on the other side of the planet, in an incompatible time zone with limited means of communication. They also had to weather my endless complaints about my hatred of flying each and every time I visited. Whether in Arizona or in Oregon, they always have a bed and home for me whenever I need it, and being with them will always be worth even the most turbulent flight.

The universe is a big place, perhaps the biggest.

Kurt Vonnegut

1

Introduction

THE UNIVERSE, AS WE OBSERVE IT TODAY, is vast, empty, and expanding. This expansion, governed by the balance of energy and matter in the Universe at any given time, hints to some “beginning,” where everything existed in a hot, dense state. As it expanded and cooled, protons and electrons eventually combined to form Hydrogen,

and photons were able to stream freely to us today. This relic light we observe gives us a snapshot of the Universe in its infancy, and from this snapshot we can develop an understanding of not only the origin of the Universe, but its full evolutionary history. This light is called the Cosmic Microwave Background (CMB), and is one of the pillars of modern cosmology.

More than fifty years of CMB observations from ground-based telescopes and space satellites led to a standard cosmological model that describes the expansion history of the Universe. The Λ CDM model postulates that the Universe began 13.7 billion years ago in a “Big Bang,” and then expanded, with the initial under- and over-densities seeding the large-scale structure we see today. It also suggests an overall composition of baryonic matter, cold dark matter, and dark energy. Dark energy currently dominates the energy content of the Universe and is accelerating the expansion of the Universe. What this model does not explain, however, is 1) the remarkable isotropy we see in the CMB, 2) the spatial flatness of the Universe, and 3) the origins of the initial perturbations.

An extension to the Λ CDM model, inflation, provides answers to these shortcomings of the standard model, by introducing a brief period of exponential growth shortly after the Big Bang. Although inflation is consistent with many existing observations, one unique prediction of inflation that has not yet been confirmed is the existence of a primordial B -mode signal in the CMB polarization. A non-zero detection of this signal would be overwhelming evidence that the inflationary paradigm accurately de-

scribes the nature of the beginning of the Universe.

The BICEP/*Keck Array* experiment is a series of polarimeters at the South Pole targeting this primordial B -mode signature. A confident detection of this signal requires a deep and precise knowledge of the systematics associated with every aspect of the measurement. In this chapter we review the standard Λ CDM model, the physics of the early Universe and generation of polarized signals in the CMB, the nature of inflation, and the current state of BB measurements.

1.1 THE STANDARD COSMOLOGICAL MODEL

The matter and energy content of the Universe plays a critical role in the expansion and dynamics of spacetime. This is elegantly captured in a single set of equations, the Einstein field equations,

$$G_{\mu\nu} - \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \quad (1.1)$$

On the left-hand side, the Einstein tensor $G_{\mu\nu}$ and the metric tensor $g_{\mu\nu}$ describe the curvature and geometry of spacetime, respectively, and Λ is the cosmological constant. On the right-hand side, the stress-energy tensor $T_{\mu\nu}$ encapsulates the energy and momentum attributed to matter and radiation (Einstein 1916). In a simplified sense, matter and energy act as sources of gravitational fields in general relativity, just as mass is the source in Newtonian gravity.

On sufficiently large scales (~ 100 Mpc), the Universe is homogeneous and isotropic

– in other words, we are in no special place in the Universe. This statement is often called the cosmological principle. Starting with only this assumption, we can describe spacetime with the following metric:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - Kr^2} + r^2(d\theta^2 + \sin^2(\theta)d\phi^2) \right]. \quad (1.2)$$

Here a is a dimensionless scale factor (which we will come back to in a moment), 3D distance is described by the coordinates (r, θ, ϕ) , and K represents the curvature of space ($K = 0$ for a flat Universe). This is the Friedmann-Robertson-Walker (FRW) metric.

We can derive the Friedmann equations using only Equations 1.1 and 1.2. The temporal component ($\mu, \nu = 0$) of the Einstein field equations takes the form

$$R_{00} - \frac{1}{2}Rg_{00} - \Lambda g_{00} = \frac{8\pi G}{c^4}\rho g_{00}. \quad (1.3)$$

Note the definition of the Einstein tensor $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}$ has been invoked on the left-hand side, where $R_{\mu\nu}$ is the Ricci tensor and R is the Ricci scalar. We have assumed that the energy density of the Universe can be modeled as a perfect fluid (a safe assumption on large scales), and therefore the temporal component of the stress-energy tensor T_{00} simply becomes ρg_{00} , where ρ is the total density. After evaluating the Ricci scalar and the temporal component of the Ricci tensor, and substituting

$g_{00} = -c^2$ from the FRW metric, we arrive at the first Friedmann equation:

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{Kc^2}{a^2} + \frac{\Lambda}{3}, \quad (1.4)$$

where we have defined the Hubble constant H in terms of the dimensionless scale factor a . A similar analysis of the spatial components of the Einstein field equations leads to the second independent Friedmann equation, also known as the acceleration equation,

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3c^2}(\rho c^2 + 3p) + \frac{\Lambda c^2}{3} \quad (1.5)$$

where p is the total pressure.

If we continue to assume a perfect fluid, and that gravity and velocities are negligible, one can use the continuity equations ($\partial T_\nu^\mu / \partial x^\mu = 0$) to show that

$$\frac{\partial \rho}{\partial t} + \frac{3\dot{a}}{a}(\rho + p) = 0. \quad (1.6)$$

For some general component of the universe, we can relate the density to the pressure using $p = w\rho c^2$. Using this in combination with Equation 1.6, we arrive at a relationship between density and the scale factor:

$$\rho(a) \propto a^{-3(1+w)}. \quad (1.7)$$

The density of matter is $\rho_m \propto a^{-3}$ and therefore has $w = 0$, whereas the density of radiation $\rho_r \propto a^{-4}$ ($w = 1/3$) scales more rapidly due to the wavelength expanding with the growth of the scale factor. This tells us that radiation dominated the energy of the Universe early, when the scale factor was small. The Universe then grew enough that it entered a matter-dominated era, with the transition occurring at a redshift of roughly $z \sim 3000$. We have more recently entered an era where expansion driven by the dark energy dominates ($\rho_\Lambda \propto \text{constant}$, $w = -1$).

If we assume a flat Universe ($K = 0$) and temporarily set aside dark energy ($\Lambda = 0$), Equation 1.4 leads to an expression for the critical density, $\rho_{\text{crit}} = 3H^2/8\pi G$. For $\rho > \rho_{\text{crit}}$, we have a closed Universe which eventually collapses; for $\rho < \rho_{\text{crit}}$ we have an open Universe that expands infinitely. From this definition, we can also define dimensionless energy densities $\Omega_i = \rho_i/\rho_{\text{crit}}$. We can use these to rewrite Equation 1.4 and obtain

$$\frac{H^2}{H_0^2} = \Omega_{0,r}a^{-4} + \Omega_{0,m}a^{-3} + \Omega_{0,k}a^{-2} + \Omega_{0,\Lambda}, \quad (1.8)$$

where H_0 is the value of the Hubble constant at the present time, and $\Omega_{0,i}$ is the present energy density of the i th component. The most recent measurement of the Hubble constant from the *Planck* satellite is $H_0 = 67.36 \pm 0.54 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ([The *Planck* Collaboration 2020c](#)), though there is recent tension of early-Universe measurements of H_0 with late-Universe measurements which give $H_0 = 73.3 \pm 0.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$ ([Verde et al. 2019](#)). In general the cosmological observations show that the

total energy density in the Universe today is roughly 68% dark energy, 27% dark matter, and 5% baryonic matter, with a small contribution (0.005%) from radiation.

1.2 THE COSMIC MICROWAVE BACKGROUND

Long before the proposal and general acceptance of the Λ CDM cosmological model, Edwin Hubble noticed a clear linear relationship between the distance of extra-galactic objects and their recessional velocity (Hubble 1929). The farther away the object, the quicker they were receding away from us. This led to the famous Hubble's law,

$$v_r = H_0 d \tag{1.9}$$

where v_r is the recessional velocity of an object some distance d away, and H_0 is the same Hubble constant discussed in the previous section. This discovery was an indication that the Universe is indeed expanding, and tracing the expansion backward in time, this hints to some point in time corresponding to the beginning of the Universe. This alone, however, was not enough to prove that there was a hot Big Bang.

It wasn't until thirty years later that proof of an early-Universe hot, dense state was incidentally discovered by Bob Wilson and Arno Penzias at Bell Labs. Initially thought to be instrument noise, they concluded that an isotropic 3.5 ± 1.5 K signal was astronomically real (Penzias & Wilson 1965). In a companion paper by Dicke, Peebles, Roll, and Wilkinson, it was concluded that this was indeed a measurement

of blackbody radiation from the Cosmic Microwave Background (Dicke et al. 1965).

Now, because of these discoveries, the CMB is one of the pillars of modern cosmology, and has provided boundless insight into the physics of the early Universe.

After the Big Bang, as the Universe was expanding, photons were tightly coupled to a fluid of protons and electrons via Thompson scattering and Coulomb scattering. Eventually the Universe expanded and cooled enough that electrons and protons combined to form Hydrogen, drastically reducing the amount of scattering and allowing the photons to free stream through the Universe. The last-scattering surface, at a temperature of ~ 3000 K at a redshift of $z \sim 1100$, is known as the time of recombination. Since temperature varies with redshift as $T \sim (1 + z)$, the temperature we observe today is $T_{\text{CMB}} = 2.7548 \pm 0.00057$ K (Fixsen 2009). Due to the rapid interactions between protons, electrons, and photons before recombination, the Universe was in thermal equilibrium, therefore the CMB is one of the best examples of a perfect blackbody in nature.

1.2.1 THE CMB ANGULAR POWER SPECTRUM

In the tightly coupled fluid of photons and electrons before recombination, there were two competing forces, gravity and photon pressure, leading to acoustic oscillations within the fluid. Electrons and photons fall into gravitational wells leading to compression, and the compressed photons then exert outward pressure, causing rarefaction. When recombination happens, we see a “snapshot” of the state of compres-




Figure 1.1: The CMB temperature anisotropies as measured by *Planck*. The monopole, dipole, and the Galaxy (shown in grey outline) have been removed. Figure from [The *Planck* Collaboration 2020a](#).

sion/rarefaction undergone by various acoustic modes at that instant. Some modes will have just collapsed for the first time at recombination, while some will be at maximum rarefaction for the first time. These under- and over-densities lead to hot and cold spots in the CMB, respectively, due to the gravitational redshift applied to the photons escaping the gravitational hills and wells. Similarly, some oscillations were at maximal rarefaction or compression for the second time, leading to hot and cold spots half the size of the larger modes. This logic applies to modes on all scales, leading to the unique temperature anisotropy distribution in the CMB we see today. An image of the CMB temperature anisotropies from *Planck* (monopole, dipole, and Galaxy removed) is shown in Fig. 1.1.

The temperature anisotropy $\Delta T(\theta, \phi)$ at some point in the sky (θ, ϕ) can be decom-

posed into spherical harmonics $Y_{\ell m}(\theta, \phi)$:

$$\Delta T(\theta, \phi) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\theta, \phi). \quad (1.10)$$

Assuming the statistical distribution of the $a_{\ell m}$ s to be Gaussian and isotropic, then all cosmological information is contained within the angular power spectra C_{ℓ} , which is implicitly defined by the variance in the $a_{\ell m}$ s:

$$\langle a_{\ell m} a'_{\ell' m'} \rangle = C_{\ell} \delta_{\ell \ell'} \delta_{m m'}. \quad (1.11)$$

Note that recent measurements have constrained primordial non-Gaussianity to very small levels ([The *Planck* Collaboration 2020b](#)), which validates this assumption. Note that because ΔT is defined as the fluctuation with respect to the mean temperature, the mean $\langle a_{\ell m} \rangle$ is, by design, zero. Also, because the CMB is a two-dimensional surface, the term “power spectrum” used throughout this dissertation refers to the angular power spectrum just derived, and not the full three-dimensional power spectrum.

Because we only get to sample one temperature sky, there is a fundamental limit to the amount of information we can learn about the underlying distribution in each mode. For example, at $\ell = 3$ there are only 7 a_{3m} s drawn from that distribution for

us to measure. This is called the cosmic variance limit:

$$\frac{\Delta C_\ell}{C_\ell} = \sqrt{\frac{2}{2\ell + 1}}. \quad (1.12)$$

The power spectrum can be evaluated by “crossing” a single map X with itself (taking the “auto-spectrum”) or by crossing two separate maps X and Y (taking the “cross-spectrum”). In general, this takes the form:

$$C_\ell^{XY} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m}^X a_{\ell m}^{Y*} \quad (1.13)$$

where $a_{\ell m}^X$ are the $a_{\ell m}$ for map X , and $*$ denotes complex conjugation. X and Y can be the E and B modes measured at a certain frequency, or two E maps measured by different experiments, for example. It is also worth noting that it is common to instead plot the power spectra as $\mathcal{D}_\ell \equiv \ell(\ell + 1)C_\ell/2\pi$, which is approximately flat vs ℓ for low ℓ modes that have not undergone full acoustic oscillations.

The CMB TT spectrum as measured by *Planck* is shown in Fig. 1.2. Many details describing the content and evolutionary history of the Universe can be extracted from this spectrum alone. Using knowledge of the sound speed of the primordial plasma, the location of the first peak tells us that the Universe is very close to spatially flat. The ratio of the first peak to the second peak constrains the amount of baryonic and dark matter. The high- ℓ tail of the spectrum is indicative of strong diffusion damping




Figure 1.2: The CMB temperature spectrum as measured by *Planck*. The blue line is the best-fit Λ CDM model. The spectrum below $\ell \sim 1600$ is cosmic variance limited. Figure from [The *Planck* Collaboration 2020a](#).

during the finite time period of recombination. For more details on the physics of the Λ CDM model learned from the CMB TT spectrum, see [Hu & Dodelson 2002](#).

1.2.2 CMB POLARIZATION

Due to the cosmic variance limit, we are quickly “running out” of constraining power on the Λ CDM model that we can extract from the TT spectrum alone. As we will see throughout the rest of this chapter, a wealth of cosmic information also exists in the polarization of the photons emitted from the CMB. To this end, many space satellite and ground-based CMB experiments over the last couple decades have prioritized CMB polarization measurements at different scales, leading to a rapid growth in our understanding of the early Universe and its evolutionary history.

Thomson scattering, the primary scattering mechanism in the primordial plasma, converts an incoming unpolarized photon into a polarized photon with polarization vector perpendicular to the line of sight. Since photons are constantly Thomson scat-

tering off electrons at all angles, the final polarization state of the scattered photon depends on the local temperature anisotropy around the scattering electron. If the electron were sitting in a uniform temperature bath (a monopole), the scattered photon would be unpolarized. If it takes place in a temperature dipole, the polarization induced by the hot and cold spots 180° apart would perfectly average out and match the polarization at 90°, resulting in no net polarization. However, if the electron sits in a quadrupole anisotropy, the scattered photon would indeed be linearly polarized (Rees 1968).

The Stokes parameters are frequently used to describe the polarization state of an electric field. For an electric field traveling in the $+\hat{z}$ direction with complex amplitudes E_x and E_y , the Stokes parameters are defined as¹:

$$\begin{aligned}
 I &= \langle E_x^2 \rangle + \langle E_y^2 \rangle \\
 Q &= \langle E_x^2 \rangle - \langle E_y^2 \rangle \\
 U &= 2\langle E_x E_y \rangle \cos \phi \\
 V &= 2\langle E_x E_y \rangle \sin \phi
 \end{aligned}
 \tag{1.14}$$

where ϕ is the phase difference between E_x and E_y , and $\langle E_x \rangle$ denotes the time average of E_x . These describe the electric field in terms of total intensity I , two orthogonal linear polarization states Q and U , and circular polarization V . No circular polarization

¹An alternative way of expressing the Stokes U parameter is $U = \langle E_a^2 \rangle - \langle E_b^2 \rangle$, where (a, b) is a Cartesian coordinate system rotated 45° with respect to (x, y) .

is generated in standard cosmology, so we set $V = 0$ here.

We can expand the Stokes linear polarization Q and U in a harmonic series, just as is done for the temperature anisotropies in Eq. 1.10. Unlike temperature, Q and U must be expanded in spin-2 weighted spherical harmonics ${}_{\pm 2}Y_{\ell m}$:

$$(Q \pm iU)(\theta, \phi) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} {}_{\pm 2}a_{\ell, m} {}_{\pm 2}Y_{\ell m}(\theta, \phi) \quad (1.15)$$

This decomposition of Q and U is not rotationally invariant, since Q and U themselves are coordinate dependent. It is more convenient to work in an invariant basis. Following [Bunn et al. 2003](#) and [Zaldarriaga & Seljak 1997](#), we rewrite the spin-2 harmonic coefficients in terms of coordinate-independent scalar coefficients $a_{\ell m}^E$ and pseudo-scalar coefficients $a_{\ell m}^B$:

$$\begin{aligned} a_{\ell m}^E &\equiv -(+2a_{\ell, m} + -2a_{\ell, m})/2 \\ a_{\ell m}^B &\equiv -i(+2a_{\ell, m} - -2a_{\ell, m})/2. \end{aligned} \quad (1.16)$$

We also utilize two combinations of the spin-2 spherical harmonics:

$$\begin{aligned} X_{1, \ell m} &\equiv (+2Y_{\ell, m} + -2Y_{\ell, m})/2 \\ X_{2, \ell m} &\equiv (+2Y_{\ell, m} - -2Y_{\ell, m})/2. \end{aligned} \quad (1.17)$$

We can express a vector of the linear polarization Q and U in terms of linear combina-

tions of the above expressions:

$$\begin{aligned}
\begin{bmatrix} Q(\theta, \phi) \\ U(\theta, \phi) \end{bmatrix} &= -\sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^E \begin{bmatrix} X_{1,\ell m}(\theta, \phi) \\ -iX_{2,\ell m}(\theta, \phi) \end{bmatrix} + a_{\ell m}^B \begin{bmatrix} iX_{2,\ell m}(\theta, \phi) \\ X_{1,\ell m}(\theta, \phi) \end{bmatrix} \\
&= -\sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^E Y_{\ell m}^E(\theta, \phi) + a_{\ell m}^B Y_{\ell m}^B(\theta, \phi),
\end{aligned} \tag{1.18}$$

where the last line has implicitly defined the orthonormal basis functions $Y_{\ell m}^E$ (E modes) and $Y_{\ell m}^B$ (B modes). E modes produce a distinct “curl-free” pattern in polarization, and B modes are similarly “gradient-free.” In total, we have four non-zero power spectra we can calculate between the T , E , and B anisotropy fields:

$$\begin{aligned}
C_{\ell}^{TT} &= \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} a_{\ell m}^{T*} a_{\ell m}^T \\
C_{\ell}^{EE} &= \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} a_{\ell m}^{E*} a_{\ell m}^E \\
C_{\ell}^{BB} &= \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} a_{\ell m}^{B*} a_{\ell m}^B \\
C_{\ell}^{TE} &= \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} a_{\ell m}^{T*} a_{\ell m}^E.
\end{aligned} \tag{1.19}$$

Because the real sky is invariant under parity transformation in the Λ CDM model, any cross between a pseudoscalar ($a_{\ell m}^B$) and a scalar ($a_{\ell m}^T$ or $a_{\ell m}^E$) must be zero². This set of cross-spectra completely describe the CMB under the Λ CDM model.

²Cosmic birefringence is one way to produce parity-violating TB and EB correlations, and is an active field of research. For example, see BK IX ([The BICEP2 and Keck Array Collaborations 2017](#)).




Figure 1.3: Diagram of density perturbations creating E -mode polarization. A photon scattered off of an electron within the quadrupolar temperature anisotropy generated by a density wave must have polarization at an angle 0° or 90° with respect to the wave vector. Figure from BICEP2 Collaboration and Kovac 2003.

As discussed earlier, the mechanism that generates polarization is Thomson scattering off of electrons in quadrupolar temperature anisotropies. These anisotropies can arise in the primordial plasma through two mechanisms. The first is density (scalar) perturbations – longitudinal waves in density that arise naturally within the Λ CDM-predicted acoustic oscillations. Photons interacting with electrons within the gravitational wells and hills at the peaks and troughs of these waves are Doppler shifted, inducing a quadrupole temperature pattern. If a density wave vector is travelling in some direction $\hat{\mathbf{k}}$, the polarization of photons scattering off these electrons must be 0° or 90° with respect to $\hat{\mathbf{k}}$. After superimposing over all density wave vectors, the unique “gradient-like” shape of E modes is the result. This is shown in Fig. 1.3. In summary, the density waves natural to the Λ CDM model can only generate E -mode polarization.

The E -mode anisotropies are a result of the same acoustic oscillations that generate T anisotropies. However, the source of E modes is the temperature anisotropy gener-

ated by the Doppler shift of the photons surrounding the electron. The Doppler shift is proportional to velocity and therefore is maximized/minimized halfway between maximum compression/rarefaction. This means the EE spectrum has a similar structure to the TT spectrum, but at a lower amplitude (the degree-scale anisotropies are roughly 1% polarized) and out of phase. E modes in the CMB were first detected by DASI (Kovac et al. 2002).

The second mechanism that can generate quadrupolar temperature anisotropies is a tensor perturbation to the metric. Such a perturbation is not predicted from the Λ CDM model alone – we must introduce inflation.

1.3 INFLATION

1.3.1 PROBLEMS WITH THE Λ CDM MODEL

As seen in the discussion on CMB temperature anisotropies, the Λ CDM model provides excellent agreement with the observed data, using relatively few, simple parameters. However, for the Universe to appear the way it does to us today – extremely isotropic and spatially flat – the initial conditions would need to be extremely “fine-tuned.” We list below in more detail a number of flaws in the Λ CDM model as it stands alone:

- *The Horizon Problem:* In order for two parts of the sky to be at the same temperature (without fine tuning), they need to have been in causal contact at some point in the evolution of the Universe. At the time of recombination,

the maximum distance light could have traveled through the age of the Universe (the horizon scale) corresponds to an angular separation of about 2° today. This means two patches of CMB sky separated by more than 2° have no reason to be the same temperature. Why is the Universe we observe today so homogeneous and isotropic?

- *The Flatness Problem:* The Universe today is observed to be very close to flat, $\Omega_k = 0.001 \pm 0.002$ ([The Planck Collaboration 2020c](#)). This alone is remarkable – given the infinite range of values between a closed ($\Omega_k > 0$) and an open ($\Omega_k < 0$) Universe, it seems very unlikely to have a value so close to zero. This problem is exacerbated when you consider Eq. 1.8, which tells us that the spatial curvature evolves more slowly ($\propto a^{-2}$) compared to matter ($\propto a^{-3}$) and radiation ($\propto a^{-4}$). In other words, for the Universe to be flat now, it must have been *very* flat in the primordial Universe. How can we explain such a flat Universe without fine tuning?
- *Source of the perturbations:* Although the Λ CDM model elegantly ties the initial perturbations to the content and growth of the Universe, it fails to explain the source of these perturbations, or why they are so small.

It is possible to solve all these problems using inflation, an expansion to the standard Λ CDM model. This extension, initially developed by [Starobinsky 1980](#), [Guth 1981](#), and [Linde 1982](#), posits a brief period of exponential growth very early in the Universe. This rapid growth naturally solves the horizon and flatness problems, as described below.

The horizon problem must be solved by bringing superhorizon modes (modes that had never been in causal contact in standard Λ CDM) into the horizon. Another way to think of this is by the comoving Hubble radius. The maximum distance a photon

can travel during the conformal time $d\tau = dt/a$ is given by:

$$\eta = \int \frac{c dt}{a(t)} = \int c \frac{1}{Ha} d\ln(a), \quad (1.20)$$

where the comoving Hubble radius is defined as $(Ha)^{-1}$. In order to bring superhorizon modes (as we see them today) into the comoving Hubble radius, we require some period of time where the Hubble radius *shrinks*. Inflation is one such way to do that.

The time derivative of the Hubble radius is given by:

$$\frac{d}{dt}(Ha)^{-1} = \frac{d}{dt}(\dot{a})^{-1} = -\frac{\ddot{a}}{(\dot{a})^2}. \quad (1.21)$$

If we require $\frac{d}{dt}(Ha)^{-1} < 0$, then we must have $\ddot{a} > 0$. If we revisit the acceleration equation (Eq. 1.5), we see that this requires negative pressure: $p < \rho c^2/3$. We've already discussed one cosmological component that has negative pressure and therefore drives accelerated expansion – dark energy. If we assume a “dark energy-like” behavior ($p = -\rho c^2$) during inflation, then $a \propto e^{Ht}$. In order to bring superhorizon modes into the comoving Hubble radius, a 60 e -fold expansion (a factor of about 10^{26}) is required.

The flatness problem is solved in a very similar way. Rearranging the first Friedmann equation (Eq. 1.4) in terms of the curvature dimensionless energy density $\Omega_k =$

$\Omega - 1$, we have:

$$\Omega_k = \frac{Kc^2}{H^2 a^2}. \tag{1.22}$$

The flatness problem was a problem of fine-tuning – how could we have a Universe so spatially flat ($\Omega_k \approx 0$)? If we follow the same logic as the solution to the horizon problem, inflation provides a natural solution. If we assume a “dark energy-like” state during inflation, then $a \propto e^{Ht}$ and the above equation reduces to $\Omega_k \propto e^{-2Ht}$. After an inflationary expansion of 60 e -foldings, this term becomes e^{-120} , which gives us the very nearly flat Universe we see today without fine-tuning. Intuitively, this is equivalent to dramatically “zooming in” on any curved surface. Regardless of how curved the Universe was pre-inflation, an expansion of this magnitude “flattens out” this curvature.

The final solution that inflation solves, and perhaps the most important one, is the origin of the perturbations seen in the CMB, which seeds the large-scale structure seen throughout the Universe today. Vacuum fluctuations, which can be thought of as virtual particle pairs creating and annihilating continuously, are brought out of contact with each other during inflation. Newly formed quantum particle pairs, which would normally immediately self-annihilate without inflation, instead become scalar and tensor perturbations to the metric. Before describing the perturbations in detail, it is useful to introduce “slow-roll” inflation.

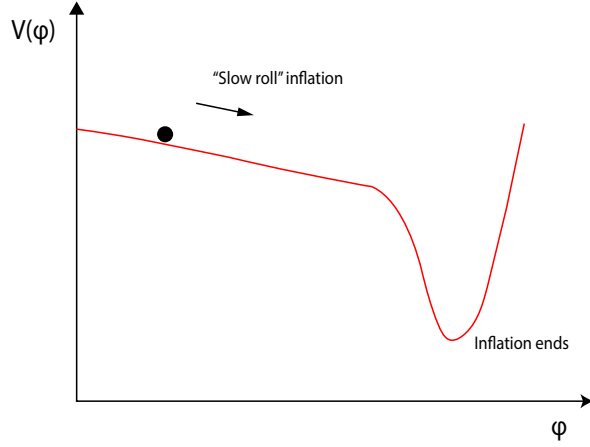


Figure 1.4: Cartoon demonstrating the dynamics of slow-roll inflation. In an analogy to classical mechanics, the scalar field ϕ slowly rolls for a long enough time to allow 60 e -fold inflation, before falling to a potential minimum, ending inflation.

1.3.2 DYNAMICS OF INFLATION

The slow-roll model of inflation assumes a single scalar field, which we will denote here as ϕ . In order to determine how spacetime evolves with ϕ , we are generally required to evaluate the action, which is the integral over the Lagrangian density, as well as the variation of the action with ϕ , and the stress-energy tensor $T_{\mu\nu}$ for a Universe dominated by ϕ (for a review see, e.g., [Baumann 2009](#)). Eventually we arrive at the pressure p and energy density ρ for a scalar field ϕ with potential energy $V(\phi)$:

$$\begin{aligned}\rho &= \frac{1}{2}\dot{\phi}^2 + V(\phi) \\ p &= \frac{1}{2}\dot{\phi}^2 - V(\phi).\end{aligned}\tag{1.23}$$

Note we have set the speed of light $c = 1$. As long as ϕ undergoes a “slow roll,” i.e. $\dot{\phi}^2 \ll V(\phi)$, then we have negative pressure $p \approx -\rho$, which creates the accelerated expansion needed to drive inflation. Eventually, after the necessary 60 e -fold expansion has passed, ϕ rolls toward a minimum potential, breaking the slow-roll condition and therefore ending inflation. This is shown pictorially in Fig. 1.4. The slow-roll parameters are defined as:

$$\begin{aligned}\epsilon &\equiv \frac{1}{16\pi G} \left(\frac{V'}{V} \right)^2 \\ \eta &\equiv \frac{1}{8\pi G} \frac{V''}{V}\end{aligned}\tag{1.24}$$

where V' denotes a derivative of V with respect to ϕ . The slow-roll conditions are $\epsilon \ll 1$ and $|\eta| \ll 1$.

As the Universe is accelerating during inflation, quantum fluctuations to the space-time metric are being injected. These scalar perturbations (denoted ξ , manifested as density waves) and tensor perturbations (denoted γ_σ , manifested as gravitational waves) have the form:

$$\begin{aligned}\langle \xi(\mathbf{k}) \xi(\mathbf{k}') \rangle &= (2\pi)^3 \delta^3(\mathbf{k} + \mathbf{k}') \frac{2\pi^2}{k^3} P_\xi(k) \\ \langle \gamma_\sigma(\mathbf{k}) \gamma_{\sigma'}(\mathbf{k}') \rangle &= (2\pi)^3 \delta_{\sigma\sigma'} \delta^3(\mathbf{k} + \mathbf{k}') \frac{2\pi^2}{k^3} \frac{1}{2} P_t(k)\end{aligned}\tag{1.25}$$

where σ labels the two polarization states of the gravitational waves (these two states also lead to the factor of 1/2 in the tensor equation), and P_ξ and P_t denote the power spectra of scalar and tensor perturbations, respectively. These power spectra in the

single-field slow-roll model inflation are:

$$\begin{aligned}
 P_\xi(k) &= \frac{1}{2\epsilon M_P^2} \left(\frac{H}{2\pi} \right)^2 \Big|_{k=Ha} \\
 P_t(k) &= \frac{8}{M_P^2} \left(\frac{H}{2\pi} \right)^2 \Big|_{k=Ha}
 \end{aligned}
 \tag{1.26}$$

where $M_P = (8\pi G)^{-1}$ is the reduced Planck mass and $k = Ha$ means they are evaluated at the horizon crossing. The spectral tilts of the power spectra of the scalar and tensor perturbations are:

$$\begin{aligned}
 n_s - 1 &\equiv \frac{d \ln P_\xi}{d \ln k} \\
 n_t &\equiv \frac{d \ln P_t}{d \ln k}.
 \end{aligned}
 \tag{1.27}$$

A scale-invariant scalar perturbation spectrum corresponds to $n_s = 1$, but slow-roll inflation predicts a value slightly less than 1, specifically $n_s = 1 - 6\epsilon + 2\eta$ (Liddle & Lyth 2000), which comes out to less than 1 due to both ϵ and η being positive, and $\epsilon \approx \eta$. The current best estimate of n_s from CMB measurements is $n_s = 0.965 \pm 0.004$ (The Planck Collaboration 2020c). A scale-invariant tensor perturbation spectrum corresponds to $n_t = 0$ (the difference in definition with n_s is standard in the literature), and slow-roll inflation predicts $n_t = -2\epsilon$. Since tensor perturbations have not yet been detected, this tilt has not been measured. It is possible to have a running of the spectral tilts – a dependence on scale k . As long as $n_s \approx 1$ and $n_t \approx 0$, we can expand

them about some pivot scale k_* :

$$\begin{aligned} n_s(k) - 1 &= n_s - 1 + \left. \frac{dn_s(k)}{d \ln k} \right|_{k=k_*} \ln(k/k_*) + \dots \\ n_t(k) &= n_t + \left. \frac{dn_t(k)}{d \ln k} \right|_{k=k_*} \ln(k/k_*) + \dots \end{aligned} \quad (1.28)$$

which leads to the following approximation of the perturbation power spectra:

$$\begin{aligned} P_\xi(k) &= A_s \left(\frac{k}{k_*} \right)^{n_s - 1 + \frac{1}{2} \left. \frac{dn_s}{d \ln k} \right|_{k=k_*} \ln(k/k_*) + \dots} \\ P_t(k) &= A_t \left(\frac{k}{k_*} \right)^{n_t + \frac{1}{2} \left. \frac{dn_t}{d \ln k} \right|_{k=k_*} \ln(k/k_*) + \dots} \end{aligned} \quad (1.29)$$

Planck has found that the running of the spectral index is consistent with zero: $dn_s/d \ln k = -0.0045 \pm 0.0067$ ([The *Planck* Collaboration 2020d](#)). Finally, we can define the tensor-to-scalar ratio as:

$$r \equiv \frac{A_t}{A_s} \quad (1.30)$$

which is evaluated at a pivot scale k_* , usually chosen to be 0.002 or 0.05 Mpc^{-1} . The values referenced in this work use 0.05 Mpc^{-1} .

1.3.3 THE SEARCH FOR PRIMORDIAL GRAVITATIONAL WAVES

Most of the predictions of inflation have been confirmed in measurements of the CMB. We see correlations on the sky on superhorizon scales, a very nearly flat Universe, perturbations that are Gaussian, and a scalar perturbation spectral tilt slightly less than




Figure 1.5: Diagram of tensor perturbations creating B -mode polarization. A photon scattered off of an electron within the quadrupolar temperature anisotropy generated by a primordial gravitational waves must have polarization at an angle $\pm 45^\circ$ with respect to with respect to the wave vector. The wave vector is slightly rotated to point out of the page, to demonstrate the unique pattern in temperature created by gravitational waves. Figure from BICEP2 Collaboration and Kovac 2003.

unity. A detection of tensor perturbations to the metric manifested as primordial gravitational waves (PGW), the only prediction not yet observed, would essentially confirm that inflation took place in the early Universe.

Early constraints on r were derived from measurements of the CMB TT spectrum – PGWs add power at large angular scales, so an excess of power over the standard Λ CDM model at low ℓ would correspond to power from tensor perturbations. The *Planck* TT measurement gives $r < 0.11$ at 95 % confidence (The *Planck* Collaboration 2014), but this constraint cannot be improved, due to the cosmic variance limit. We must turn to CMB polarization for tighter constraints.

In Section 1.2.2, we show that polarization can only be generated due to Thomson scattering of photons off of electrons in quadrupolar temperature anisotropies, and that density perturbations only generate E -mode polarization. PGWs, which are uniquely predicted from inflation, can also create quadrupolar temperature anisotropies.

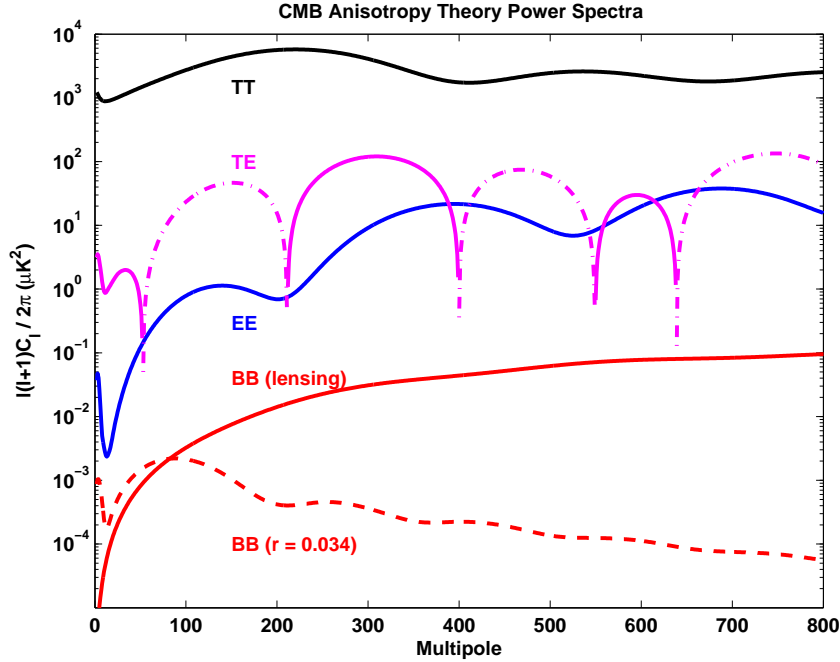


Figure 1.6: Theory power spectra of the CMB anisotropies. These were generated with CAMB (Lewis et al. 2000) using cosmological parameters from *Planck* 2018. The amplitude of the primordial BB spectrum shown corresponds to the BK18 upper limit $r < 0.034$ at 95% confidence (The BICEP2 and Keck Array Collaborations 2021). The unpolarized TT anisotropies are orders of magnitude brighter than polarized EE and BB . The TT and EE spectra are sourced by the same density perturbations, except EE samples the velocity of the plasma while TT samples the state of compression/rarefaction, causing TT and EE to be out of phase. For this same reason, a nonzero TE correlation exists (magenta; the dot-dash line is negative). The lensing BB spectrum is a smoothed version of the EE spectrum, and is brighter than the (current best upper limit) primordial BB spectrum at most scales, showing the need for delensing.

As shown in Fig. 1.5, the symmetry of this quadrupole ($m = \pm 2$) is uniquely different than the quadrupole from density perturbations ($m = 0$). A photon scattering off an electron in an $m = \pm 2$ quadrupole must be polarized at an angle $\pm 45^\circ$ with respect to the wave vector, creating a unique B -mode (“curl-like”) pattern. Because density perturbations cannot create this pattern, a detection of primordial B modes in the polarized CMB would provide direct evidence for inflation.

Although the prospect of identifying the primordial B -mode signature in the CMB is very exciting, it is far from easy. A few of the challenges associated with a detection of primordial B modes includes:

- The amplitude of the primordial BB spectrum is orders of magnitude dimmer than both the TT and EE spectra, as shown in Fig. 1.6. Even with no other external factors considered, exquisite noise levels are required to measure a signal this small.
- Gravitational lensing of primordial E modes creates a B -mode pattern on the sky that is not primordial. This lensing BB spectrum, first measured by SPT (Hanson et al. 2013), followed by POLARBEAR (Polarbear Collaboration 2014), is brighter than the primordial BB spectrum (again see Fig. 1.6) at most scales and therefore must be accounted for in the analysis. Partial removal of this lensing signal, called delensing, has been demonstrated to improve constraints on r in a joint analysis between BICEP/Keck Array and SPTpol (BICEP/Keck Collaboration & SPTpol Collaboration 2021).
- Polarized Galactic foregrounds, primarily thermal dust and synchrotron emission, also generate B modes. These must be removed by making observations at multiple frequencies to isolate the signal from each component (dust, synchrotron, PGWs) separately.

All of these challenges together require targeted measurements of the polarized CMB at degree angular scales (where the primordial BB spectrum is expected to peak) at multiple frequencies with an instrument with very low statistical noise and exquisite control of systematics. The BICEP/Keck Array program has taken on this challenge, and produced world-leading constraints on r from B -mode measurements.

Fig. 1.7 shows a collection of published constraints and measurements of the CMB BB spectrum as of early 2021. The two most recent constraints from BICEP/Keck




Figure 1.7: Collection of published CMB BB spectrum measurements. Theory values are also shown for the lensing BB spectrum (solid red), primordial BB spectrum (dashed red), and the sum of both (dotted red). The recent BK18 data are shown after removing Galactic foregrounds, while the other points do not have any subtraction. Figure from BK18 (*The BICEP2 and Keck Array Collaborations 2021*).

Array are referred to frequently throughout this dissertation. The BK15 data set, which included all data from BICEP2 and *Keck Array* through the 2015 season, including observations at 95, 150, and 220 GHz, gave a 95 % confidence upper limit of $r < 0.09$ from B modes alone, down to $r < 0.07$ after including *Planck* (*The BICEP2 and Keck Array Collaborations 2018*). After including additional data through 2018, adding three years of 95 GHz data from BICEP3 and more *Keck Array* observations at 220 GHz, this constraint improves to $r < 0.034$ (*The BICEP2 and Keck Array Collaborations 2021*). A performance-driven forecast of BICEP/*Keck Array* sensitivity and $\sigma(r)$ constraints, which include realistic observing efficiency, is shown in Fig. 1.8. The projections are made by scaling real achieved performance in the BK14 data set (*The*




Figure 1.8: Projected sensitivity of ongoing and planned BICEP/Keck Array receivers. *Top:* Graphic showing the frequencies (color) and throughput (thickness) of individual telescopes. *Middle:* Map depths as a function of time for each frequency band, with the black X's denoting published achieved performance. *Bottom:* Sensitivity to r after marginalizing over the parametric foreground model. Future sensitivity heavily relies on delensing (dashed grey line). The dotted red line shows an idealized case with no foregrounds or lensing.

BICEP2 and Keck Array Collaborations 2016b). For more details on the forecasting framework, see Buza 2019.

1.4 THESIS OUTLINE

In this dissertation, we focus on measurements of the bandpass and beam response of BICEP/Keck Array instruments, the uncertainty on those measurements, and their systematic impact on the final cosmological parameter constraints. We begin in Chapter 2 by describing the BICEP/Keck Array experiment, going over each component of

the telescopes and the overall program strategy. In Chapter 3 we describe the spectral bandpass measurements made with a Fourier Transform Spectrometer, and assess the bias on r due to uncertainties in the bandpasses. We detail the various beam map measurement procedures and the far-field beam map analysis pipeline in Chapter 4, with a thorough analysis of the far-field beams in Chapter 5. In Chapter 6 we introduce specialized simulations that quantify the expected amount of $T \rightarrow P$ leakage in real BICEP/*Keck Array* data, including analysis of $T \rightarrow P$ leakage in the BK18 data set. We develop a set of simple, beam map-derived metrics and present optical models of small-aperture telescope beams in Chapter 7, followed by concluding remarks in Chapter 8.

Throughout this dissertation, we frequently refer to the following publications from the BICEP/*Keck Array* program:

- **BK I** – “BICEP2 I: Detection of B -Mode Polarization at Degree Angular Scales by BICEP2” ([The BICEP2 Collaboration 2014a](#))
- **BK II** – “BICEP2 II: Experiment and Three-Year Data Set” ([The BICEP2 Collaboration 2014b](#))
- **BK III** – “BICEP2 III: Instrumental Systematics” ([The BICEP2 Collaboration 2015](#))
- **BK IV** – “BICEP/*Keck Array* IV: Optical Characterization and Performance of the BICEP2 and *Keck Array* Experiments” ([The BICEP2 and *Keck Array* Collaborations 2015a](#))
- **BK V**, or the **BK13** Data Set – “BICEP/*Keck Array* V: Measurements of B -mode Polarization at Degree Angular Scales and 150 GHz by the *Keck Array*” ([The BICEP2 and *Keck Array* Collaborations 2015b](#))

- **BK VI**, or the **BK14** Data Set – “BICEP/*Keck Array* VI: Improved Constraints on Cosmology and Foregrounds from BICEP2 and *Keck Array* Cosmic Microwave Background Data with Inclusion of 95 GHz Band” ([The BICEP2 and Keck Array Collaborations 2016b](#))
- **BK VII** – “BICEP/*Keck Array* VII: Matrix based E/B Separation applied to BICEP2 and the *Keck Array*” ([The BICEP2 and Keck Array Collaborations 2016b](#))
- **BK X**, or the **BK15** Data Set – “BICEP/*Keck Array* X: Constraints on Primordial Gravitational Waves Using *Planck* WMAP, and New BICEP/*Keck Array* Observations through the 2015 Season” ([The BICEP2 and Keck Array Collaborations 2018](#))
- **BK XI** – “BICEP/*Keck Array* XI: Beam Characterization and Temperature-to-Polarization Leakage in the BK15 Data Set” ([The BICEP2 and Keck Array Collaborations 2019](#))
- **BK XIII**, or the **BK18** Data Set (in prep) – “BICEP/*Keck Array* XIII: Constraints on Primordial Gravitational Waves using *Planck*, WMAP, and BICEP/*Keck* Observations through the 2018 Observing Season” ([The BICEP2 and Keck Array Collaborations 2021](#))

As of the writing of this dissertation (end of April 2021), the BK18 publication, although very mature, is still in preparation. Many results from BK18 are shown in this dissertation, including final BB spectra and parameter constraints. At the expected time of release of this dissertation, BK18 will be published, and any differences seen in the results presented here and in the BK18 paper, if any, are expected to be very minor.

The BICEP/*Keck Array* 220 GHz frequency band is analyzed in great detail throughout this dissertation. Starting in 2016, there are two slightly distinguished bands

Band	Band Center	Observation Years (in <i>Keck</i>)	BK15 Nomenclature	BK18 Nomenclature
220	224 GHz	2016-2019	–	Coadded together to “220”
230	232 GHz	2015-2019	“220”	Coadded together to “220”

Table 1.1: Nomenclature of the two slightly distinguished frequency bands near 220 GHz. In most analyses presented here, the 220 and 230 GHz bands³ will be treated separately, except in the context of final power spectra and likelihood results for BK15 and BK18.

near 220 GHz which are denoted throughout this dissertation as 220 and 230 GHz. In most calibration analyses, these bands are treated separately. However, at the power spectrum stage of analysis of the real data (and all simulations), these two bands are coadded together into a single band which is then labeled as 220 GHz. See Table 1.1. When discussing final results for BK15 or BK18 (namely power spectra and likelihood results), the “220 GHz band” refers to all data in either band coadded together. This distinction will be made clear whenever necessary throughout the dissertation.

³The internal BICEP/*Keck Array* pipeline refers to these two bands as 210 and 220.

*The greatest mystery the universe offers is not
life but size.*

Stephen King, *The Dark Tower*

2

The BICEP / *Keck Array* Experiment

THE BICEP / *Keck Array* EXPERIMENT is a series of polarimeters at the Amundsen-Scott South Pole Station searching for the signature of primordial gravitational waves in the polarized CMB. The program has a rich history of projects, starting with BICEP1 in 2006 up to the present day, with both BICEP3 and BICEP Array currently

making observations. Each of these telescopes are designed to maximize sensitivity at degree angular scales, where the theoretical inflationary BB spectrum peaks. Other goals and design features that have remain unchanged throughout the experiment include:

- Use a compact, on-axis refractor design to maximize sensitivity at degree angular scales and minimize polarization systematics,
- Continuously observe a small, low-foreground patch of sky from a stable, dry site to produce deep, low-noise maps,
- Utilize boresight rotation and multiple levels of shielding to reduce beam systematics,
- Use a modular design to allow quick, easy detector replacement as upgrades become available.

These features of the program have led to world-leading constraints on the tensor-to-scalar ratio r derived from BB spectra. The success of these strategies also provide critical inputs to various design choices of future iterations of small-aperture telescopes (SAT), including CMB-S4 SATs (Carlstrom et al. 2019). A brief summary of the frequency coverage and detector count in deployed and future BICEP/Keck Array receivers is shown in Table 2.1. In this chapter, we consider in detail each component of these receivers, with emphasis on the growth in technology over the past few years of BICEP/Keck Array receivers.

Experiment	Years in Operation	Nominal # of Detectors	Notes
BICEP1 100 GHz 150 GHz 220 GHz	2006-2008	50 44 4	
BICEP2 150 GHz	2009-2012	512	First receiver in program to use TES detectors.
<i>Keck Array</i> 100 GHz 150 GHz 220/230 GHz 270 GHz	2011-2019	288 512 512 512	Five receivers each very similar to BICEP2.
BICEP3 100 GHz	2015 - present	2,560	Throughput is equivalent to all five <i>Keck</i> receivers
BICEP Array 30 / 40 GHz 100 GHz 150 GHz 220 / 270 GHz	2020 - present	192 / 300 4,056* 7,776* 8,112* / 12,288*	BA1 (30/40 GHz) deployed Jan 2020. BA2 and BA3 (100 and 150 GHz) planned for Jan 2022. BA4 (220/270 GHz) planned for Jan 2023.

Table 2.1: Brief deployment history of BICEP/*Keck Array* telescopes. *This is the projected detector count for BICEP Array receivers currently in development.

2.1 DETECTORS

Modern telescopes require a means of converting incident photons into an electric signal that can be processed into data. Along with the source signal, there is Poisson noise and all other sources of noise coming from imperfections of the detection method. For most CMB experiments, bolometers are used as the key detector elements in the focal plane units (FPU), though some experiments vary in the way they couple the incident light to the bolometers.

In general, a bolometer consists of some absorber with heat capacity C coupled to a thermal bath T_{bath} via a thermal link with conductance G . For any form of incident power P coupling to the bolometer, the bolometer heats up by an amount $\Delta T = \Delta P/C$, and that power is dissipated away with a time constant $\tau = C/G$. The

detectors used in BICEP1 were neutron transmutation-doped germanium (NTD Ge) thermistors. The light coupled to corrugated feed horns which then dumped power to the NTD Ge thermistors via an absorbing polarized mesh (Yoon et al. 2006). Starting with BICEP2, the BICEP/*Keck Array* program has used transition-edge sensors (TES) coupled to free space via a phased array of planar slot antennas, which are described below.

A TES bolometer is a superconductor which is held at its superconducting transition, the temperature where the resistance sharply drops to zero. Within this narrow transition region, the resistance varies sharply with temperature. Because the resistance curve has $dR/dT > 0$, it is voltage biased to provide negative electrothermal feedback – incident radiation increases the resistance, lowering the current traveling through the TES. The current supplied via the feedback loop to bring the TES back to its nominal resistance is a measure of the signal from that radiation (Irwin 1995).

For CMB science observations (low detector loading), the BICEP/*Keck Array* program (after BICEP1) uses titanium TES detectors, which have a superconducting transition at $T_c \simeq 500$ mK. For high-loading calibration measurements (such as FTS measurements in Chapter 3 or FFBM measurements in Chapter 4), we bias on the aluminum TES, with a transition temperature $T_c \simeq 1.2$ K, which is in series with the titanium TES. A detailed image of the TES is included in Fig. 2.1. One input line delivers the bias voltage to the TES island, and a separate line from the antenna network carries the radiation signal, depositing heat to the island via a lossy gold me-




Figure 2.1: Composite figure showing a 150 GHz detector tile in a BICEP2 / *Keck* style focal plane. a) A single 4" detector tile, housing 8×8 polarized detector pairs. b) Image of the back of the focal plane, with the niobium backshort removed. All BICEP2 and *Keck Array* focal planes hold four detector tiles. c) One detector pair, showing the phased array of planar slot antennas which couple light to the TES islands. d) One of the TES detectors (top left) and the band-defining microstrip filter (top right). e) Zoomed view of a thermally isolated TES bolometer, including the titanium and aluminum TESs in series, as well as the gold meander which acts as a termination resistor. Figure from [Kuo et al. 2008](#).

ander. The island heat capacity C varies with the thermal mass of the gold microstrip, and the thermal conductance G varies with the geometry of the legs that thermally isolate the island ([BICEP/Keck Array and SPIDER Collaborations 2015](#)).

Light is collected using a phased array of planar slot antennas, shown for one detector pair in the bottom left of Fig. 2.1 ([Kuo et al. 2008](#)). There are 288 antennas each aligned vertically or horizontally, defining the two polarization orientations. The

Receiver	2012	2013	2014	2015	2016	2017	2018	2019
BICEP2	D2 (150)							
Keck rx0	E1 (150)	E1 (150)	E6 (95)	E6 (95)	E10 (220)	E10 (220)	E10 (220)	E10 (220)
Keck rx1	E5 (150)	E5 (150)	E5 (150)	E1 (230)	E1 (230)	E1 (230)	E1 (230)	E4c* (150)
Keck rx2	E2 (150)	E2 (150)	E9 (95)	E9 (95)	E11 (220)	E11 (220)	E11 (220)	E11 (220)
Keck rx3	E7 (150)	D2 (150)	D2 (150)	E2 (230)	E2 (230)	E2 (230)	E2 (230)	E2 (230)
Keck rx4	E3b (150)	E4c (150)	E4c (150)	E4c (150)	E4c (150)	E12 (270)	E13 (270)	E13 (270)
BICEP3					B3 (95)	B3 (95)	B3 (95)	B3 (95)

Table 2.2: A history of the FPU’s deployed in the BICEP/*Keck Array* program between 2012 and 2019. FPU identifiers are listed, and the detector frequency in GHz is listed in parentheses. *E4c was retrofitted with microwave multiplexed readout and SLAC Microresonator Radio Frequency warm electronics (Cukierman et al. 2020).

incident light waves are collected and summed coherently and carried by planar microstrip lines, which then passes through a band-defining, 3-pole microstrip filter (with band width 25 % of the nominal band center). After passing through the filter (also shown in Fig. 2.1), the light is deposited onto the bolometer island.

Starting with BICEP2 the BICEP/*Keck Array* program has deployed a wealth of TES detector arrays, at frequency bands of 30, 40, 95, 150, 220, 230, and 270 GHz. Table 2.2 shows the history of these FPU’s for BICEP2, *Keck Array*, and BICEP3 between 2012 and 2019. Not shown is the 30/40 GHz FPU deployed in BICEP Array in the 2019-2020 austral summer. From this table alone, one can see the importance of high-frequency data to driving progress on constraining r , due to the need to measure and isolate the polarized dust signal.

2.2 DETECTOR READOUT

In order to read out the very small signals generated by hundreds or thousands of TES bolometers, we amplify and multiplex (MUX) the signals using superconducting quantum interference devices (SQUIDs). The SQUIDs used in the BICEP/*Keck Array* program are made by NIST (de Korte et al. 2003), which provide multiple stages of low-noise amplification. The SQUIDs are arranged for time-domain multiplexed readout, which greatly reduces the number of electrical wires at the sub-Kelvin stage, lowering the heat load on the detectors and the complexity of the feedthroughs needed at the cryostat interfaces.

BICEP2 and *Keck Array* use the “MUX07a” and “MUX09s” generation of NIST-developed SQUIDs, respectively. Each features three stages of amplification. Each TES bolometer is inductively coupled to a single first-stage SQUID (SQ1). 33 SQ1s, coupling to 16 detector pairs (leaving one “dark SQUID”) are fed to a single second-stage SQUID (SQ2) via a summing coil. Each SQ1 is sequentially turned off and on by sending a bias signal – when no bias is applied, that SQ1 remains superconducting and its readout signal bypasses the summing coil. See Fig. 2.2 for a simplified schematic. The SQ1s and SQ2s exist on a single MUX chip on the focal plane at 250 mK. 16 MUX chips are needed to read out a single BICEP2 / *Keck* focal plane. The primary difference between the MUX07a and MUX09s SQUIDs is a new gradiometric design which drastically reduces the pickup from external magnetic fields

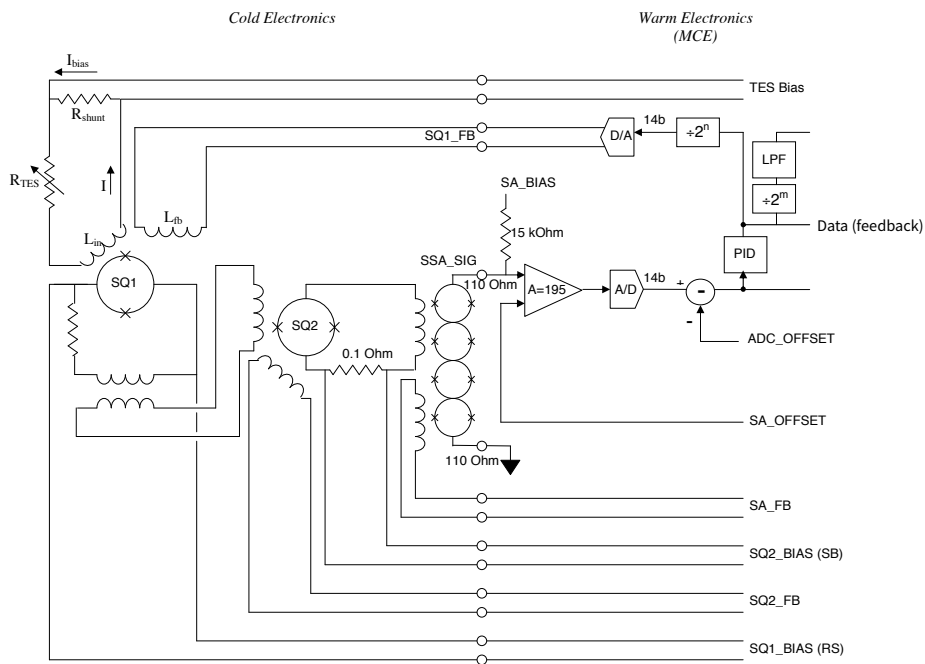


Figure 2.2: Block diagram of the time-domain SQUID-based multiplexed readout used by BICEP2 and *Keck Array*. The components within the cryostat (at 4 K or below) are in the left half of the diagram; components are room temperature are in the right half. The three stages (SQ1, SQ2, SSA) are visible. All biasing and read-out are provided by the MCE which is mounted directly to the cryostat. Image modified from the MCE Wiki: www.e-mode.phas.ubc.ca/mcewiki.

(Stiehl et al. 2011).

BICEP3 uses a newer “MUX11d” generation of SQUIDs, which removes the SQ2 intermediate stage, and takes advantage of flux-activated superconducting-to-normal resistive switches, called “row-selects” (Irwin et al. 2012). Each row-select is coupled to a SQ1 and is designed to switch at a critical current that is twice that of the SQ1, allowing it to share a bias line with the SQ1. Each row is addressed by applying a flux that drives one row-select to normal, while also biasing its companion SQ1. At the same time, the flux sent to all other row-selects in that column keeps them superconducting, shorting out the bias signal of all other SQ1s. This eliminates the need for the extra summing coil and SQ2 used in previous generations of time-domain SQUID readout.

The final stage of the SQUID amplification stage is SQUID series array (SSA), which is mounted on the fridge bracket at 4 K. The signal is carried from the FPU to the SSAs using superconducting niobium-titanium twisted-pair cables manufactured by Tekdata¹. The voltage response of a SQUID is periodic with magnetic flux Φ , so during standard observations we apply an active feedback loop to keep amplification in a linear regime. Each stage of the SQUID readout must be biased so that each amplifier is at its optimal $dR/d\Phi$ or $dV/d\Phi$ point, to ensure proper stability and linear amplification of the TES signal. This process of “SQUID tuning” must be done after every fridge cycle – for a review of the tuning procedure for *Keck Array* and BICEP3,

¹Tekdata: www.tekdata-interconnect.com

see [Kernasovskiy 2014](#) and [Grayson 2016](#), respectively.

2.3 OPTICS

All telescopes in the BICEP/*Keck Array* program are compact, on-axis refractors. The small aperture size allows for sensitivity at angular scales where the primordial B -mode spectrum is expected to peak ($\ell \sim 100$) while minimizing the size and complexity of the optics and other cryogenically cooled components. Keeping the optical components cold (4 or 50 K) drastically reduces the thermal load onto the detectors, improving noise performance. The compact size also allows for boresight rotation, which is critical for minimizing systematics in the CMB measurement.

A schematic of the optics used in BICEP2/*Keck Array*, BICEP3, and BICEP Array is shown in Fig. 2.3. There are a number of differences in the optics between these receivers, the most notable being the aperture size, 264 / 520 / 550 mm for *Keck*, BICEP3, and BICEP Array, respectively. BICEP Array and BICEP3 also feature faster optics and a wider instantaneous field of view than *Keck*. A summary of the optical parameters is shown in Table 2.3. In each receiver, the lenses are cooled to 4 K during observations, and various infrared (IR) filters are installed at 4 and 50 K. A low-loss window is used to hold vacuum. Below, we discuss the details of each element, with particular emphasis on BICEP3 and *Keck*, due to the focus on BICEP3 / *Keck* data throughout this dissertation.

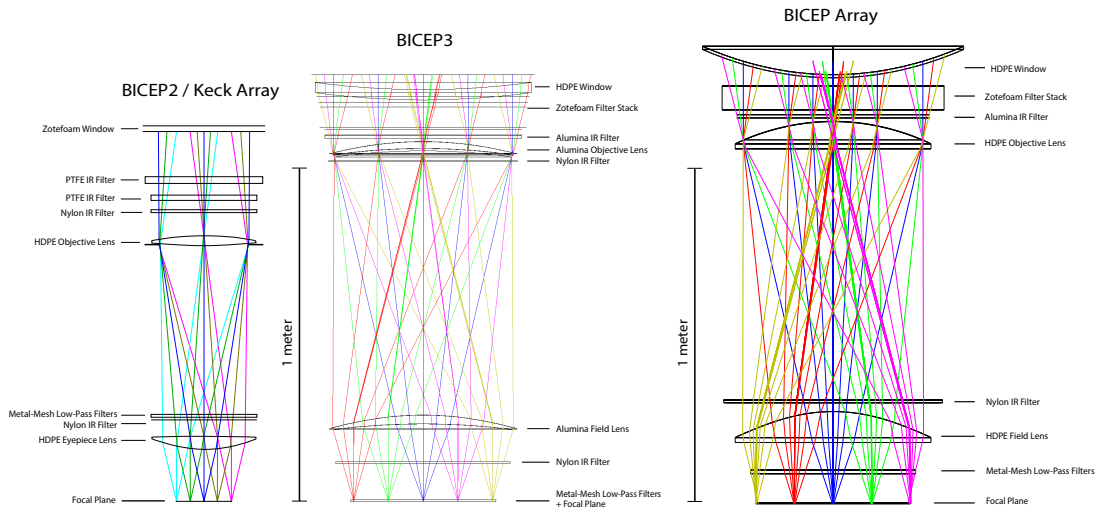


Figure 2.3: Schematic of the optics in BICEP2 / *Keck*, BICEP3, and BICEP Array telescopes. The individual optical elements are labelled and drawn roughly to scale. The BICEP3 Zotefoam filter stack replaced a metal-mesh reflective filter stack in the 2016-2017 austral summer. The BICEP3 window was replaced with a slightly larger BICEP Array-compatible window (also HDPE) in the 2018-2019 austral summer. The larger aperture and faster optics of BICEP3 result in a dramatically increased throughput compared to *Keck*. The optics of BICEP Array are similar to BICEP3 in size and throughput.

Experiment	Aperture Size	Field of View	f/#
BICEP2 / <i>Keck Array</i>	264 mm	15°	2.4
BICEP3	520 mm	27°	1.7
BICEP Array	550 mm	29.6°	1.57

Table 2.3: Optical parameters for BICEP2 / *Keck*, BICEP3, and BICEP Array.

The windows must be strong enough to hold vacuum, and not so thick that the amount of thermal emission raises the internal loading to unacceptable levels. For *Keck*, the windows are 32 cm in diameter and 12 cm thick, made of nitrogen-expanded polyethylene foam (Zotefoam² HD30) and bonded to the aluminum frame with Sty-cast 2850 epoxy. For BICEP3 the 67 cm window is instead made of a 3.2 cm thick high-density polyethylene (HDPE) plastic, which has higher structural rigidity than HD30. The BICEP3 window is anti-reflection (AR) coated with Teadit 24GRD³ expanded polytetrafluoroethylene (PTFE). In all receivers, a membrane of biaxially-oriented polypropylene (BOPP) is sealed above the window, which is pressurized by dry nitrogen gas, acting as an environmental shield and sublimating away any accumulated snow.

IR power enters the cryostat at ~ 35 W and ~ 100 W for *Keck* and BICEP3, respectively, indicating a strong need for IR filters to reduce both the loading on the focal plane and the work required by the pulse tube cooler to remove the heat. In *Keck*, all the IR filters are absorptive. There are two PTFE filters (0.5 inch and 1.3 inch thick) and a 3 mm-thick nylon filter at 50 K, as well as a 5.2 mm nylon filter at 4 K. For the larger IR load in BICEP3, aggressive filtering is critical to not overload the cryocooler. A stack of metal-mesh IR-reflective filters were installed at 50 K just below the window for the 2016 observing season. In the 2016-2017 austral summer,

²Zotefoam: www.zotefoams.com

³Teadit: www.teadit.com

this was replaced with a stack of 10, 1/8 inch thick HD30 foam layers each separated by 1/8 inch. This reduced the loading on the 50 K stage from 19 W to 12 W (Kang et al. 2018). In addition, BICEP3 also has a 10 mm alumina ceramic filter heat sunk to 50 K, and two nylon filters at 4 K, one below each lens. Finally, each receiver also uses a metal-mesh low-pass edge filter (Ade et al. 2006), with a cutoff nominally just above the high-frequency edge of the frequency band to reject blue leaks. This filter is located above the eyepiece lens in *Keck* or mounted directly to the detector modules in BICEP3. When deployed with BICEP3 in 2016, anomalous features in the band-pass measurements were eventually attributed to delamination of these filters. In the 2016-2017 BICEP3 upgrade, these delaminated filters were replaced, eliminating the spectral anomalies (see Section 3.2.2 and Kang et al. 2018).

The 4 K lenses are made of HDPE in *Keck* and 99.6% pure alumina ceramic in BICEP3. Alumina was chosen for BICEP3 due to its higher index of refraction, allowing a thinner size with less curvature needed compared to HDPE. Alumina also has higher thermal conductivity, preventing thermal gradients across the lens. A 26.4 cm (52 cm for BICEP3) aperture stop is mounted directly beneath the objective lens. It is made of Stycast-loaded HR-74 Eccosorb microwave absorber. The positioning and curvature of the lenses were optimized using Zemax⁴ optical ray-tracing software. The optimal focus is placed on an annulus of detectors at a median distance from the center of the focal plane.

⁴Zemax: www.zemax.com

2.4 CRYOSTAT

In order to measure very small cosmological signals with photon-noise limited TES detectors, we must minimize any non-cosmological radiation coupling to the detectors, including those emanating from within the telescope itself. To this end, we vacuum pump and cryogenically cool everything within the telescope enclosure, including the detectors, optics, and shielding architecture. BICEP2 used liquid helium to cool down to 4 K, but subsequent receivers moved away from liquid cryogens in order to eliminate the high cost and bulkiness associated with transportation and storage.

Here we focus on the *Keck Array* and BICEP3 cryostats, which are shown in Fig. 2.4. Both feature a 300 K vacuum jacket with enclosed shells at 50 K and 4 K. A Cryomech⁵ pulse tube cryocooler continuously cools the shells to 50 K/4 K, with the pulse tube head attached at the base of each receiver (within the “doghouse” enclosure). The different shells are thermally isolated from each other via carbon fiber supports at the bottom and titanium supports (so-called “boomerangs”) at the top. Between the 50 K and 4 K stage are 20 layers multi-layer insulation (MLI) which reduces radiative heat transfer. A layer of Amuneal Amumetal 4K⁶ is wrapped around the 4 K stage to reduce magnetic pickup in the SQUIDS.

The pulse tube cryocoolers cannot cool the detectors to the sub-Kelvin temper-

⁵Cryomech, Inc.: www.cryomech.com. *Keck* uses the PT410, while BICEP3 and BICEP Array uses a PT415.

⁶Amuneal: www.amuneal.com





Figure 2.4: *Keck* and BICEP3 cryostat architecture. All individual elements are labelled and roughly to scale (the two models are not to scale with each other - BICEP3 is twice the diameter of a *Keck* cryostat). Both use pulse tube cryocoolers to cool to 4 K and a three-stage sorption refrigerator to cool the detectors to operating temperature (~ 250 mK). The top image is from [The BICEP2 and Keck Array Collaborations 2015b](#) and the bottom is a modified version of that from [Grayson et al. 2016](#).

atures needed to bias the TESs on the superconducting transition. Each receiver houses a three-stage ($^4\text{He}/^3\text{He}/^3\text{He}$, colloquially called “Helium-10”) sorption refrigerator capable of cooling the detectors to near or below 250 mK. The *Keck* (and now BICEP Array) fridges are provided by Duband (Duband & Collaudin 1999) and the BICEP3 fridge is provided by Chase Research Cryogenics⁷. The fundamental process of the closed-cycle adsorption fridge is the adsorption and release of gas from active charcoal pumps, which are controlled via a series of resistive switches and heaters. ^4He is condensed by heating the charcoal, releasing adsorbed ^4He gas, raising the pressure and boiling point. Upon contacting a “condensation point” which is heat sunk to a sufficiently low temperature, the ^4He is condensed. Liquid ^4He collects into a reservoir via gravity, and the external heating of the charcoal is disabled, “pumping” on the ^4He and lowering the pressure and temperature (~ 1.4 K). The ^4He reservoir pre-cools the neighboring ^3He (inter-cooler, or IC) stage via a thermal link, and a similar procedure lowers the IC stage to 350 mK. Finally, the IC stage is linked to and cools the ultra-cold (UC) stage, which condenses ^3He at roughly 230 mK. The UC stage is heat-sunk to the focal plane through a stainless steel block, which acts as a passive thermal filter. The IC stage is heat-sunk to a niobium cylinder (the so-called “spittoon”) which provides further magnetic shielding around the focal plane.

⁷Chase Research Cryogenics Ltd: www.chasecryogenics.com

2.5 TELESCOPE MOUNT AND OBSERVATORY SITE

The BICEP/*Keck Array* program operates from two separate buildings at the Amundsen-Scott South Pole Station, which are separated by ~ 200 m. The Martin A. Pomerantz (MAPO) building currently (2020+) holds BICEP Array, after housing *Keck Array* from 2011-2019. It was also the location of DASI (2000-2003) and QUAD (2005-2007). The Dark Sector Laboratory (DSL) has hosted BICEP3 since 2015, after previously holding BICEP1 (2006-2008) and BICEP2 (2008-2011). Fig. 2.5 shows a photo of both buildings, with a view of the BICEP Array mount in MAPO and the BICEP mount in DSL. The mount structures and shielding architecture, as described below, are very similar between the two.

Both telescope mounts allow motion around three axes: azimuth, elevation, and boresight rotation (colloquially called the “theta” or “deck” axis, DK for short). Azimuth (AZ) and elevation (EL) dictate the point being observed on the sky, while boresight rotation changes the projection of the beams on the sky, allowing us to measure different polarization components on the same point on the sky and to perform systematics checks.

Even at a flat and cold (~ 250 K) observing site like the South Pole, light exiting the telescope aperture can and will couple to the ground and contaminate the detector signal, if unmitigated. For this reason, all BICEP/*Keck Array* receivers use two levels of warm, external shielding – a co-moving forebaffle attached directly to the




Figure 2.5: Photo of MAPO and DSL at Amundsen-Scott South Pole Station. The nearest building is MAPO, with the ground shield surrounding the BICEP Array mount (formerly the DASI mount) clearly visible. The far building to the right is DSL, with the BICEP3 ground shield (right side of building) and 10 m dish of the South Pole Telescope (left side of building) visible.

telescope aperture, and a stationary ground shield around the base of the mount. The fore baffles are absorptive, lined with sheets of Eccosorb (HR-10 for *Keck*, a combination of HR-25 and AN-75 for BICEP3), which controls the response from extended sidelobes. The upper lip of the forebaffle is rolled outward to prevent diffraction over a hard edge. The shape of the forebaffle is designed with two key constraints in mind – the maximum opening angle of the receiver must be unobstructed, and there must be no line of sight connecting any point on the window with the edge of the ground shield when at the lowest CMB observation elevation.

The ground shield around the base of the mount is made of aluminum (BICEP3) or plywood coated in aluminum tape (*Keck*/ BICEP Array), which reflects any rays that

have diffracted off the forebaffle to cold sky. Overall, this enforces the criterion that a ray must diffract twice in order to couple to the ground. The ground shields, like the forebaffles, have a curved lip at the edge. In practice, some azimuth-fixed signal is still seen in the detector timestreams, which is filtered away on one-hour (per-scanset) timescales.

Along with the flat, cold local environment, observing from the South Pole provides a number of considerable advantages, both scientific and programmatic. This site has a rich history of CMB science over the past few decades, proving to be one of the most ideal locations for millimeter and sub-millimeter telescopes (Kovac & Barkats 2007). The South Pole sits atop the Antarctic Plateau, at an elevation of roughly 2800 m, leading to a very dry atmosphere. The lack of precipitable water vapor yields relatively low atmospheric opacity, and the stability over the eight-month winter reduces atmospheric noise. Additionally, from the South Pole we can observe the same patch of sky as the Earth rotates, allowing continuous integration from a fixed elevation angle. The station is run by the United States Antarctic Program (USAP), which provides infrastructure, housing, and power needed to maintain the telescope and crew. A series of satellite uplinks allows for daily transfer of telescope data, which allows regular data monitoring and expedites any debugging that may be necessary during standard observations.

2.6 OBSERVATION STRATEGY

As discussed so far in this chapter, the BICEP/*Keck Array* experiments have grown in size and complexity as the technology has improved over the years. However, due to the proven success of continuous observation from the South Pole over the past 10-15 years, the observing strategy has remained largely the same. We observe the same 400-600 sq. deg. patch of sky, which varies with the size of the instantaneous field of view of the experiment (see Table 2.3). The patch, centered at (RA = 0 hr, dec = -57.5°), lies in the “Southern Hole,” which is far from the galactic plane and has very low polarized foregrounds.

Various program-specific terminology is used to describe discrete “units” of observations at different timescales. This terminology and their associated definitions are listed below:

- **Elnods:** These are elevation nods that scan up-down-up (or down-up-down) in elevation at a fixed azimuth, with a throw of 1.28°. These take place before and after each scanset (see below) and are primarily used for relative gain calibration.
- **Halfscans:** This is a single scan of the telescope 64.4° in one direction in azimuth, at a constant elevation. The speed of each scan is 2.8°/sec.
- **Scansets:** This is a collection of roughly 50 pairs of halfscans (all at fixed elevation). These are bracketed on each end by elnods, and take 50 minutes to complete. These are colloquially referred to as “tags.”
- **Phases:** These are collections of 6-10 scansets, where the elevation and azimuth are slightly adjusted after every pair of scansets to account for the sky rotation.

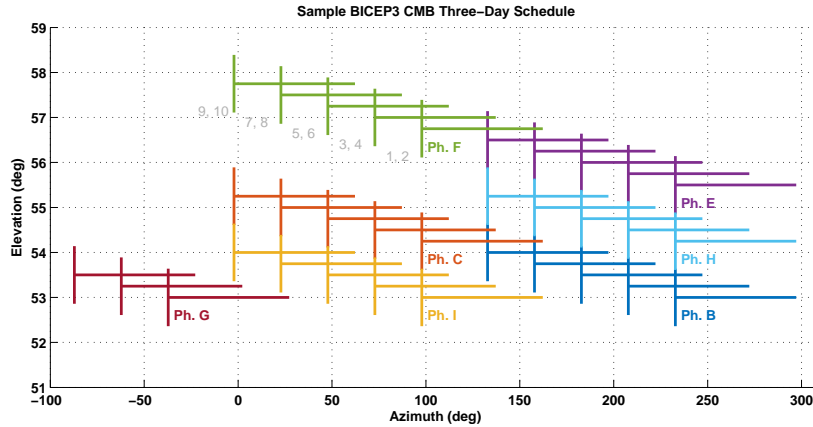


Figure 2.6: Example of a standard BICEP3 CMB schedule. The vertical stripes mark the elevation motion from the elevation nodes that bracket each scanset, and the horizontal stripes denote the azimuth extent of that scanset. Each colored set corresponds to a single phase, which is labeled below the first scanset. The scanset ordering is labeled for Phase F (the pattern is the same for each phase, except for Phase G which only has six scansets).

- **Schedules:** A collection of multiple phases of data, all at a constant boresight rotation angle. The azimuth and elevation are changed between phases to expand sky coverage. A schedule takes roughly 2-3 days to complete (however long it takes to deplete the liquid helium in the UC stage of the sorption fridge). The fridge is cycled (usually twice, back-to-back) and the boresight rotation angle is changed between schedules.

A typical BICEP3 three-day CMB schedule is shown in Fig. 2.6. The azimuth range of each scanset and the elevation range of each elnod is shown by the horizontal and vertical lines, respectively. A standard *Keck* CMB schedule from 2016 and onward is very similar to this BICEP3 schedule, except the phases only go through phase F (BICEP3 has a longer fridge hold time by one day, allowing more phases). Prior to 2016, the azimuth and elevation were adjusted after every single scanset, instead of every pair of scansets. Such an observing pattern can be seen in [The BICEP2 Collaboration](#)

2014b.

From 2013 on, *Keck Array* observed the CMB over eight evenly-spaced DK angles (68° , 113° , 158° , 248° , 293° , 338° , 383°). BICEP Array, which replaced *Keck Array* in the austral summer of 2020-2021, also uses these same DK angles. Due to mechanical limitations of the pulse tube cooler in the mount, BICEP3 cannot achieve full 360° rotation and therefore only observes at four DK angles (23° , 68° , 203° , 248°). In principle, this could reduce the beneficial cancellation of certain temperature-to-leakage modes compared to the eight-DK angle strategy, but for BICEP3 this effect seems small (see Section 5.7).

Infinity itself looks flat and uninteresting. Looking up into the night sky is looking into infinity – distance is incomprehensible and therefore meaningless.

Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

3

Bandpass Measurements and Systematics

IN ORDER TO CONFIDENTLY DETECT primordial B modes from the CMB, we must be able to identify and separate other signals coming from polarized sources between us and the CMB. The most prominent foregrounds at degree angular scales – polarized dust and synchrotron emission – have spectral profiles that are different from each

other and from the CMB. This means observations at many different frequencies are required to perform foreground component separation. It is therefore imperative that we have precision measurements of the end-to-end spectral response (bandpass) of our receivers.

Bandpass measurements also provide valuable feedback on other possible systematic contamination in our polarization data. Before and after each scanset (50 minutes of CMB observations), we perform elevation nods in order to calibrate the relative gain of the detectors. This is not a perfect method – differences in bandpass shapes between two co-located detectors will couple to the difference between the atmospheric and the CMB spectral emission, resulting in temperature-to-polarization leakage. This effect is accounted for in deprojection, which is discussed in detail in Section 6.1.2. However, since deprojection relies on an externally-provided temperature map, leakage could still arise from differences between the external temperature map and the true temperature sky. Bandpass measurements can be used to quantify this effect – previous work has shown this to be small¹, but this could become significant in future CMB experiments.

As a ground-based experiment, we must choose our frequency bands such that they optimize foreground separation while also avoiding contamination from the atmosphere, including oxygen and water spectral lines. Even a small sensitivity to such lines could increase detector loading to unacceptable levels. With the addition of 30

¹In the case of unpolarized CO emission, for example, the expected leakage is $r < 10^{-4}$.

and 40 GHz frequency bands in BICEP Array (Moncelsi et al. 2020), the BICEP/Keck Array experiment now observes at six different frequencies spanning all available transmission windows between 0 and 300 GHz, above which the atmospheric transmission decays until reopening at infrared and optical frequencies.

Whenever a new detector array is deployed for CMB observations, we perform *in-situ* measurements of the end-to-end instrument spectral response using a field-deployable Fourier Transform Spectrometer (FTS). In this chapter, we discuss the measurement procedure, analysis methods and results, and the impact of systematic and statistical uncertainty of these bandpass measurements on the final cosmological parameter estimates.

3.1 FOURIER TRANSFORM SPECTROMETER

3.1.1 THE INTERFEROMETER

Our FTS is a Martin-Puplett interferometer (Martin & Puplett 1970), which uses a polarized wire grid to split the input beam. The input port (in a time-reverse sense) is a rectangular aperture 5 inches wide. A polarized input grid couples the input signal to the interferometer axis. The signal is then split by a polarized beamsplitter, a free-standing wire grid with the polarization axis rotated 45° with respect to the input grid. The two interferometer arms hold a stationary roof mirror and a second

roof mirror on an Velmex BiSlide² automated translation stage. We use an Emix23 magnetic tape encoder on the moving arm for high-precision recording of the position of the moving mirror. These roof mirrors rotate the polarization angle of the split beam by 90°, allowing the beams to recombine after passing through the beamsplitter in the reverse direction. Once recombined, an off-axis parabolic mirror redirects the light to a polarized source wire grid, which splits the beam to ambient Eccosorb HR-10 temperature microwave absorber (300 K, or 250 K if *in-situ* at the South Pole) or a container of LN2-soaked HR-10 (77 K). The difference between the two source temperatures determines the strength of the interferometer signal. See Fig. 3.1 for a photo of the FTS with each component highlighted. For a detailed calculation on the electromagnetic signal propagation through a Martin-Puplett interferometer, see [Lesurf 1990](#).

3.1.2 COUPLING TO THE RECEIVER

Because the FTS aperture is small compared to the size of the beam coming out of the receiver aperture, and because the beams have significantly diverged between the 4 K aperture stop and the FTS aperture, the FTS must be able to translate while mounted to the receiver in order to couple to all of the detector beams while minimizing the frequency-dependent effect of beam truncation. We use a pair of parallel-mounted Velmex BiSlide automated translation stages to allow free translation in two

²Velmex: velmex.com/Products/BiSlide




Figure 3.1: Image of the FTS at CalTech for pre-deployment lab testing, with individual components highlighted. The detector beams are coupled to the interferometer via the input grid, and then split by the polarized wire grid beamsplitter in the center of the figure. The difference in path length between the two arms is modulated by the moving mirror in the upper arm. After reflecting off the roof mirrors, the beam is recombined at the beamsplitter and redirected via an off-axis parabolic mirror. The beam is split again at the source grid, and terminates on either ambient temperature absorber (~ 300 K, below the source grid in the figure) or absorber soaked in LN2 (~ 77 K). The difference in temperature between the two sources determines the strength of the interferometric signal.

dimensions. The FTS rests atop this stage, which is then mounted to the receiver aperture. The input grid also has freedom to rotate in two dimensions, by using a goniometer and rotation stage to fine-tune the coupling between the FTS and detector beams entering the FTS at different angles. The degrees of freedom provided by the XY translation stage and the input grid rotation allow for optimal illumination of any subset of detectors within the receiver focal plane. However, there are still systematic uncertainties associated with non-idealities in this coupling – this is further discussed in Section 3.3.1.

3.1.3 MEASUREMENT PROCEDURE

FTS measurements are done on all detector arrays after they are installed in a receiver for CMB observations. This is usually done at the South Pole, either in the lab (e.g. the MAPO building) or after the receiver has been installed in the mount. Fig. 3.2 shows an image of the FTS mounted on a *Keck* receiver in the mount.

There are a number of operational choices to make when taking FTS data, including the speed of the moving mirror, the number of "half-scans" made with the moving mirror, and the rotation angle of the FTS with respect to the detector polarization axes. The nominal choices and reasoning for each of these is discussed below:

- *Moving mirror speed*: the mirror speed determines the Nyquist frequency f_{Nq} , which is half the sample rate f_s . The sample rate of our measurements is $f_s = f_{GCP} c / (2v_{\text{mirror}})$, where f_{GCP} is the GCP "fast mode" data sampling rate,




Figure 3.2: Image of the FTS mounted on a *Keck* receiver in the DASI mount. The FTS can be translated in two dimensions to optimize the illumination of any subset of detectors in the focal plane. Measurements are usually taken at multiple rotation angles (with respect to the stationary receiver) as a systematics check.

which is ~ 150 Hz for BICEP3 and ~ 180 Hz for *Keck*. We nominally use 2 mm/s as the moving mirror speed, which gives a generous $f_{Nq} \sim 5$ THz.

- *Number of half-scans*: a half-scan in this context is the travel of the moving mirror from one end to the other. We normally take four half-scans – two in each direction. This allows us to check for scan direction-dependent systematic effects and for repeatability between scans in the same direction. More half-scans is more time-consuming, but allows for better averaging down of noise.
- *FTS rotation angle*: the FTS can be mounted such that the input wire grid is aligned parallel/perpendicular or at an angle 45° with respect to the detector polarization axes. We normally choose 45° so that we can measure both polarization detectors equally, but we also take a second measurement at 135° or 225° as a systematics check.

The spectral resolution of the FTS, $\Delta f = c / \Delta L$ is set by the maximum achievable path difference ΔL between the two arms of the interferometer. In our case, $\Delta L \sim 20$ inches, which gives a resolution of $\Delta f \sim 0.5$ GHz. This could be improved by bringing the stationary mirror closer to the beamsplitter – however, we choose to keep the stationary mirror at a distance roughly halfway between the minimum and maximum moving arm range. This allows us to have two-sided interferograms in each half-scan, providing a valuable check on the interferometer alignment. Moving the stationary mirror closer would improve Δf by a factor of 1-2 but will lead to one-sided interferograms in the data.

With all the hardware parameters chosen, the measurement procedure is straightforward. The FTS is mounted onto the receiver aperture, the LN2 container is filled, and the detectors are biased on the Al transition. A set of sample detectors is illu-

minated, and interferograms are monitored in real time while the alignment of the interferometer is fine-tuned. Once the sample interferometers are maximally symmetric, we then take data over a series of FTS “pointings,” where one pointing illuminates a grid of roughly 4x4 detector pairs. At each pointing, the desired number of half-scans are taken, and the XY translation stage and input grid goniometer and rotations stage are adjusted for the next pointing. For 4” x 4” wafers (*Keck*, BICEP3) we typically do five pointings per tile. The process is repeated for the desired number of FTS rotation angles. In principle, this process can be fully automated (up to an occasional LN2 refill), since all the motorized stages are electronically controlled. We normally choose to have 1-2 people monitoring the interferograms in real time to ensure data quality.

3.2 ANALYSIS & RESULTS

3.2.1 FROM INTERFEROGRAMS TO SPECTRA

The raw data take the form of per-detector interferograms, and a simultaneous timestream of the moving mirror encoder. The goal is to obtain high signal-to-noise spectra that represents each detector’s response to a beam-filling source with a Rayleigh-Jeans spectrum.

For each detector, the interferograms are identified, polynomial filtered, and windowed. The polynomial filter is usually third order, and only necessary for data taken

on the mount – thermal drifts in the ambient load temperature lead to drifts in the interferograms which are at a much lower frequency than the detector bandpass, but interfere with the algorithm that identifies the white light fringe (WLF). The WLF is identified by choosing a point near the expected true value (which is known due to the repeatability of the encoder), then calculating the interferogram symmetry at each point in a small neighborhood around the initial guess. The point with maximal symmetry is taken to be the WLF. The windowing is done after the identification of the WLF. Windowing comes at a cost of spectral resolution, but is essential for a smooth interferogram with reduced “ringing” in the spectral domain. We typically use a Blackman-Harris window, though a triangle window has been used for the lower-frequency detectors.

A unitary Fourier transform is applied to the interferograms, and the real and imaginary parts are calculated and kept separately, as a secondary means to check the impact of interferogram misalignment on the spectra. The spectra from all half-scans are averaged together for each detector to form one spectrum per FTS pointing and rotation angle. The optimal pointing is the one that yields the highest signal-to-noise spectrum (although the pointing with the most symmetric interferogram can also be used, and often corresponds to the highest S/N spectrum). All optimal per-detector spectra can then be averaged together to find the receiver-averaged bandpass.

At this point a correction is applied to account for the fact that the FTS aperture is not beam-filling. The correction is done in frequency space by dividing the spectra

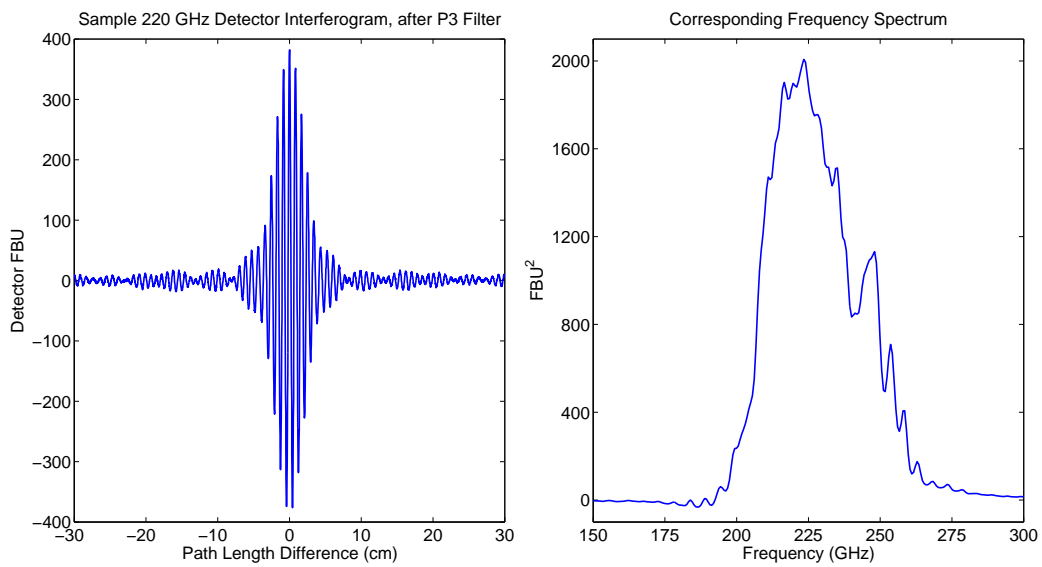


Figure 3.3: Example interferogram (left) and spectrum (right) for a 2017 *Keck* rx0 detector at 220 GHz. The interferogram is in raw detector feedback units (FBU) with only a third-order polynomial filter applied. The spectrum is a result of taking a Fourier transform of the interferogram after applying a window function. The aperture correction has not yet been applied.

by f^a , where a depends on both the FTS aperture size and the nominal beam size at the receiver aperture – for all BICEP/*Keck Array* bands, $1 < a < 2$. Full discussion on the calculation of this correction and the associated uncertainty is shown in Section 3.3.1.

3.2.2 ARRAY-AVERAGED SPECTRA AND THEIR APPLICATIONS

Fig. 3.4 shows the BICEP/*Keck Array* array-averaged spectra overlaid with the South Pole transmission spectrum. The spectra have had the aperture correction applied, meaning they are the response to a beam-filling source with a Rayleigh-Jeans spectrum. This figure also demonstrates the similarity in the *Keck* 220 and 230 GHz bands – in most analyses (including most work in this dissertation), these two bands are treated entirely separately, though in the primary CMB analysis, they are merged together at the power spectrum stage. This merge has been shown to have negligible effect on the final cosmological parameter estimation.

Our multicomponent likelihood analysis takes all BICEP/*Keck Array* maps and external maps (such as *Planck* and WMAP), and evaluates the likelihood of a given CMB + foreground model given all auto- and cross-spectra derived from those maps. The array-averaged spectra are used in the calculation of the expected signal of each component map. More explicitly, the signal from a given component (e.g. polarized dust) is scaled from a nominal frequency to the desired bandpass, given some frequency scaling law for that component (for dust this is usually a greybody spectrum).




Figure 3.4: *BICEP/Keck Array* array-averaged spectra overlaid with a South Pole atmospheric transmission spectrum typical of a winter day (black). The plotted bands are peak-normalized, and include BICEP Array 30 and 40 GHz (dark red), BICEP3 95 GHz (red), and *Keck* 150 GHz (green), 220/230 GHz (light/dark blue) and 270 GHz (magenta). The spectra have been corrected such that they are the response to a beam-filling, Rayleigh-Jeans source. Bands not included in this plot are the dual-band BICEP Array 30/40 GHz tile, and the *Keck* 95 GHz spectrum (very similar to BICEP3 95 GHz). In most analyses, the 220 and 230 GHz bands are treated separately, though in the primary CMB analysis they are merged together in the maps \rightarrow power spectra stage. The atmospheric spectrum was calculated using *am* (Paine 2019).

Therefore an error in the estimation of the bandpass leads to an error in the predicted signal for the component, which has immediate impact on the cosmological parameter estimation. This is discussed further in Section 3.3.3.

The averaged spectra can be used to identify problems in the receiver optical chain or in the detectors themselves. One distinct example of this is the diagnosis of delaminated low-pass edge (LPE) filters in the BICEP3 2016 observing season. As shown in Fig. 3.5, the BICEP3 spectra showed a series of “spikes” and “dips” that varied from module to module. Since these features were invariant with FTS pointing or rotation angle, they were determined to be optically real. These fringes in the spectral domain corresponded to a physical distance that was approximately the same as the

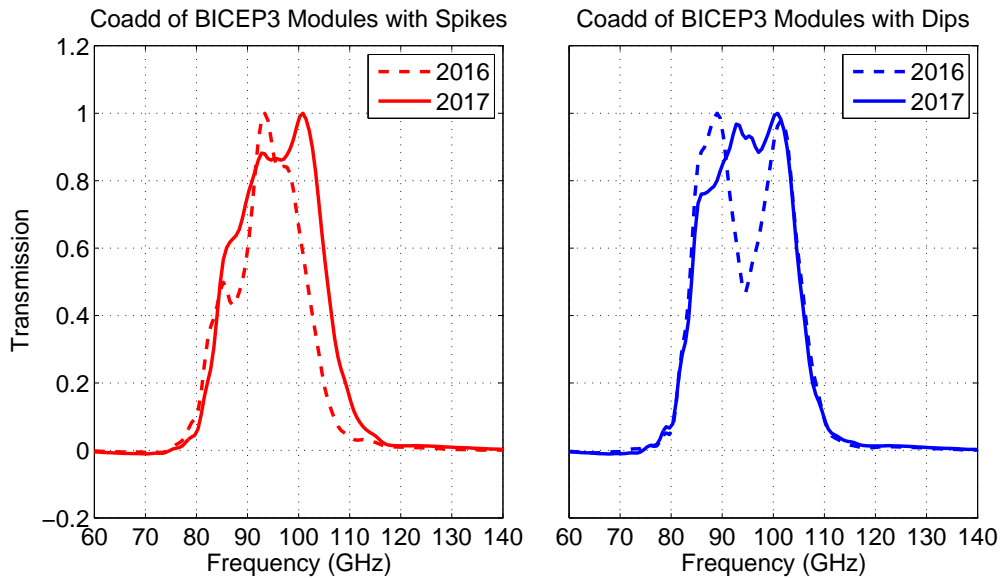


Figure 3.5: Peak-normalized spectral response of BICEP3 in 2016 (dashed) and 2017 (solid), separately coadded to include modules that demonstrated "spikes" (left) and "dips" (right) in the 2016 spectra. These features in the 2016 spectra were consistent across different FTS pointings and rotation angles, and were eventually traced to interactions with optical cavities formed by delaminated low-pass edge filters just above the detector modules. These features were eliminated in the 2017 spectra after replacing the delaminated filters in the 2016-2017 austral summer.

distance between the detectors and the LPE filters. Upon opening the receiver in the 2016-2017 austral summer, delaminated filters were indeed found above most of the modules. After replacing these filters and cooling back down, the spectra showed no fringing features.

3.2.3 BANDPASS PARAMETERS

For each per-detector spectrum $S(\nu)$, we compute the band center $\langle \nu \rangle$ and the bandwidth $\Delta\nu$. The band center is defined as

$$\langle \nu \rangle = \int \nu S(\nu) d\nu, \tag{3.1}$$

and the bandwidth is defined as

$$\Delta\nu = \frac{(\int S(\nu) d\nu)^2}{\int S(\nu)^2 d\nu}. \tag{3.2}$$

These statistics are accumulated for each detector, and can be used to check for 1) patterns that depend on the detector position in a tile for FPU or 2) systematics that depend on the hardware or measurement, such as half-scan direction or FTS rotation angle. Table 3.1 shows the band centers and bandwidths over detectors representative of each of the BICEP/*Keck Array* observing bands.

3.3 UNCERTAINTIES AND IMPACT ON LIKELIHOOD ANALYSIS

Non-idealities in the FTS instrument and measurement procedure lead to systematic uncertainties associated with the resulting spectra and spectral parameters. In this section, we use both modeling and real data to examine how uncertainty in the coupling, alignment, and overall repeatability can change the shape of the spectra or shift

Band	Data Sets Used In Array Average	Band Center (GHz)	Band Width (GHz)
Keck 95	Keck Rx0+2 on Mount Jan 2014	95.76 ± 0.74	25.37 ± 1.25
Keck 150	Keck Rx1+4 on Mount Jan 2014	153.04 ± 1.94	41.71 ± 1.11
Keck 220	Keck Rx0+2 on Mount Jan 2017	224.39 ± 3.28	53.93 ± 4.65
Keck 230	Keck Rx1+3 in MAPO Jan 2015	232.32 ± 2.06	49.18 ± 2.64
Keck 270 (FPU E12)	Keck Rx4 in MAPO Jan 2017	277.69 ± 2.11	67.56 ± 3.74
Keck 270 (FPU E13)	Keck Rx4 in MAPO Jan 2018	276.74 ± 2.19	67.53 ± 1.77
BICEP3 (before LPE filter swap)	BICEP3 on Mount Jan 2016	94.38 ± 1.31	23.40 ± 2.41
BICEP3 (after LPE filter swap)	BICEP3 on Mount Jan 2017	95.90 ± 1.52	25.73 ± 1.27
BA1 30	BA1 in MAPO Jan 2020	29.67 ± 0.98	7.61 ± 0.25
BA1 40	BA1 in MAPO Jan 2020	40.07 ± 1.61	12.12 ± 0.92
BA1 30 Dualband	BA1 in MAPO Jan 2020	30.63 ± 1.69	16.94 ± 0.83
BA1 40 Dualband	BA1 in MAPO Jan 2020	38.79 ± 2.63	13.75 ± 1.21

Table 3.1: Band centers and bandwidths for each of the BICEP/Keck Array observing bands, presented as (Mean ± Scatter) over all detectors in the listed data sets. The detector scatter is not indicative of the measurement uncertainty, which is discussed in Section 3.3.

the band center values. We then implement band center uncertainties in the multi-component likelihood analysis and quantify their effect on the estimate of the tensor-to-scalar ratio r .

3.3.1 COUPLING TO THE FTS

The power P absorbed by one single-moded, polarization-sensitive detector is given by

$$P = \frac{1}{2} \int d\nu \eta(\nu) A_{\text{eff}}(\nu) \int \int d\theta d\phi B(\nu, \theta, \phi) I(\nu, \theta, \phi), \quad (3.3)$$

where η is the dimensionless optical efficiency, the frequency-dependent fraction of photons coupling to the detector; A_{eff} is the receiver effective area (m^2); B is the dimensionless peak-normalized spatial angular response function (the beam); I is the source spectral radiance per unit frequency ($\text{W m}^{-2} \text{sr}^{-1} \text{Hz}^{-1}$), which is assumed to be unpolarized (hence the factor of 1/2). We want to relate the spectrum we measure

with the FTS to $\eta(\nu)$. If we assume for the moment that the source is beam-filling, and has a Rayleigh-Jeans (RJ) spectrum $I(\nu) \sim \nu^2$, then Eq. 3.3 becomes:

$$P = \frac{1}{2} \int d\nu \eta(\nu) I(\nu) A_{\text{eff}}(\nu) \Omega_{\text{beam}}(\nu) \sim \frac{1}{2} \int d\nu \eta(\nu) \nu^2 [A\Omega](\nu) \quad (3.4)$$

where $\Omega_{\text{beam}} = \int \int d\theta d\phi B(\nu, \theta, \phi)$. We have introduced the etendue, or throughput, $A\Omega$, and here define it to be $A\Omega = F(\nu)\lambda^2$, where $F(\nu)$ is some dimensionless coupling efficiency. In the case of a truly beam-filling RJ source, $F(\nu) \sim \text{const.}$ and we find that the recovered FTS spectrum is just $\eta(\nu)$, up to some frequency-independent normalization.

However, the FTS aperture is 12 cm wide, sitting atop a *Keck* receiver aperture with a diameter of 26 cm (or 52/55 cm, for BICEP3 / BICEP Array), meaning the FTS is not beam-filling. The fraction of the total beam power that couples to the FTS depends on the frequency, so we effectively have $F(\nu) \sim \nu^a$, where a depends on the size of the beam coming out of the receiver aperture and how well the beam couples to the FTS aperture.

A nominal value for the aperture correction factor a can be found via modeling. Starting with a Gaussian beam at the plane of the FTS aperture (whose beam width is taken from near-field beam map measurements from that receiver), we calculate the coupling fraction given the FTS aperture size. We then perturb the beam frequency, changing the beam width by a known amount, and re-calculate the coupling fraction.

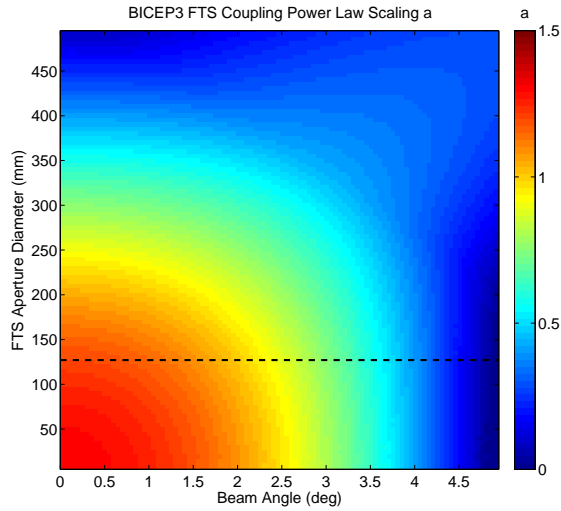


Figure 3.6: BICEP3 FTS coupling aperture correction factor a as a function of FTS aperture diameter and off-axis beam angle. The true diameter of the FTS aperture size (127 mm) is marked by the dashed line. As the FTS diameter approaches the BICEP3 aperture diameter (520 mm), the correction factor approaches zero, which is expected when there is no frequency-dependent truncation. For an on-axis beam ($\theta = 0$), the correction factor predicted by this model is $a = 1.20$.

From the ratio of the coupling fractions, we can extract a . This can also be done for off-axis beams – beams that slightly miss the FTS aperture by some angle θ will have a different frequency-dependent coupling than beams that perfectly align with the FTS aperture. The result of this model applied to BICEP3 is shown in Fig. 3.6. For a given frequency band we use a single value of a to correct all spectra.

This model can also be used to predict a measurement uncertainty associated with possible off-axis beam truncation. Each FTS pointing illuminates a roughly 4x4 grid of detector pairs, so most detectors will not have perfect on-axis coupling. We can compute the maximum angle a detector may be off-axis within a given point, then compute the error in band center associated with using the on-axis value of a . The

Band	Nominal (on-axis) Correction Factor a	Band Center Error for $\theta = 2.5^\circ$
Keck 95	0.91	0.5 %
Keck150	1.03	0.6 %
Keck 220	1.03	0.5 %
Keck 230	1.07	0.4 %
Keck 270	1.08	0.5 %
BICEP3 95	1.20	0.2 %
BA1 30	1.03	0.3 %
BA1 40	1.23	0.4 %

Table 3.2: List of nominal FTS aperture correction factors for each BICEP/Keck Array band, and band center uncertainties for a detector beam that is off-axis from the FTS aperture by $\theta = 2.5^\circ$. That value of θ is chosen based on the spacing of FTS pointings in a standard measurement. These numbers are predicted from a model assuming a Gaussian beam at the receiver aperture with FWHM taken from real near-field beam map data. The array-averaged spectra are divided by ν^a before being used in the multicomponent likelihood analysis.

nominal (on-axis) aperture correction factors and band center uncertainties from off-axis coupling are shown in Table 3.2. The predicted uncertainties from this model are $\lesssim 0.6\%$

This model is not perfect – the main shortcoming is that it assumes the beam still couples perfectly to the interferometer axis, even if it gets truncated at the FTS aperture at a strong angle. In reality, we expect many detector beams will travel through the interferometer at some angle, which can lead to more frequency-dependent distortions in the resulting spectrum if the deviation from the interferometer axis is significant compared to the wavelength. We can explore the impact of this possible variation by looking at the band centers derived from different FTS pointings for the same detector. We take the difference between the band center of a detector’s “best” pointing (highest signal or most symmetric interferogram) and compare it to the band center of its second best pointing. The data from the best pointing should have bet-

ter coupling and alignment with the interferometer axis compared to the second best pointing, giving some quantification of the effect of being off-axis. This is done for BICEP3 in the top right of Fig. 3.7. The uncertainty in the real data (0.23 GHz, or 0.24 % of the band center) is very similar to the uncertainty predicted by the model (0.2 %).

3.3.2 MEASUREMENT SYSTEMATICS

As with any other calibration apparatus, the FTS measurement is not perfect. As we've already seen in the previous section, imperfect coupling between the detectors and the FTS aperture and interferometer can noticeably change the shape of the resulting spectra. We also must consider misalignments in the interferometer itself, polarization-dependent effects, and the inherent repeatability from measurement to measurement. Thankfully, during the data-taking process we use multiple configurations to allow us to probe and quantify these effects.

Polarization-dependent systematics can be studied by comparing spectra taken at different FTS rotation angles. One such systematic is the efficiency of the polarized wire grids. An ideal grid would completely reflect the field component parallel to the wires and completely transmit the perpendicular component. Measuring the band-pass at two angles probes the frequency-dependence and imperfections of this separation. A beamsplitter aligned at some angle other than 45° will also lead to different signal between the two rotation measurements. In the BICEP3 2017 FTS campaign,

we measured the spectra of both polarization detectors with the FTS at 45° and at 135° . When comparing the band centers between the two angles on a per-detector basis, we find an uncertainty of 0.18 GHz (0.2%) in this case, as seen in the top right of Fig. 3.7.

After being shipped to the South Pole, the individual components of the interferometer must be re-installed and aligned before any real measurements. Therefore, the alignment is never guaranteed to be exactly the same between any two data sets. Misalignments lead to asymmetry in the interferograms, which will lead to a nonzero imaginary component of the spectrum, and possibly half-scan direction dependent effects. At each FTS pointing, we nominally take four half-scans – two in each direction. The bottom-left histogram of Fig. 3.7 shows an uncertainty of 0.27 GHz (0.3%) in the case of BICEP3 2017 FTS data.

Similarly, if we compare two half-scans moving in the *same* direction, there should be no systematic differences between the two measurements. The alignment and coupling will be as identical as possible, allowing for a quantification of the overall repeatability of the FTS measurement. Looking at the band center differences for BICEP3 2017 FTS data between two parallel half-scans (bottom right of Fig. 3.7), we see an uncertainty of 0.24 GHz (0.2%). We consider this to be an estimate of the statistical uncertainty, which is at a very similar level to the systematic uncertainties discussed so far.

The distributions of the band center shifts for BICEP3 in Fig. 3.7 are roughly Gaus-

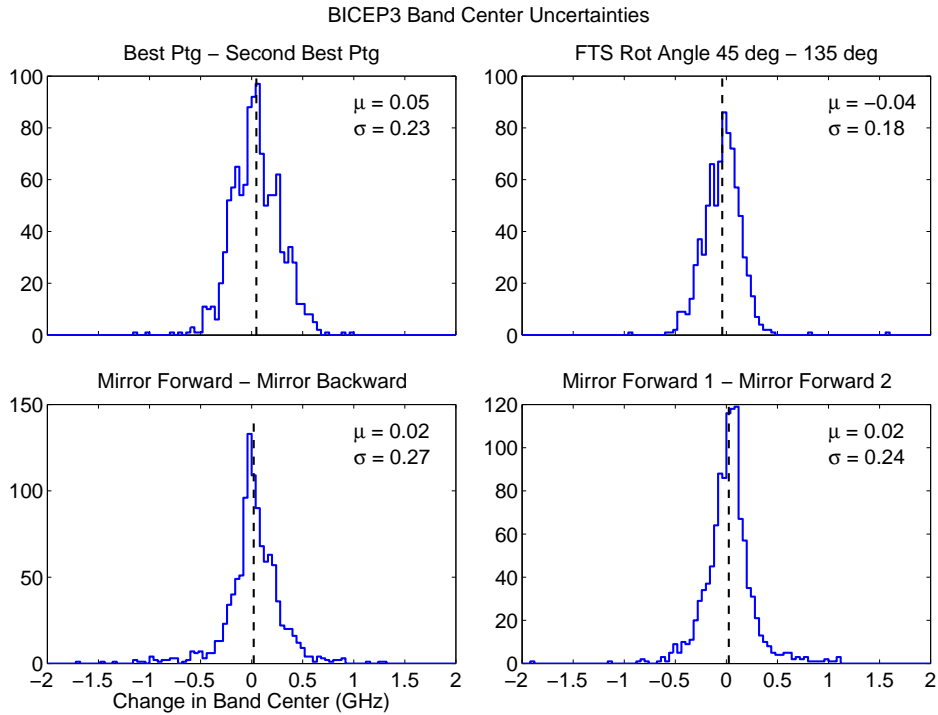


Figure 3.7: BICEP3 per-detector shifts in band center when considering various systematics associated with the FTS measurement. The mean (μ) and standard deviation (σ) in GHz for each distribution are shown in text in the top right of each histogram. *Top left:* comparing the “best” and second-best FTS pointing (usually determined by spectrum signal-to-noise), probing the effect of imperfect coupling to the interferometer axis. *Top right:* comparing two FTS rotation angles separated by 90° , probing wire grid imperfections or other polarization-dependent effects. *Bottom left:* comparing moving mirror half-scan direction, which is sensitive to the alignment of the interferometer optical elements. *Bottom right:* comparing two half-scans that move in the same direction in the same data set, which we consider to be an estimate of the overall repeatability. The data here was taken in 2017, after the LPE filter replacement.

sian and centered on zero, which gives us confidence that our band center estimates are not uniformly biased in either direction (this is also true of the other frequency bands, not shown). Therefore we use the width of these distributions as estimates of the band center uncertainty. However, it is important to note that the uncertainties associated with different FTS pointings, rotation angles, and mirror scan directions are likely not independent. For example, a misaligned beamsplitter could exacerbate poor coupling of the beam to the interferometer axis, while also creating polarization-specific distortions in the spectrum. This could be further studied using improved modeling by, for example, using ray-tracing simulations through the interferometer with ideal and non-ideal coupling and alignment.

Table 3.3 lists the same band center uncertainties as shown in Fig. 3.7 for all BICEP3 and *Keck* observing bands. In general, the systematic uncertainties are comparable to the statistical uncertainties, which are both subdominant to the detector-to-detector scatter. Now we are left to choose which number to use (for each band) as a single band center uncertainty in the maximum-likelihood search described in the next section. Due to possible covariances between the systematic shifts described in Table 3.3, we should not simply add the systematic uncertainties in quadrature. Since the systematic and statistical uncertainties are small compared to the detector-to-detector scatter, we nominally choose an upper limit of 2% band center error, which is above the scatter in each band shown in Table 3.3. As we will see, this leads to a relatively small bias on r .

Band	FPU Scatter	Δ Pointing	Δ Rotation Angle	Δ Mirror Scan Dir	Statistical Uncertainty
Keck 95	0.74 (0.8 %)	0.12 (0.1 %)	N/A*	0.15 (0.2 %)	0.16 (0.2 %)
Keck 150	1.9 (1.3 %)	0.25 (0.2 %)	N/A*	0.14 (<0.1 %)	0.15 (<0.1 %)
Keck 220	3.3 (1.5 %)	0.50 (0.2 %)	0.92 (0.4 %)	0.68 (0.3 %)	0.70 (0.3 %)
Keck 230	2.0 (0.9 %)	1.1 (0.5 %)	1.2 (0.5 %)	0.58 (0.2 %)	0.65 (0.3 %)
Keck 270	2.1 (0.8 %)	0.60 (0.2 %)	1.4 (0.5 %)	0.40 (0.1 %)	0.50 (0.2 %)
BICEP3	1.5 (1.6 %)	0.23 (0.2 %)	0.18 (0.2 %)	0.27 (0.3 %)	0.24 (0.2 %)

Table 3.3: FTS band center detector-to-detector scatter and uncertainties for each BICEP3 and *Keck Array* observing band, all listed in GHz and % of nominal band center. In general, Δ Pointing probes uncertainty in the detector-to-interferometer coupling, Δ Rotation Angle is sensitive to polarization-dependent systematics, and Δ Mirror Scan Direction probes interferometer alignment. However, it is unlikely each of these uncertainties are completely independent. These uncertainties are used in the multicomponent likelihood analysis to determine the bias on r due to band center error (see Section 3.3.3). *Due to time constraints, FTS measurements at multiple rotation angles were not done for these data sets.

3.3.3 IMPACT ON COSMOLOGICAL PARAMETER RECOVERY

We can use the per-band uncertainties on the band center in the multicomponent likelihood analysis to determine the resulting bias on r , with the goal of keeping that bias far below the experiment statistical uncertainty.

One method to quantify the effect of misestimated bandpasses in the multicomponent likelihood is to consider a uniform shift in the measured spectra in frequency by some static amount. This “band center shift” method is our approach in both the BK15 and BK18 data sets, but in principle the spectral shape could be distorted by, for example, applying some frequency-dependent change $S(\nu) \rightarrow S(\nu)(\nu/\nu_0)^\alpha$, where ν_0 and α are some physically-motivated parameters informed by FTS measurements. This approach could be taken in future work as the constraints on r improve, or as the FTS measurement uncertainties improve in future generations of telescopes.

In general, the bandpasses are used to scale the expected signals from external or

simulated foreground maps from their nominal frequency to the desired band, assuming some spectral scaling relationship for each component (such as a power law or a greybody). A shift in the band center of any of these bandpasses therefore changes the expected signal from each component. In this work we focus on the polarized dust component, which is relatively bright and not spectrally flat at 220 GHz, where the BK18 data set has many receiver-years of observations. There are three different approaches to consider: 1) rescale the dust amplitude using a new, shifted band center in each observing band, and ignore it in the likelihood analysis; 2) rescale the dust amplitude, and attempt to detect it in the likelihood using a series of nuisance parameters; 3) do not rescale the dust amplitude, and attempt to detect band center errors in the real data as-is using nuisance parameters.

For the BK15 data set, we considered all three scenarios. Four nuisance parameters were used: a fractional band center shift in each of the three observing bands (95, 150, and 220 GHz) and a fourth band center shift that is completely correlated across all three bands. The model we chose to include in the BK15 publication ([The BICEP2 and Keck Array Collaborations 2018](#)) is the third scenario listed above – do not manually rescale the dust amplitudes, but include the nuisance parameters in the likelihood analysis and observe the change in the recovered tensor-to-scalar ratio r , dust amplitude A_d , and synchrotron amplitude A_s . The 1D posteriors resulting from a sampling using Markov chain Monte Carlo for this scenario is shown in Fig. 3.8 (for more details on the multipcomponent likelihood analysis, see Section 6.2). The change




Figure 3.8: One-dimensional posteriors of the tensor-to-scalar ratio r , dust amplitude A_d , and synchrotron amplitude A_s for various alternative analyses from BK15. The baseline results are in black, and the band center error analysis are in blue dashed lines. In the band center analysis, four nuisance parameters were included in the likelihood, including one uncorrelated fractional band center error for each observing band (95, 150, and 220 GHz), and one additional fractional error completely correlated in all three bands. The priors used on each of these nuisance parameters were Gaussian with a mean of zero and a standard deviation of 0.02, as informed by FTS measurements at the time. This is a reproduction of Fig. 16 from BK15 ([The BICEP2 and Keck Array Collaborations 2018](#)).

seen in the final 1D posteriors after including band center errors is small, leading to a $\Delta\sigma_r < 1 \times 10^{-4}$.

For the BK18 analysis, again we are considering all three scenarios previously listed, with the expectation that the shift in r in any case will remain small, given the BK15 results. Here we present the results of maximum-likelihood (ML) searches after injecting 2% band center shifts in each BK18 band (BICEP3 95 GHz and *Keck* 95, 150, and 220 GHz) and including no nuisance parameters. Note that the ML search returns a single number corresponding to the best-estimate of the maximum likelihood value, not a full posterior distribution. The combination of band center shifts that lead to the worst-case bias in r (in the positive and negative direction) are shown in Fig. 3.9. When the band centers of BICEP3 and *Keck* 95 GHz are fractionally shifted down 2%

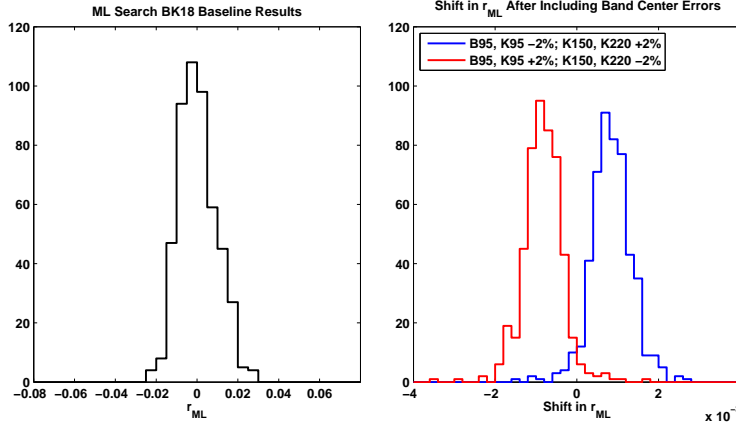


Figure 3.9: BK18 ML search results after including band center errors. *Left:* Histogram of the recovered value of r_{ML} from a search over 499 realizations of the BK18 baseline scenario (for full details see Section 6.3). *Right:* Histogram of shift in r_{ML} after including band center shifts in each frequency band. The fractional band center shifts that give the worst negative bias are +2 % in BICEP3 and *Keck* 95 GHz and -2 % in *Keck* 150 and 220 GHz (red), resulting in $\Delta r_{ML} = -8.3 \pm 4.9 \times 10^{-4}$. The shifts that give the worst positive bias are -2 % in BICEP3 and *Keck* 95 GHz and +2 % in *Keck* 150 and 220 GHz (blue), resulting in $\Delta r_{ML} = 8.4 \pm 5.0 \times 10^{-4}$. All other band center error combinations result in shifts that lie between these bracketing cases, which are subdominant to the BK18 statistical uncertainty of $\sigma(r) = 0.009$ (The BICEP2 and *Keck* Array Collaborations 2021).

and the band centers of *Keck* 150 and 220 GHz are shifted up 2 %, the shift in the recovered ML value of r is $-8.3 \pm 4.9 \times 10^{-4}$. When the opposite shift is applied to each band the change in the ML value of r becomes $8.4 \pm 5.0 \times 10^{-4}$. All other band center error combinations result in shifts in r that are between these bracketing cases. All the observed shifts are subdominant to the BK18 statistical uncertainty of $\sigma(r) = 0.009$ (The BICEP2 and *Keck* Array Collaborations 2021).

*Don't let anyone rob you of your imagination,
your creativity, or your curiosity. It's your place
in the world; it's your life. Go on and do all you
can with it, and make it the life you want to
live.*

Mae Jemison

4

Beam Measurements and Analysis

Pipeline

A PRECISION MEASUREMENT of B -mode power at degree angular scales requires exquisite control of systematics. Many of the most prominent systematics in modern CMB tele-

scopes – temperature-to-polarization ($T \rightarrow P$) leakage and extended sidelobe response, for example – rely on detailed understanding of the detector angular response in order to constrain them. Many of the design choices made when building these telescopes, including the lens material, the optical coupling to the detectors, and the external and internal baffling structures, are all heavily influenced by their impact on controlling this angular response.

In the millimeter regime, it is common to consider a detector an emitter of photons, rather than a collector. This is permitted because of the reciprocity theorem, which states that Maxwell’s equations are time-invariant. This is why the detector angular response is called the “beam.” The signal output from the detector at any given time is the spatial convolution of the beam with the temperature pattern of whatever it is observing – the atmosphere, the ground, astrophysical foregrounds, or the CMB. Therefore, not only are beam measurements required for systematic control, they are also needed to convert detector timestreams into CMB maps, and eventually into power spectra and cosmological parameter constraints.

Throughout the history of the program, *BICEP/Keck Array* has prioritized beam measurements of all detectors that are used for on-sky observations. In this chapter, we discuss the procedure and low-level results of these measurements, which includes near-field beam mapping (NFBM), far-sidelobe mapping (FSL), and far-field beam mapping (FFBM). Emphasis is placed on the far-field measurements and analysis, due to their wide variety of applications in CMB analysis and systematics control dis-

cussed in later chapters. After discussing the FFBM measurement procedure, we walk through the FFBM analysis pipeline, paying careful attention to possible systematic contamination from the analysis itself, followed by a brief discussion on each of the other beam measurements listed above.

4.1 FAR-FIELD BEAM MAPPING PROCEDURE

The electromagnetic signal as emitted from a receiver (in a time-reverse sense) has two distinct “zones” – the near field and the far field. In the near field, the evolution of the waves as they propagate outward is complex, and measurements in this area are very sensitive to the phase of the radiation in the plane of measurement. Once the waves travel far enough from the receiver, they can be considered standard plane waves, where a change in the electric field generates a change in the magnetic field and vice-versa, infinitely. The soft criterion that separates these two zones is $2D^2/\lambda$, where D is the diameter of the antenna element (in this case the receiver aperture) and λ is the radiation wavelength. Once beyond that distance, the electromagnetic wave will virtually be the same as it is at infinity.

Some experiments take maps of bright sources fixed to the celestial sky, such as the planets or nearby stars. Our beams are too large to map our detectors using those objects. However, the small-aperture design of our telescopes begets a very manageable far field distance of 100-200 m. BICEP3 and *Keck* (now BICEP Array) exist in separate

Band	$2D^2 / \lambda$
Keck 95	44 m
Keck 150	71 m
Keck 220/230	104-108 m
Keck 270	129 m
BICEP3	171 m
BICEP Array 30/40	61-79 m
BICEP Array 95*	192 m
BICEP Array 150*	303 m
BICEP Array 220/270*	444-545 m

Table 4.1: List of far-field distances for all *Keck*, BICEP3, and BICEP Array observing bands, assuming an aperture diameter of 26.4, 52, and 55 cm, respectively. Simulations discussed in Section 7.3 show the effect of mapping a source from a distance below this criterion, which will become significant for future BICEP Array receivers. *Not currently deployed; assumes a nominal band center.

facilities that are spaced roughly 200 m apart, allowing us to map a source mounted on the opposite building from the observing receiver. Table 4.1 shows a list of the far-field distances for all *Keck*, BICEP3, and BICEP Array observing bands. Future BICEP Array receivers at 150 and 220/270 GHz will have a distance significantly larger than 200 m, however. The impact of being under the far-field criterion on the final beam shape is explored in simulations presented in Section 7.3.

4.1.1 SOURCES

The source most commonly used for far-field beam mapping is a large chopped thermal source that is affixed to a mast and raised ~ 10 m above the roof of the observatory. The choppers currently used were made in 2015, with the chopper blade and enclosure made of a composite material composed of carbon fiber, low-density foam,




Figure 4.1: *Left:* Image of *Keck Array* and BICEP3 in Feb 2017 performing simultaneous beam mapping campaigns, with redirecting mirrors and thermal sources visible. *Middle:* The thermal chopper undergoing lab tests with the enclosure off. *Right:* The same chopper with enclosure on, sealed with Zotefoam HD30. Figure from BK-XI (*The BICEP2 and Keck Array Collaborations 2019*).

and Nomex honeycomb fiber. The blade is coated in microwave absorber (HR-10) and the enclosure is covered with aluminum tape to redirect non-coupling rays to the cold sky. The dual-blade chopper spins as fast as 9 Hz (chop rate of 18 Hz). The source chops between ambient-temperature HR-10 (~ 250 K) and a 45° aluminum mirror that redirects to zenith (~ 10 K). Zotefoam HD30 seals the aperture during operation in order to allow millimeter waves to transmit to the source while preventing excessive vibrations from the wind. The chopper with the enclosure on and off can be seen in the right two panels of Fig. 4.1.

Another source that has been used for FFBM observations is the high-powered broad-spectrum noise source (BSNS) which is also used for measurements of the polarization angle and cross-polar response and of the far-sidelobe response (Section 4.4). The BSNS is composed of a 50Ω SMA termination whose thermal radiation is amplified, upconverted, and bandpass filtered to the desired frequency range. The signal is coupled to free space via a WR-6 horn waveguide. For a full description of the source,

see [Bradford 2012](#) and [Cornelison et al. 2020](#). A series of attenuators provides the flexibility to map out the main beam, mid-sidelobes, and far-sidelobes independently with high signal-to-noise.

The chop signal is recorded with a motor encoder for the chopper; the BSNS chop is controlled with a PIN switch via an externally provided TTL signal. In either case, the chop reference is sent to the opposite building via an optical fiber, and plugged into the housekeeping system for that receiver. The chop reference is recorded simultaneously with the detector timestreams and used for demodulating the signal.

4.1.2 REDIRECTING FLAT MIRROR

Due to the presence of a large ground shield around each telescope, and because we cannot make long observations below an elevation of 45° , a large aluminum honeycomb mirror is used to redirect light coming from the telescope at zenith to the source at the opposite building. A separate mirror is used each for *Keck* (1.8 x 2.7 m) and BICEP3 (1.7 x 2.5 m). The first beam mapping campaign for BICEP Array reused the same mirror for *Keck*, but an improved mirror may be developed in the coming 1-2 years. For *Keck* and BICEP3, the forebaffles are removed from the receivers while beam mapping due to space constraints with the mirrors.

For BICEP3, the mirror has good enough coverage to ensure that each detector beam out to a radius of $> 2^\circ$ couples to the source¹. This is true for all boresight

¹Coverage out to a 2° radius is the minimum, and only true for far-edge detectors. For

rotation angles. For *Keck*, it is unfeasible to have a mirror large enough to simultaneously redirect the light for all five receivers. The *Keck* mirror therefore only covers 2-3 receivers at one time, and requires multiple positions in order to fully map the detectors at all angles. The nominal series of *Keck* beam mapping schedules covers a range of ten uniformly-spaced boresight rotation angles each over two mirror positions. Care must be taken in analysis to ensure that only rays coupling from the detector to the source via the mirror are kept in the component beam maps – this is discussed in greater detail in Section 4.2.4.

In January 2018, photogrammetry measurements were taken on both mirrors *in-situ*, to ensure the curvature and surface roughness was at acceptable levels. This was done by applying numerous calibration stickers over the surface the mirror, taking photographs at many angles, and using the photogrammetry software Photomodeler² to produce surface curvature data. The results showed a maximum deviation of ± 0.8 mm for the *Keck* mirror and ± 0.3 mm for the BICEP3 mirror, both of which featured an overall parabolic curvature and no detectable small-scale surface roughness. The results for *Keck* are shown in Fig. 4.2. The impact of this small curvature on the overall beam shape derived from far-field beam maps is discussed in Section 7.3.

most detectors in the FPU, the coverage exceeds 10° from the main beam.

²Photomodeler: www.photomodeler.com




Figure 4.2: Photogrammetry of the *Keck* far-field beam map redirecting mirror *in-situ*. *Left:* photo of the mirror with calibration stickers, which are marked and identified by the photogrammetry software Photomodeler. *Right:* resulting analysis of the surface flatness of the mirror. Parabolic curvature appears to dominate over small-scale surface roughness, but the deviation from flat is no larger than ± 0.8 mm. Results for the smaller BICEP3 mirror (not shown) indicate a maximum deviation of ± 0.3 mm.

4.1.3 SCAN STRATEGY

There are a number of experimental choices to make when designing the far-field beam mapping observation schedules. These decisions must be informed by the physical limitations of the measurement (fridge hold time, telescope field of view) and the desired binning strategy used in beam map analysis. Each of these decisions and their associated constraints are listed below:

- **Azimuth scan speed:** Combined with the GCP fast mode sampling rate (~ 150 Hz for BICEP3 and ~ 180 Hz for *Keck*), we choose to have multiple chop cycles across the main beam in each azimuth scan. However, we also want to scan fast enough to prevent the overall schedule length from becoming too long. We use $0.4^\circ/\text{sec}$ for *Keck* $0.8^\circ/\text{sec}$ for BICEP3, and $0.6^\circ/\text{sec}$ for BICEP Array.
- **Azimuth range:** This is limited by the telescope field of view. In order to guarantee that each detector sees the source out to a radius of at least 2° , we use an azimuth throw of 22.8° (*Keck*), 30° (BICEP3), and 32° (BICEP Array).
- **Elevation range and step size:** Again, the range is limited by the field of

Experiment	Az Scan Speed	Az Range	El Step Size	El Range
<i>Keck Array</i>	0.4°/sec	22.8°	0.05°	20°
BICEP3	0.4°/sec	30°	0.05°	30°
BICEP Array	0.6°/sec	32°	0.05°	24°

Table 4.2: Far-field beam mapping schedule parameters for each experiment.

view, but the step size is dictated by the desired pixelization of the beam maps. To guarantee at least two hits in each 0.1° square pixel, we use an elevation step size of 0.05° for all experiments. The elevation range is 20° for *Keck*, 30° for BICEP3, and 24° for BICEP Array.

Because we conservatively mask the beam maps where the main beam intersects the ground, we use multiple boresight rotation angles to ensure complete coverage of each detector beam. A summary of these choices is shown in Table 4.1.

4.2 FAR-FIELD BEAM MAP ANALYSIS PIPELINE

Differences in the beam response between two co-located, orthogonally polarized detectors leads to temperature-to-polarization leakage that can bias final cosmological parameter constraints. Although deprojection and cancellation from coadding over multiple boresight rotation angles drastically reduces this leakage, the remaining difference power (the undeprojected residuals) still bias the constraint on r (this is discussed in much greater detail in Chapter 6). The undeprojected residuals in general have a complex shape, and often have an amplitude $< 0.5\%$ of the main beam peak power. For these reasons, it is critical that the pipeline that takes in raw beam map timestreams and generates final composite beam maps injects minimal systematics. In

the case of unavoidable systematic effects (e.g. the pixel window function associated with beam map pixel size), we must quantify, and if possible remove, this contamination.

In this section, we walk through the FFBM pipeline analysis steps, starting with raw timestreams and ending with composite beam maps and array-averaged beam profiles. Many aspects of this pipeline were improved in the past few years, in an effort to minimize systematics associated with certain analysis choices. We highlight the changes associated with this upgrade, pay careful attention to possible systematic effects, and note room for future improvement. Many of these analysis steps are also applied in other beam map measurements discussed later in this chapter.

4.2.1 DEMODULATION

Modulation of the source for any kind of beam measurement is a powerful tool to separate the signal coming from the source from the response to any other object the beam may intercept. These stray signals may come from the ambient-temperature ground, the atmosphere (whose temperature varies with airmass), or other buildings and structures in the local vicinity. As discussed in Section 4.1.1, each source used for observation is modulated in some way, with a chop reference that is recorded simultaneously with the real data. Ideally, by isolating the signal of the chopped source to a single frequency with a known waveform, we could confidently construct our beam maps using signal that is entirely from the chopped source (see Fig. 4.3 for a real ex-

ample). However, there are multiple avenues for systematics to enter the analysis at this stage, and it is critical to contain them here – imperfect demodulation can impact both the beam shape and noise characteristics in ways that cannot be corrected further in the pipeline.

The chop reference in all of our sources is a binary square wave, and one demodulated sample is given for each full period of the square wave. The “cosine” (in-phase) component is evaluated by taking the (uniform-weighted) mean of all samples when the chop reference is “high” and subtracting the mean of the samples in the second half of the preceding “low” phase and the first half of the subsequent low phase. The “sine” (quadrature) component uses the same kernel, 90° out of phase from the cosine component. Assuming perfect modulation at an infinitely fast chop rate, the cosine component would perfectly match the source signal and the sine component would be zero. Given that the fast mode data sampling rate used in beam mapping is $\sim 150\text{-}180$ Hz, if we want to ensure that we have multiple samples in each high and low phase of the chop reference, we must limit our chop rate to be 20 Hz or slower.

In some cases, the chop reference may not perfectly reflect how the source is truly being chopped. The most common form of this is single-value spikes or glitches in the binary chop reference that do not exist in the detector timestreams. This can cause changes in the beam shape, or create “hot pixels” in the binned beam maps. Either of these can interfere with the 2D elliptical Gaussian fits, cause an otherwise good beam to be cut, or lead to inaccurate estimates of $T \rightarrow P$ leakage in the specialized

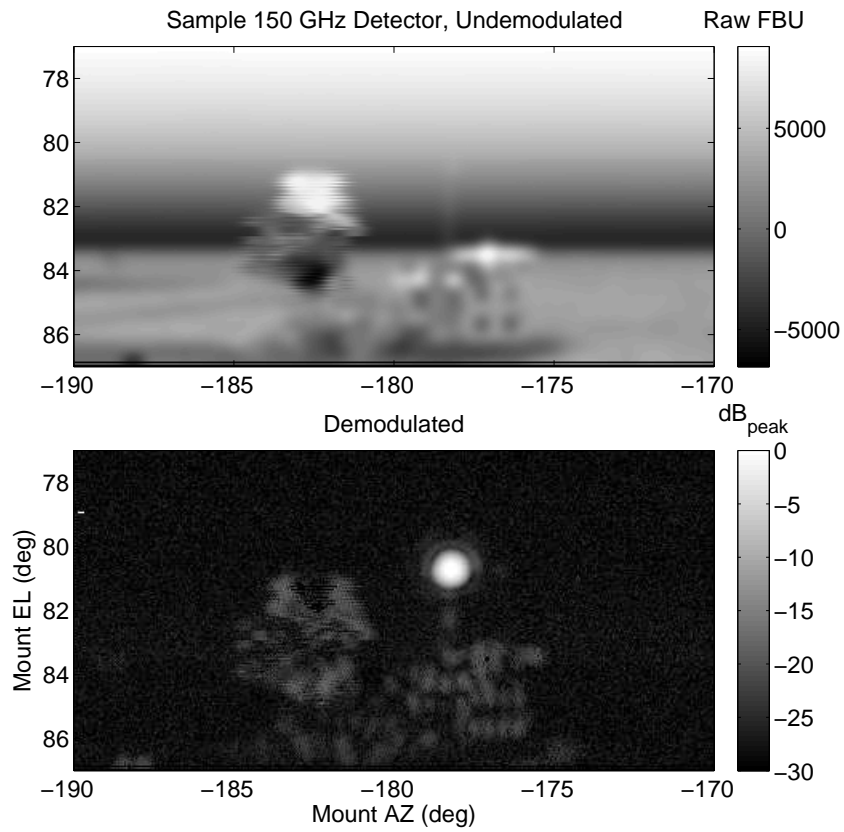


Figure 4.3: A full component beam map from a sample *Keck* 150 GHz detector, before (top) and after (bottom) demodulation. Contamination from the Dark Sector Lab building, the South Pole Telescope 10 m dish, and the mast holding the source are clearly visible before demodulation. After demodulating at 16 Hz, the contamination is drastically reduced but still present at low levels (note the top plot has a linear color scale while the bottom plot has a logarithmic color scale). This is why we choose to mask regions of known contamination.

$T \rightarrow P$ leakage simulations. One key improvement made to the pipeline involves “cleaning” these imperfections in the chop reference, by constructing an ideal square wave using the peak frequency of the Fourier transform of the real chop reference, and rotating the constructed square wave by a low-pass filtered version of the phase difference between the constructed and real chop reference. This effectively removes high-frequency anomalies, and also gives sub-sample precision of the cleaned chop reference. This sub-sample precision allows us to properly identify which side of a chop transition a sample belongs to when the real chop transition happens between samples. An example of a demodulated timestream with a raw and cleaned chop reference is shown in Fig. 4.4.

By choosing a uniform-weighted time-domain kernel, we have implicitly assumed that the physical chop mechanism is instantaneous. For the PIN switch-operated BSNS this may be very close to true, but for the large thermal chopper we expect a phase gradient along the azimuth direction due to the finite swing time of the chopper blades. This effect is small – the phase shift induced with respect to the chop reference varies by ~ 1.5 radians across the main beam in the case of BICEP3³. This affects both polarized detectors in a detector pair equally, leaving no effect on the pair-difference beams. However, it does make the measured beam shape slightly ($< 1\%$) more elliptical in the azimuth direction in each component beam map. Future im-

³This is done by looking at “phase” beam maps, made by binning values of $\arctan(\text{sine component} / \text{cosine component})$ instead of binning the cosine component.

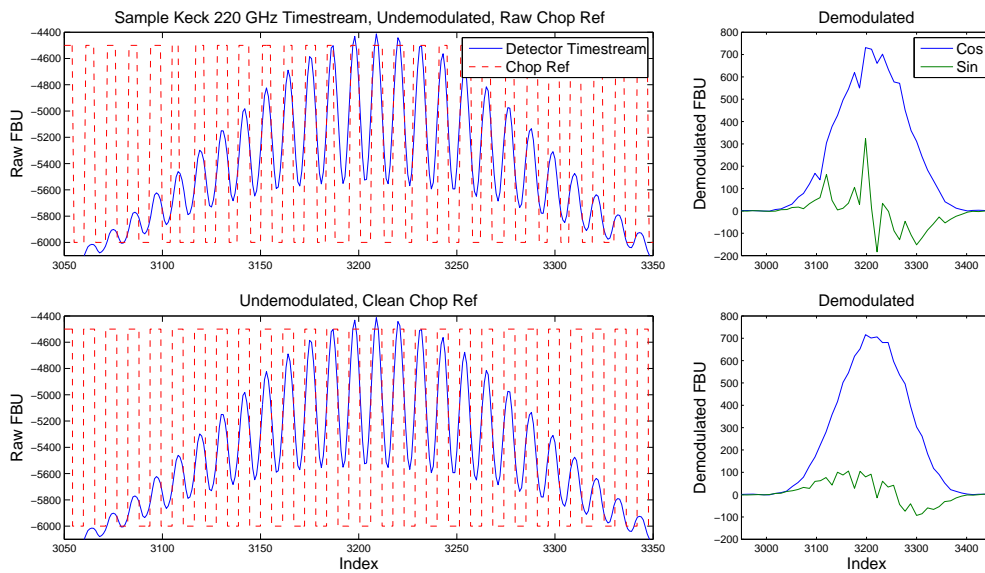


Figure 4.4: A beam map timestream of a sample *Keck* 150 GHz detector, showing the improvement in demodulated signal after cleaning the chop reference. The timestream shown is a portion of one azimuth scan where the main beam illuminates the thermal chopper. *Top left:* Raw detector timestream and raw binary chop reference, scaled for presentation. *Top right:* Cosine (in-phase) and sine (quadrature) components of the demodulated signal using the raw chop reference. *Bottom left:* Same raw detector timestream, with cleaned binary chop reference with same scaling. *Bottom right:* Demodulated signal using clean chop reference. Cleaning the chop reference made noticeable improvements – see near samples 3110 and 3190, for instance. The resulting cosine component is much smoother, and the sine component is closer to zero.




Figure 4.5: Diagram of the instrument-fixed beam map coordinate system as described in the text. Each pixel P is at a location (r, θ) from the receiver boresight B . The (x', y') coordinate is defined from the location of P and reference to the $\theta = 0$ ray. Figure from BK XI ([The BICEP2 and Keck Array Collaborations 2019](#)).

provements can be made to the demodulation routine to correct for this effect, by applying a position-dependent phase shift to each sample of the chop reference before demodulation.

4.2.2 BEAM MAP COORDINATE SYSTEM

Prior to binning into maps, the coordinates are calculated for the post-demodulation, downsampled timestreams. The coordinate system is independent of the instrument orientation with respect to the sky. A pixel P containing two orthogonally polarized detectors has a defined location (r, θ) with respect to the boresight B (the telecentric

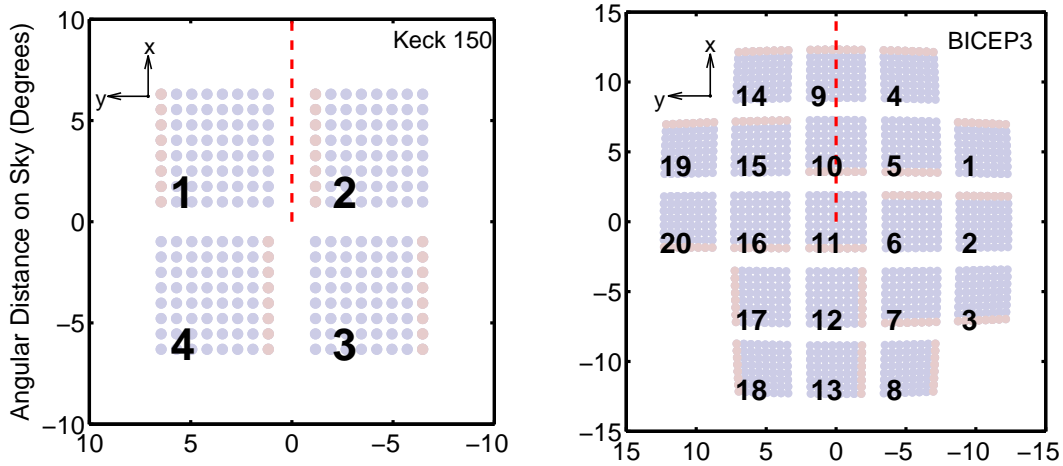


Figure 4.6: Projections of the *Keck* (left) and BICEP3 (right) FPUs on the sky. The $\theta = 0^\circ$ ray is shown as a red dashed line. The tile numbers are shown on each tile. Each detector pair is shown as a single point, with the first column of each tile highlighted in red (this helps show the relative orientation of each tile as installed in the FPU). The *Keck* FPU shown is a 150 GHz array, which is identical in detector number and spacing to the 220, 230, and 270 GHz arrays (the *Keck* 95 GHz arrays feature 6x6 detector pairs per tile instead of 8x8).

axis), where r is the radial distance from B and θ is the counterclockwise angle looking out from the telescope towards the sky from the $\theta = 0^\circ$ ray. This is shown in a diagram in Fig. 4.5. The $\theta = 0^\circ$ ray is fixed to the instrument: for *Keck* we choose this to be the line that divides Tiles 1 and 2, and for BICEP3 we choose it to be the line that runs along Tiles 11, 10, and 9 (see Fig. 4.6).

First, the raw mount encoder values are converted to horizon azimuth (AZ), elevation (EL), and boresight angle (DK) timestreams by using the telescope pointing model that is updated roughly every month with star pointing observations. The mount (AZ, EL, DK) coordinate timestreams are converted to per-detector (x' , y')

coordinate timestreams, which requires precise knowledge of the detector pointings as well as the location and orientation of the source and redirecting mirror. However, small inaccuracies in those parameters at this stage result in a small translation from the true beam centers to the measured beam centers, which is fit and removed after the fitting stage, so *exact* measurements of the mirror and source parameters are not strictly required⁴.

Prior to the pipeline upgrade, these coordinates were calculated after the binning step, by rotating and interpolating the binned maps from (AZ, EL, DK) to (x' , y') coordinates. This interpolation effectively smoothed the beam, which had minimal effect on the larger beam sizes (*Keck* 95 and 150 GHz) in the most important science ℓ bins, but had noticeable impact for smaller beams, such as *Keck* 220 GHz and BICEP3 95 GHz. By calculating all coordinate timestream before binning, all interpolation steps are eliminated from the pipeline.

4.2.3 2D ELLIPTICAL GAUSSIAN FITS

Once the detector timestreams have been demodulated and the coordinate timestreams have been derived, we choose to do the 2D elliptical Gaussian fits at this stage. Fitting before the binning step ensures that the resulting Gaussian parameters are not biased by the pixel window function. The basis used for the Gaussian fits has been

⁴However, precise knowledge of the source and mirror parameters can help with other auxiliary beam map analyses, such as deriving absolute detector pointings or mapping deformities in the mirror.

described in many previous publications – here we closely follow [St Germaine et al. 2020](#).

The standard 2D elliptical Gaussian function has six free parameters:

$$B(x) = \frac{1}{\Omega} e^{-\frac{1}{2}(x-\boldsymbol{\mu})^T \Sigma^{-1}(x-\boldsymbol{\mu})}, \quad (4.1)$$

where $x = (x', y')$ is the beam map coordinate, $\boldsymbol{\mu} = (x_0, y_0)$ is the beam center, Ω is the normalization, and Σ is the covariance matrix, parametrized as:

$$\Sigma = \begin{pmatrix} \sigma^2(1+p) & c\sigma^2 \\ c\sigma^2 & \sigma^2(1-p) \end{pmatrix}. \quad (4.2)$$

σ is the beamwidth, and p and c are plus and cross ellipticity, respectively. A Gaussian with a major axis oriented along the x_0 or y_0 axes has $+p$ or $-p$ ellipticity, respectively, and one with a major axis oriented along $y' = \pm x'$ has $\pm c$ ellipticity. Differential parameter estimates are defined to be that parameter for the A -polarization detector minus that for the B -polarization detector. For example, differential ellipticity plus is defined as $dp = p_A - p_B$. In general, we do not use the gain Ω or absolute pointings as derived from these fit parameters. The absolute gain and pointing of our detectors for CMB analysis (when biased on the Titanium transition) are found by taking cross-correlations with *Planck*.

Imperfect fitting can happen in a couple ways. “Spikes” in the detector timestreams

(singular, highly negative or positive values) can bias a fit drastically. The occurrence of these spikes has been significantly reduced by the chop reference cleaning method described in Section 4.2.1. Another is from poor initial guesses in the fitting algorithm. Regardless of the performance of the chosen algorithm, a starting point that is orders of magnitude away from the true value will result in a failed or poor fit. Simply using prior knowledge of the expected beam size, amplitudes, and ellipticity can avoid this problem.

The current fitting function uses the Nelder-Mead method, but the `minuit` function minimization tool⁵ has also obtained similar results. This analysis returns per-detector, per-schedule Gaussian parameters describing the beam shapes. From these parameters, we make automatic data-quality cuts, and can analyze spatial patterns and trends across the receiver FPU to verify the optical performance. The results of these fits for BICEP/*Keck Array* receivers are presented in Section 5.2. At this stage, the (x', y') coordinate timestreams are shifted such that the pair centroid (the point exactly between the fit centers of the *A* beam and the *B* beam) is centered at the origin.

4.2.4 BINNING AND MASKING

The next step is to take the demodulated detector timestreams and bin them into maps, allowing for easy coadding and averaging. The resulting per-detector, per-

⁵`minuit` was originally developed by CERN, root.cern.ch/download/minuit.pdf

schedule beam maps are referred to as “component” maps throughout this dissertation.

For all detectors across all bands, the beams are binned onto the same common grid, with 0.1° square pixels, with a pixel centered at $(x', y') = (0, 0)$. Combined with our FFBS observation strategy detailed in Section 4.1.3, this enforces at least two hits per map pixel. The value in a given map pixel is simply the uniform-weighted mean of all samples within that pixel. The observation band with the smallest beam in BICEP/*Keck Array* is currently the *Keck* 270 GHz band, with a Gaussian $\sigma = 0.120^\circ$. The pixel window function does not significantly impact the beam profile in the relevant multipole ranges for CMB science, but as we move to future receivers such as BICEP Array 270 GHz (which will have $\sigma \simeq 0.060^\circ$ beam width), we will likely need to decrease this map pixel size, and use finer elevation steps in the beam map observation schedules⁶.

At this point, a mask is applied to each component map to remove the ground, the mast holding the source, and the South Pole Telescope 10 m dish (for *Keck* beam maps). For *Keck*, a second mask is applied, to remove any rays from the detector that do not couple to the mirror and source. This mask is calculated geometrically given the angle, size, and position of the mirror, and varies with the receiver position in the

⁶The pixel window function is theoretically calculable and can be partially removed in harmonic space, but 1) as the beam width becomes small enough compared to the map pixel width, the loss of information at relevant multipole ranges becomes irrecoverable, and 2) since a correction cannot be applied in beam map space, this has negative consequences on the efficacy of, say, the $T \rightarrow P$ leakage simulations where the beam map and temperature map convolution is done in real space.

drum of the mount. A small margin is applied to make the mask stricter, in order to account for distortions that occur when a detector main beam is partially, but not completely, off the edge of the mirror. A diagram of the mirror mask and a demonstration of the mirror and ground masks are shown in Fig. 4.7

There is some room for improvement in the masking. In particular, the sun is unaccounted for in either of these masks (beam mapping always occurs in the austral summer when the sun is above the horizon). Over the course of an 8-hour schedule, the sun can create a swath of unusable data in a component beam map. In some cases, the sun happens to be close to the source and therefore appears near the main beam in the component map, and its presence triggers an automatic cut (see the following Section). In other cases, it intersects the main beam far enough away from the source that it eludes the automatic cuts, but still affects the mid-sidelobe area of the map. In principle, using the known location of the sun versus time along with knowledge of the mirror and mount orientation, the sun could be effectively removed from the component maps while still keeping the rest of the map usable.

4.2.5 CUTS AND COMPOSITING

With a collection of component beams and corresponding 2D elliptical Gaussian fits for each detector, we can now stack them together to create composite beam maps, which are the highest signal-to-noise representation of a detector's true beam response. As highlighted in numerous different topics throughout this dissertation, the per-

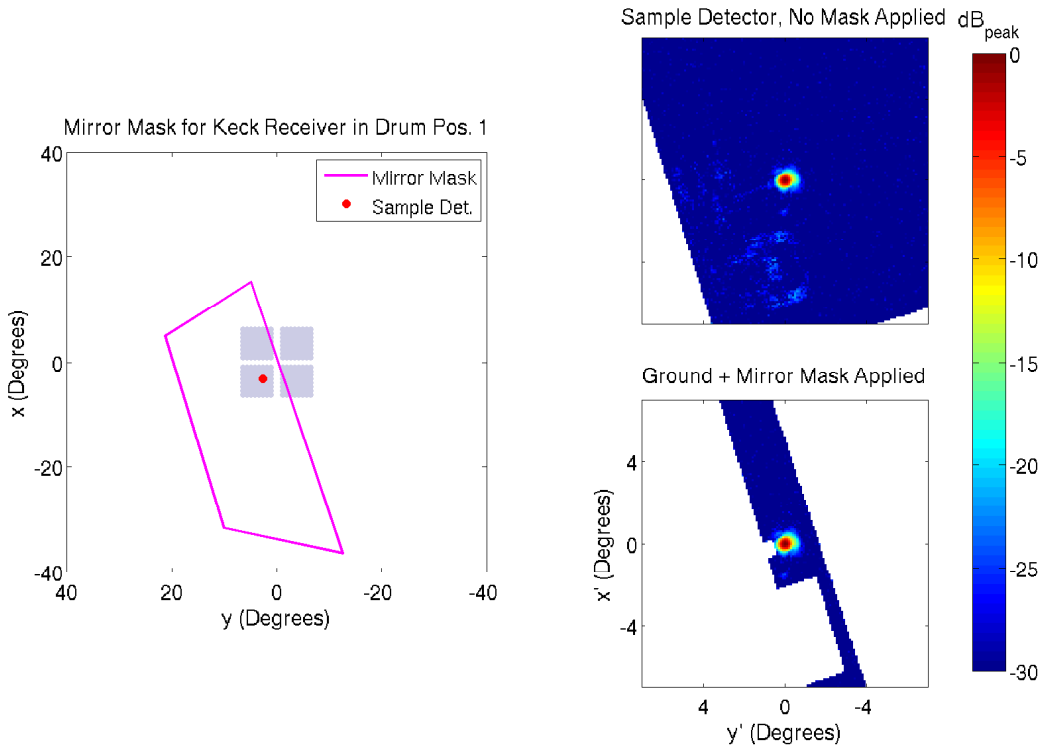


Figure 4.7: *Left:* Projection of a *Keck* 220 GHz FPU onto the sky with the mask from the redirecting mirror overlaid. These masks vary with mirror position and the receiver position in the mount – here the receiver is in drum position 1 (the location of Rx0 when the mount boresight rotation angle $DK = 166^\circ$) and the mirror is in the “back” position. The red dot highlights the sample detector whose beam maps are shown in the right-hand plots. *Right:* Sample component beam map, with no mask (top) and both the mirror and ground masks (bottom) applied. The ground mask cuts out the visible structure in the $y' \lesssim 0^\circ$ region (including the SPT 10 m dish), and the mirror mask cuts out rays in the $y' \gtrsim 4^\circ$ region which do not couple to the mirror and source. Note in this case the receiver is not directly under the mirror – this example was chosen to illustrate how the mirror mask can identify usable regions of the beam for detectors only partially illuminating the mirror.

detector composite beam maps have great utility for estimating $T \rightarrow P$ leakage, assessing overall optical performance, validating optical models, and more. However, we must be judicious when choosing which component maps to include in a composite stack.

The task of separating “good” component beam maps from “bad” ones grows more challenging with the increase in detector count in the BICEP/*Keck Array* program over the past few years. BICEP3 and *Keck Array* each have ~ 2500 detectors, and each year we take ~ 40 - 70 beam schedules that map each of those detectors. At this point we can no longer manually check each component beam map by eye – we must rely on a series of automated cuts that can reliably identify detectors whose beams show unacceptable qualities. These automated cuts include:

- The main beam must couple to the redirecting mirror and source. In Section 4.2.4, we demonstrated the mirror mask calculation used to remove rays that do not couple to the mirror and source. If the center of the main beam is removed by this mask, the component map is cut.
- The detector must be responsive, and have reasonable fit parameters. Cuts on the relative pointing, beam width, amplitude, and ellipticity are frequency band-dependent.
- Both detectors in the polarized pair must be responsive and have reasonable fit parameters. We generally do not keep the component map of an A -polarization detector if the B -polarization component map is poor.
- The detector must remain responsive with reasonable noise levels away from the main beam. In some cases, a component beam map will have good response in the inner $< 2^\circ$, but will fall off transition and become unresponsive as the elevation in the beam mapping raster changes. This can lead to unphysical leakage estimates derived from the pair difference if not removed.

- The fit residual must be reasonable. In rare cases, the 2D elliptical Gaussian fit may appear reasonable but is actually a poor fit to the data. This could happen, for example, if the fitting algorithm does not sufficiently explore the 6D parameter space before converging to its final value. We catch these by cutting component beam maps that have an unusually high (beam - fit) residual.

Once cuts have been applied to the component maps, we can coadd them together to create composite maps. This is done by stacking all good composite maps on the common (x', y') grid, then taking the median in each map pixel. Because the masked component maps were taken at a variety of boresight rotation angles, the resulting composite map usually has reliable signal out to $> 10^\circ$ from the main beam. For most detectors, the noise floor near the main beam ($r < 2^\circ$) can often reach as low as -40 dB from the main beam peak. The composite map and all component maps for a sample BICEP3 detector from the 2017 FFBM observations are shown in Fig. 4.8.

At this stage we also construct “split” beam maps, which are made by randomly separating all component maps for a given detector into two halves with an equal number of components each, then taking the difference. These are a good estimate of the noise in a given composite beam map at small radii, and are used to get uncertainties on the $T \rightarrow P$ leakage estimates derived in simulations presented in Chapter 6. These become poor noise estimates at larger radii, however, due to the varying spatial coverage limited by the coupling to the redirecting mirror, especially for *Keck* and BICEP Array.

In principle, the mean could be used instead of the median in each map pixel when

BICEP3 2017 Detector 607 Composite FFBM




Figure 4.8: Composite beam map of a sample BICEP3 95 GHz detector using component beam maps from the February 2017 observing campaign. The coverage in the composite map benefits from taking beam map observations at a wide variety of boresight rotation angles. Although the beam maps plotted here are truncated to a radius of 2° from the main beam for a clear view of the main beam shape, this composite map extends out to radius of 20° . This is a typical number of component maps used in a single-year composite map for BICEP3, but for *Keck* and BICEP Array the typical number is much lower (~ 10) due to the limited coverage of the redirecting mirror.

compositing, but in many cases the noise in the composite beam ends up being dominated by a handful of component maps with noise levels that largely vary with elevation, especially for the frequency bands that are more sensitive to the atmosphere (*Keck* 220 and 270 GHz). This is a targeted area of improvement in future work – getting clean composite and noise maps using the mean and standard deviation, respectively, would give us more reliable noise estimates for a composite beam out to any radius, regardless of the varying coverage.

4.2.6 BEAM WINDOW FUNCTION AND PIPELINE TRANSFER FUNCTION

One critical application of the per-detector composite beam maps is generating per-band, array-averaged beam window functions $B(\ell)$. These are essential for converting measurements of the CMB into true physical estimates of the on-sky signal. To illustrate, consider the temperature of the sky as observed by a telescope decomposed into spherical harmonics:

$$T_{\text{obs}}(\theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} B_{\ell m} a_{\ell m} Y_{\ell m}(\theta, \phi) \quad (4.3)$$

where $B_{\ell m}$ is the beam window function. To obtain the true values of the $a_{\ell m}$, the beam smoothing must be undone. This beam correction is done in practice in the calculation of the bandpower window functions, described in Section 6.1.4.

For each frequency band, we coadd all per-detector composite beams, exactly anal-

ogous to how the per-detector composites are constructed from the component beams. Before taking the 2D Fourier transform, we take the azimuthal average of the coadded beam (effectively averaging over all m in each ℓ bin, converting $B_{\ell m} \rightarrow B(\ell)$). The real part of the 2D Fourier transform is averaged in radial bins, obtaining a 1D beam window function $B(\ell)$. Presentation of the beam window functions for each BI-CEP/*Keck Array* band is deferred to Section 5.3.

At this step, we correct for the finite size of the aperture of the thermal chopper used in far-field beam map observations. For a given measured beam map, the observed beam is actually a convolution of the true detector beam response with a uniform circular aperture of diameter D . Because convolution in real space is mathematically equivalent to multiplication in Fourier space, we can correct for the aperture by dividing the observed $B(\ell)$ by the Fourier transform of the uniform circular aperture of the thermal chopper:

$$B(\ell)_{\text{true}} = B(\ell)_{\text{obs}} / \frac{2 J_1(\ell\theta_{\text{chop}})}{\ell\theta_{\text{chop}}} \quad (4.4)$$

where J_1 is the Bessel function of the first kind and θ_{chop} is the angular radius of the thermal chopper aperture as seen from the telescope. For a thermal chopper with $D = 24$ inches at a distance ~ 200 m away, the angular radius is $\theta_{\text{chop}} \simeq 0.087^\circ$. This correction increases the observed $B(\ell)_{\text{obs}}$ by 1.2 % at $\ell = 200$ and 7.1 % at $\ell = 500$.

To quantify the effect of the end-to-end FFBM analysis pipeline, we run simulated

observations of an “input beam” with any desired shape or characteristics. This input beam is observed with a scan strategy that matches real FFBM observations, and then is put through the same analysis pipeline used to analyze the real beam maps. The resulting beam window function can be compared to the true input value to estimate the transfer function.

The results of one of these simulations is shown in Fig. 4.9. The input beam is an ideal Airy disk with an intensity pattern given by

$$I(r) = \left(\frac{2 J_1(\pi D \sin(r) / \lambda)}{\pi D \sin(r) / \lambda} \right)^2 \quad (4.5)$$

where D is the aperture diameter, λ is the wavelength, and $r = \sqrt{x^2 + y^2}$. This input beam was chosen to resemble a *Keck* 220 GHz beam ($D = 0.264$ m), although in reality our far-field beams are not quite Airy disks, since we under-illuminate our apertures to minimize sidelobes. The input was modulated with a square wave, using a chop reference taken directly from real observations. Gaussian noise was injected at a level -30 dB from the main beam peak. The beam was observed at many different boresight angles and the usual ground and mirror masks were applied (although no ground signal was simulated here). Simulated component beam maps were then coadded into composite maps, and an array-average beam window function was calculated as with real data. This beam window function is compared to the known analytic function corresponding to the input Airy disk in the right-hand plot of Fig. 4.9. The

upgrade to the FFBM pipeline shows dramatic improvement in the fidelity of the end-to-end measurement, although it is still not perfect. Ideally, the binning step would be the dominant discrepancy between the input and simulated beam window function (i.e., the transfer function would approximately equal the pixel window function at all scales). This is true until roughly $\ell \sim 200$, where the pipeline transfer function (green line in Fig. 4.9) starts to deviate from the empirical pixel window function (blue line). This difference is less than 5 % below $\ell = 500$, but still indicates some effective beam smoothing in the pipeline.

These new simulations of FFBM observations are generally a powerful tool and potential area for further work. The transfer function can be compared using different scan strategies, modulation functions, and beam shapes, which could directly influence the real observation strategies chosen for FFBM observations.

4.3 NEAR-FIELD BEAM MAPS

Near-field beam maps (NFBM) are measured in a plane just above the window of the aperture, well within the far-field criterion for any frequency band. While far-field measurements are sensitive to the beam amplitude in the detector plane, measurements of the near-field beam pattern are primarily sensitive to the phase. They are one of the first calibration measurements performed immediately after a receiver cooldown, due to their value in checking the overall health of the detectors and opti-




Figure 4.9: Simulation of beam map observations to quantify the analysis pipeline transfer function. *Left:* The input synthetic beam map, which is an ideal Airy disk with a main beam width roughly matching a *Keck* 220 GHz beam. *Middle:* The composite beam resulting from observing the input beam map with a realistic scan strategy at multiple boresight rotation angles, and putting it through the full analysis pipeline as described in this chapter. Gaussian noise was injected at a -30 dB level relative to the main beam peak. *Right:* Comparison of the beam window function as output from this simulation against the true analytic window function of the input beam. The blue line shows the effect of only binning the input map into standard 0.1° square pixels, with no other operations, effectively quantifying the pixel window function. The red and green lines compare the output beam window functions of the analysis pipeline before and after the upgrades described throughout this chapter, respectively. Despite significant improvement, there is still a $< 5\%$ difference between the new pipeline transfer function and the pixel window function at $\ell = 500$. See the text for a full description of the simulation.

cal chain.

We simultaneously map the near-field response of all detectors in a receiver by mounting the apparatus on top of the aperture window, and scanning a chopped thermal source enclosure over the entire field of view. The apparatus is composed of two Velmex BiSlide automated translation stages, allowing freedom to scan in X and Y. The source enclosure contains a heated ceramic source chopped by a Thorlabs MC2000 optical chopper. A series of improvements were made to the thermal source in 2018-2019 – the ceramic is now coated in concrete-carbon mixture and thermal conductive putty, instead of a silicone-carbon mixture, allowing a safe increase in operating temperature from 225°C to 395°C. The variac power supply was replaced with a PID temperature controller, allowing quicker heating times and stable temperature conditions over 20+ hours. Improved heat sinking and insulation was installed to prevent the source enclosure itself from increasing noticeably in temperature. The aperture of the source enclosure is large enough such that any detector within 15° of the axis perpendicular to the aperture couples to the heated ceramic. Two varieties of NFBM apparatus exist – a smaller 20x20 inch mapper for the *Keck*-sized apertures, and a larger 30x30 inch mapper (called the “Mega Mapper”) for BICEP3 and BICEP Array. The setup and source enclosure can be seen in Fig. 4.10.

As with most calibration measurements, the detectors are biased on the Aluminum superconducting transition (~ 1.2 K) so that the high-temperature source does not saturate the detectors. The length of a NFBM observation varies with aperture size and




Figure 4.10: Near-field beam map measurement setup. *Left:* 30x30 inch NFBM apparatus installed on the BI-CEP Array 30/40 GHz receiver BA1 for initial testing at CalTech. *Right:* Close-up of the thermal source enclosure with front cover removed, showing the optical chopper, heated ceramic element, and insulation.

desired signal-to-noise (~ 1 -12 hours). The source moves in increments and pauses, usually for 4 – 10 seconds, at each step in the raster to allow sufficient integration over noise. The detector timestreams, optical chopper reference signal, and X-Y encoder positions are all recorded simultaneously for the duration of the measurement.

Features seen in the near-field beam maps have been used to diagnose imperfections in the detectors and optics numerous times throughout the *BICEP/Keck Array* program. Non-ideal beam pointing at the focal plane in some detectors leads to significant “beam steer” in the near field, when the main beam is severely truncated by the aperture stop. High beam steer is likely due to variation in the thickness of the interlayer dielectric near the edge of a tile. Small features in the near-field beam maps have also been seen, which arise from secondary reflections off the 4 K filters that are then refocused near the aperture plane. These small features contain $< 0.1\%$ of the total integrated beam power. For more details on these features, see [The BICEP2 and Keck Array Collaborations 2015a](#).

A set of example NFBMs taken on BICEP3 after the 2017 upgrade are shown in Fig. 4.11. The NFBM on the left is typical of most BICEP3 maps. The middle and right plots show sample detectors with a low and high amount of excessive edge taper, respectively. Although these few detectors with high edge taper are cut from CMB analysis and therefore do not contribute weight to the final CMB maps, special cases like these can be used to feed back on optical model design, and look for correlations with other optical performance metrics, such as optical loading from coupling to exter-




Figure 4.11: Sample near-field beam maps for three different detectors in BICEP3 taken in the 2016-2017 austral summer. All are plotted in linear, peak-normalized, demodulated feedback units. The x and y axes mark the distance in the plane of the near-field beam mapper, which sits ~ 0.5 m above the aperture. The samples shown include a nominal NFBM (left), a NFBM with a slightly anomalously high edge taper (middle), and a NFBM with very high edge taper (right). The very high edge taper such as that shown on the right is seen only in ~ 5 (out of 1200) optically-coupled detector pairs. Optical modeling of anomalous edge taper and comparisons against other optical performance metrics are discussed in Chapter 7.

nal warm baffling, far-field ellipticity, and far-field beam mismatch. This is explored further in Chapter 7.

4.4 FAR-SIDELOBE MAPS

All BICEP/Keck Array receivers use two levels of warm baffling, to ensure that any ray must diffract twice to couple to the ground. The first is a comoving absorptive baffle mounted directly to the window of the receiver. The second is a large, fixed ground shield around the base of the mount. Therefore, any excess power in the far sidelobes (roughly defined as the region $\gtrsim 20^\circ$ from the main beam) should couple to ambient-temperature absorber or redirect via the ground shield to cold sky. However, this power still increases the loading of the detectors, and if polarized, could still

lead to leakage that may be difficult to constrain, depending on where it terminates. We therefore take measurements using a high-powered noise source to map the far-sidelobe region of the detector beams.

Constructing a mirror large enough to couple the entire beam of a single detector to the source at the opposite building is an infeasible task, so we are forced to take a different approach compared to the FFBM measurement procedure. We use the mast on the same building, which rises far above the lip of the ground shield. Since the source is only roughly 10 m away in this case, it is technically not in the far field, but we expect the impact of this effect on the far-sidelobe amplitude and shape to be small. The source is the broad-spectrum noise source (BSNS) described in Section 4.1.1, which has two variable attenuators to give the high (~ 70 dB) dynamic range needed for this measurement.

A typical FSL schedule takes 380° scans in azimuth, with an elevation range of 34° in 0.5° steps, all repeated over five boresight rotation angles. The measurement is often repeated with both the comoving forebaffle on and off, to verify the amount of power intercepting the forebaffles. A waveguide twist can be placed before the BSNS output horn that couples to free space, in order to take measurements at two orthogonal source polarizations. Fig 4.12 shows a photo of a FSL measurement being done for the BICEP Array 30/40 GHz receiver with the forebaffle off.

Three different power settings are used to map out the entire beam at each boresight rotation angle, where the power settings are changed by adjusting the attenua-




Figure 4.12: BICEP Array 30/40 GHz receiver BA1 performing far-sidelobe observations in March 2021. Note the forebaffle is removed for this observation. The measurement is made by rastering over a broad-spectrum noise source mounted on a mast on the same building as the receiver, which peaks far above the lip of the ground shield. Photo courtesy of Brandon Amat.




Figure 4.13: Sample far-sidelobe map of a BICEP3 detector, made by stitching together measurements from all three power settings and coadding maps made at both source polarization orientations. The forebaffle was on for this measurement, so this is a representation of the true beam response on sky as it is during CMB observations. The main beam and extent of the BICEP3 aperture is clearly seen. The feature just below the main beam in the map is the “buddy beam”, formed by reflection at the surface of a 50 K filter (after both 4 K lenses) which travels back to the focal plane at a location exact opposite the boresight of the main detector, then reflects again and emerges as a well-formed beam on sky (see Fig. 5.9 for a diagram)..

tion in the BSNS. A “low” power setting maps the main beam, “medium” maps the mid sidelobes, and “high” maps the far sidelobes. The maps made with each setting are stitched together to create a single map for each detector. The maps at each polarization (made with and without the waveguide twist installed) can be coadded together to create an effective unpolarized far sidelobe map. Buddy beams are discussed in Section 5.5.

Far sidelobe maps of *Keck* receivers in 2012-2013 revealed a bright ring feature

roughly 20° from the main beam. It was soon discovered that these ring features were due to specular reflections off the interior of the upper optics tube, despite being covered in microwave absorber. All receivers eventually had interior baffling installed to prevent these specular reflections, and subsequent FSL measurements showed these rings features were eliminated (Karkare 2017).

More recently, far sidelobe measurements of BICEP3 initially showed that the “buddy beam” amplitudes – beams formed by specular reflection off a 50 K filter, which travels back through the optics and emerges as a well-formed beam on the opposite side of the boresight – were anomalously high in Tile 1 compared to the other nineteen tiles. This feature was later confirmed in independent FFBM measurements. The large buddy beam amplitudes in Tile 1 were due to the lack of a detector tile at the location opposite the boresight. Instead of another tile, Mylar polyester film (used for magnetic and thermal shielding) reflects the signal back through the optics, leading to larger and widely varying buddy beam amplitudes. Because of this optical anomaly associated specifically with this tile, Tile 1 is removed from the final CMB maps in the BK18 data set.

*Astronomy? Impossible to understand and
madness to investigate.*

Sophocles

5

Far-Field Beam Characterization

Armed with the beam measurement procedures described in the previous chapter, we can now begin to characterize the beam patterns of BICEP/*Keck Array* receivers. The results of from the far-field beam analysis discussed in Section 4.2 have a wide variety of applications that are discussed throughout the rest of this dissertation. These applications include:

- **Array-averaged beam window functions:** These are used to smooth any

external map used for simulations, and to evaluate the per-bandpower beam suppression correction applied to the power spectra.

- **Per-detector composite beams and beam statistics:** These are the high-fidelity beams made by compositing all measurements together for a single detector, and the associated 2D elliptical Gaussian fit parameters. These have many applications, including validation of detector performance and of the end-to-end optical chain.
- **Beam model feedback and validation:** Optical models are a valuable tool for designing telescope optics and predicting their performance, but these models require validation against real measurements. The vast library of archival BICEP/*Keck Array* far-field beam maps is well-suited for such validation of small-aperture telescope beam models.
- **Simple metric evaluation and correlations:** These metrics can be evaluated using only beam measurements or models, and can help understand the relation between, for example, the aperture stop illumination and the differential beam power in the far field. This is particularly valuable for experiments in the design phase, where full timestream simulations cannot be done.
- **Temperature-to-polarization leakage simulations:** perhaps the most important use of FFBM data, real beam maps are used in full-timestream simulations of a temperature-only, noiseless sky to predict the amount of leakage in a real data set.

Chapter 7 discusses beam models and metrics, and Chapter 6 presents the $T \rightarrow P$ leakage simulations and results. In this chapter, we focus on the array-averaged beam window functions, and beam statistics, and other notable features seen in the beam maps. We emphasize the results of beam measurements on receivers that observed the CMB in 2016, 2017, and 2018, i.e. the new receiver-years included in the BK18 data set. This includes three years of BICEP3 at 95 GHz and twelve receiver-years of *Keck* 220/230 GHz observations.

5.1 ALTERNATE POLARIZATION AXIS DEFINITION

Before discussing specific results, it is important to note that the polarization axes A and B used throughout this dissertation and previous BICEP/*Keck Array* publications are fixed to the frame of the detector tile, and these tiles may be installed in the focal plane at any orientation. For a reminder of the orientation of the tiles as installed in *Keck* and BICEP3, see Fig. 4.6. For *Keck* and BICEP2 (which had the same tile orientations as *Keck*), Tiles 1 and 2 have the same orientation, and Tiles 3 and 4 are rotated 180° with respect to Tiles 1 and 2. This means the A and B polarization axes are consistently oriented within a focal plane and between focal planes (since parallel and antiparallel are the same to a linearly polarized wave). This is not true for BICEP3, which has tiles installed at 0° , 90° , 180° , and 270° . In particular, BICEP3 Tiles 17, 12, 18, 13, and 8 have their A polarization axes aligned parallel to the $\theta = 0$ axis, as they are in *Keck*, but the other fifteen tiles have their A polarization axes rotated by 90° (or 270°) with respect to $\theta = 0$. Here we introduce an alternate polarization axis definition, V and H :

$$(V, H) = \begin{cases} (A, B), & \text{Keck and BICEP3 Tiles 17, 12, 18, 13} \\ (B, A), & \text{All other BICEP3 Tiles} \end{cases} \quad (5.1)$$

These alternate axes are fixed to the instrument and independent of tile orientation. The use of these axes as opposed to the standard (A, B) will be specially marked in every case throughout this work.

5.2 2D ELLIPTICAL GAUSSIAN FIT PARAMETERS

Here we present the results of the 2D elliptical Gaussian fit procedure described in Section 4.2.3, in the form of tables and FPU plots. In the tables, we show the statistics for one receiver-year at a time, presented as (FPU Median) \pm (FPU Scatter) \pm (Measurement Uncertainty). The “FPU Median” is the median over all good detectors (detectors that contribute weight to the final CMB map) in that receiver, and the “FPU Scatter” is an estimate of the detector-to-detector spread within that receiver. The uncertainty of a parameter estimate for a single detector is the spread in fit values over all component maps that form that detector’s composite map. The “Measurement Uncertainty” presented in the tables is the median of those per-detector spreads over all detectors in the FPU. In general, the FPU scatter is larger than the measurement uncertainty, meaning we can confidently quantify trends and variations across the FPU presented in the FPU plots.

The beam parameter estimates from 2016, 2017, and 2018 are found in Table 5.1, Table 5.2, and Table 5.3, respectively. All differential beam parameters in these tables are evaluated as $A - B$. The beamwidths σ have been corrected for the finite size of

Parameter	BICEP3 (95 GHz)	Keck Rx0 (220 GHz)	Keck Rx1 (230 GHz)
$\sigma(^{\circ})$	$0.163 \pm 0.002 \pm 0.001$	$0.139 \pm 0.002 \pm 0.003$	$0.134 \pm 0.002 \pm 0.003$
p	$0.013 \pm 0.021 \pm 0.021$	$0.001 \pm 0.014 \pm 0.035$	$0.000 \pm 0.017 \pm 0.033$
c	$-0.004 \pm 0.017 \pm 0.021$	$-0.008 \pm 0.031 \pm 0.040$	$-0.003 \pm 0.023 \pm 0.035$
$d\sigma(^{\circ})$	$0.000 \pm 0.0001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp	$-0.007 \pm 0.016 \pm 0.004$	$-0.009 \pm 0.007 \pm 0.003$	$-0.016 \pm 0.009 \pm 0.005$
dc	$0.000 \pm 0.007 \pm 0.003$	$-0.003 \pm 0.003 \pm 0.002$	$0.001 \pm 0.005 \pm 0.004$
$dx(^{\prime})$	$-0.06 \pm 0.13 \pm 0.04$	$-0.13 \pm 0.26 \pm 0.05$	$0.30 \pm 0.24 \pm 0.06$
$dy(^{\prime})$	$-0.03 \pm 0.16 \pm 0.04$	$-0.17 \pm 0.21 \pm 0.05$	$-0.49 \pm 0.20 \pm 0.05$
Parameter	Keck Rx2 (220 GHz)	Keck Rx3 (230 GHz)	Keck Rx4 (150 GHz)
$\sigma(^{\circ})$	$0.139 \pm 0.002 \pm 0.003$	$0.135 \pm 0.002 \pm 0.004$	$0.201 \pm 0.002 \pm 0.004$
p	$0.000 \pm 0.021 \pm 0.032$	$-0.001 \pm 0.018 \pm 0.039$	$0.005 \pm 0.019 \pm 0.028$
c	$0.001 \pm 0.034 \pm 0.032$	$0.000 \pm 0.027 \pm 0.043$	$0.003 \pm 0.014 \pm 0.030$
$d\sigma(^{\circ})$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp	$-0.013 \pm 0.004 \pm 0.002$	$-0.018 \pm 0.004 \pm 0.004$	$-0.020 \pm 0.004 \pm 0.003$
dc	$-0.002 \pm 0.004 \pm 0.002$	$0.004 \pm 0.004 \pm 0.004$	$-0.003 \pm 0.002 \pm 0.003$
$dx(^{\prime})$	$-0.01 \pm 0.19 \pm 0.06$	$-0.46 \pm 0.16 \pm 0.05$	$0.23 \pm 0.38 \pm 0.10$
$dy(^{\prime})$	$-0.20 \pm 0.20 \pm 0.06$	$0.40 \pm 0.11 \pm 0.05$	$-0.13 \pm 0.32 \pm 0.10$

Table 5.1: 2D elliptical Gaussian fit parameters from measurements taken prior to the 2016 observing season. Differential parameters are evaluated as $A - B$.

the thermal source aperture. The change in these measurements from year-to-year is generally smaller than the measurement uncertainties, with a couple notable exceptions: between the 2016/2017 observing seasons, the BICEP3 low-pass edge filters were replaced after they delaminated during the 2016 season, slightly changing the beam characteristics. Additionally, the 270 GHz receivers in 2017 and 2018 are different focal plane units (Rev. E12 and E13, respectively).

For BICEP3, the differential beam parameters are also evaluated using the alternate polarization axes $V-H$ instead of the traditional $A-B$. These are shown in Table 5.4. The most notable change is in differential ellipticity plus, dp , which is positive for all tiles in this alternate axis definition. The root cause for this uniformly positive dp is

Parameter	BICEP3 (95 GHz)	Keck Rx0 (220 GHz)	Keck Rx1 (230 GHz)
$\sigma(^{\circ})$	$0.161 \pm 0.003 \pm 0.001$	$0.140 \pm 0.003 \pm 0.003$	$0.135 \pm 0.002 \pm 0.002$
p	$0.008 \pm 0.023 \pm 0.021$	$0.002 \pm 0.015 \pm 0.032$	$0.003 \pm 0.019 \pm 0.023$
c	$-0.010 \pm 0.020 \pm 0.019$	$-0.008 \pm 0.033 \pm 0.037$	$-0.003 \pm 0.022 \pm 0.026$
$d\sigma(^{\circ})$	$0.000 \pm 0.0001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp	$-0.003 \pm 0.011 \pm 0.002$	$-0.009 \pm 0.007 \pm 0.002$	$-0.016 \pm 0.010 \pm 0.004$
dc	$-0.003 \pm 0.004 \pm 0.002$	$-0.004 \pm 0.003 \pm 0.002$	$0.001 \pm 0.005 \pm 0.003$
$dx(^{\prime})$	$-0.06 \pm 0.12 \pm 0.05$	$-0.16 \pm 0.26 \pm 0.04$	$0.39 \pm 0.25 \pm 0.04$
$dy(^{\prime})$	$0.00 \pm 0.13 \pm 0.05$	$-0.17 \pm 0.23 \pm 0.03$	$-0.48 \pm 0.23 \pm 0.04$
Parameter	Keck Rx2 (220 GHz)	Keck Rx3 (230 GHz)	Keck Rx4 (270 GHz)
$\sigma(^{\circ})$	$0.140 \pm 0.002 \pm 0.002$	$0.136 \pm 0.002 \pm 0.003$	$0.116 \pm 0.002 \pm 0.003$
p	$0.003 \pm 0.022 \pm 0.020$	$0.000 \pm 0.017 \pm 0.036$	$0.003 \pm 0.020 \pm 0.022$
c	$-0.001 \pm 0.040 \pm 0.027$	$0.002 \pm 0.034 \pm 0.043$	$-0.002 \pm 0.041 \pm 0.036$
$d\sigma(^{\circ})$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.001 \pm 0.001 \pm 0.001$
dp	$-0.012 \pm 0.005 \pm 0.002$	$-0.018 \pm 0.004 \pm 0.002$	$-0.015 \pm 0.006 \pm 0.005$
dc	$-0.002 \pm 0.004 \pm 0.002$	$0.004 \pm 0.004 \pm 0.003$	$-0.011 \pm 0.002 \pm 0.004$
$dx(^{\prime})$	$-0.02 \pm 0.19 \pm 0.04$	$-0.50 \pm 0.17 \pm 0.04$	$0.69 \pm 0.18 \pm 0.04$
$dy(^{\prime})$	$-0.21 \pm 0.22 \pm 0.04$	$0.41 \pm 0.12 \pm 0.04$	$0.26 \pm 0.23 \pm 0.04$

Table 5.2: 2D elliptical Gaussian fit parameters from measurements taken prior to the 2017 observing season. Differential parameters are evaluated as $A - B$.

Parameter	BICEP3 (95 GHz)	Keck Rx0 (220 GHz)	Keck Rx1 (230 GHz)
$\sigma(^{\circ})$	$0.163 \pm 0.003 \pm 0.001$	$0.140 \pm 0.003 \pm 0.003$	$0.135 \pm 0.002 \pm 0.002$
p	$0.000 \pm 0.022 \pm 0.010$	$0.001 \pm 0.015 \pm 0.025$	$0.002 \pm 0.016 \pm 0.022$
c	$-0.007 \pm 0.021 \pm 0.012$	$-0.009 \pm 0.031 \pm 0.030$	$-0.003 \pm 0.018 \pm 0.026$
$d\sigma(^{\circ})$	$-0.001 \pm 0.0001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp	$-0.003 \pm 0.011 \pm 0.003$	$-0.010 \pm 0.006 \pm 0.002$	$-0.016 \pm 0.009 \pm 0.004$
dc	$-0.003 \pm 0.004 \pm 0.003$	$-0.004 \pm 0.003 \pm 0.002$	$0.001 \pm 0.004 \pm 0.003$
$dx(^{\prime})$	$-0.06 \pm 0.12 \pm 0.06$	$-0.12 \pm 0.27 \pm 0.04$	$0.38 \pm 0.25 \pm 0.04$
$dy(^{\prime})$	$0.01 \pm 0.13 \pm 0.06$	$-0.16 \pm 0.22 \pm 0.03$	$-0.49 \pm 0.22 \pm 0.04$
Parameter	Keck Rx2 (220 GHz)	Keck Rx3 (230 GHz)	Keck Rx4 (270 GHz)
$\sigma(^{\circ})$	$0.139 \pm 0.002 \pm 0.002$	$0.136 \pm 0.002 \pm 0.003$	$0.115 \pm 0.002 \pm 0.003$
p	$0.003 \pm 0.018 \pm 0.019$	$0.002 \pm 0.017 \pm 0.029$	$0.001 \pm 0.021 \pm 0.028$
c	$0.005 \pm 0.031 \pm 0.025$	$0.005 \pm 0.031 \pm 0.034$	$0.004 \pm 0.038 \pm 0.035$
$d\sigma(^{\circ})$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp	$-0.013 \pm 0.004 \pm 0.001$	$-0.018 \pm 0.004 \pm 0.002$	$-0.016 \pm 0.006 \pm 0.005$
dc	$-0.002 \pm 0.004 \pm 0.002$	$0.004 \pm 0.004 \pm 0.002$	$-0.008 \pm 0.010 \pm 0.005$
$dx(^{\prime})$	$-0.01 \pm 0.19 \pm 0.04$	$-0.46 \pm 0.18 \pm 0.04$	$-0.78 \pm 0.15 \pm 0.03$
$dy(^{\prime})$	$-0.21 \pm 0.20 \pm 0.04$	$0.40 \pm 0.10 \pm 0.03$	$0.23 \pm 0.26 \pm 0.03$

Table 5.3: 2D elliptical Gaussian fit parameters from measurements taken prior to the 2018 observing season. Differential parameters are evaluated as $A - B$.

Parameter	BICEP3 (95 GHz) 2016	BICEP3 (95 GHz) 2017	BICEP3 (95 GHz) 2018
$d\sigma_{VH}(\circ)$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$	$0.000 \pm 0.001 \pm 0.001$
dp_{VH}	$0.012 \pm 0.010 \pm 0.004$	$0.008 \pm 0.008 \pm 0.002$	$0.008 \pm 0.008 \pm 0.003$
dc_{VH}	$0.004 \pm 0.005 \pm 0.004$	$0.003 \pm 0.003 \pm 0.002$	$0.003 \pm 0.003 \pm 0.003$
$dx_{VH}(\prime)$	$0.03 \pm 0.14 \pm 0.04$	$0.04 \pm 0.14 \pm 0.05$	$0.03 \pm 0.14 \pm 0.06$
$dy_{VH}(\prime)$	$0.07 \pm 0.14 \pm 0.04$	$0.01 \pm 0.13 \pm 0.05$	$0.00 \pm 0.13 \pm 0.06$

Table 5.4: BICEP3 differential fit parameters using $V - H$, as defined in Eq. 5.1. The most notable change (compared to the traditional $A - B$) is in differential ellipticity plus, which is now positive across all tiles. This indicates that the differential ellipticity originates from the modules themselves, and is likely due to the slot antenna geometry.

likely related to the detector modules, specifically the geometry of the slot antennas themselves. Although the beam for a single polarized detector emerging from the focal plane is formed predominantly from the sum of ideal point sources at the location of each slot antenna, the shapes of the individual antennas (which are very asymmetric) are expected to contribute some residual ellipticity that is aligned with the antenna orientation. This leads to differential ellipticity which is parallel/perpendicular to the A/B axes, which is what we observe.

Among all differential parameters presented in these tables, differential pointing is generally the most significant and leads to the most amount of $T \rightarrow P$ leakage, if unaccounted for in the analysis¹. Differential pointing can arise a number of different ways. Nonuniform niobium contamination can lead to an effective vertical gradient in the film layer of the detector wafer, which can steer the beams off boresight, where the amount of steer may be different between both detector polarizations. This would produce differential pointing predominantly in the vertical direction (parallel

¹Differential gain is also generally dominant, but this is not measured in beam maps taken on a different TES transition (aluminum).

to the x' axis for *Keck*), since the signal combination in the horizontal direction of the microstrip network happens instantaneously, whereas a larger phase gradient is accumulated before the vertical component signal is combined (see [BICEP/Keck Array and SPIDER Collaborations 2015](#) and [O’Brient et al. 2012](#)). This is just one way to induce differential pointing in the far field – another example is imperfections or birefringence in the optics coupling to pointing mismatch in the near field beam patterns. In general, trends in the differential pointing across an entire focal plane are indicative of problems with the optics, while clear tile-to-tile variations indicate detector fabrication imperfections.

The variation in differential pointing across the focal plane for each receiver in the 2016-2018 observing seasons is shown in Fig. 5.1 (*Keck*) and Fig. 5.2 (BICEP3). There is clear repeatability in these measurements from year to year in receivers that were optically unchanged (*Keck* Rx0/1/2/3 2016-2018 and BICEP3 2017-2018). The variation within the BICEP3 FPU from tile to tile is somewhat random, but is likely due to the optical chain and not the detectors, because of the consistency from 2016 to 2017 despite many detector modules being replaced or moved to a different position in the 2016-2017 austral summer. We emphasize, however, that the differential beam parameters for BICEP3 are overall much improved compared to *Keck*, as can be seen in the tables and figures shown throughout this section.

Differential ellipticity measurements are of particular importance – they are the only deprojection coefficients whose values are taken directly from FFBM data, rather

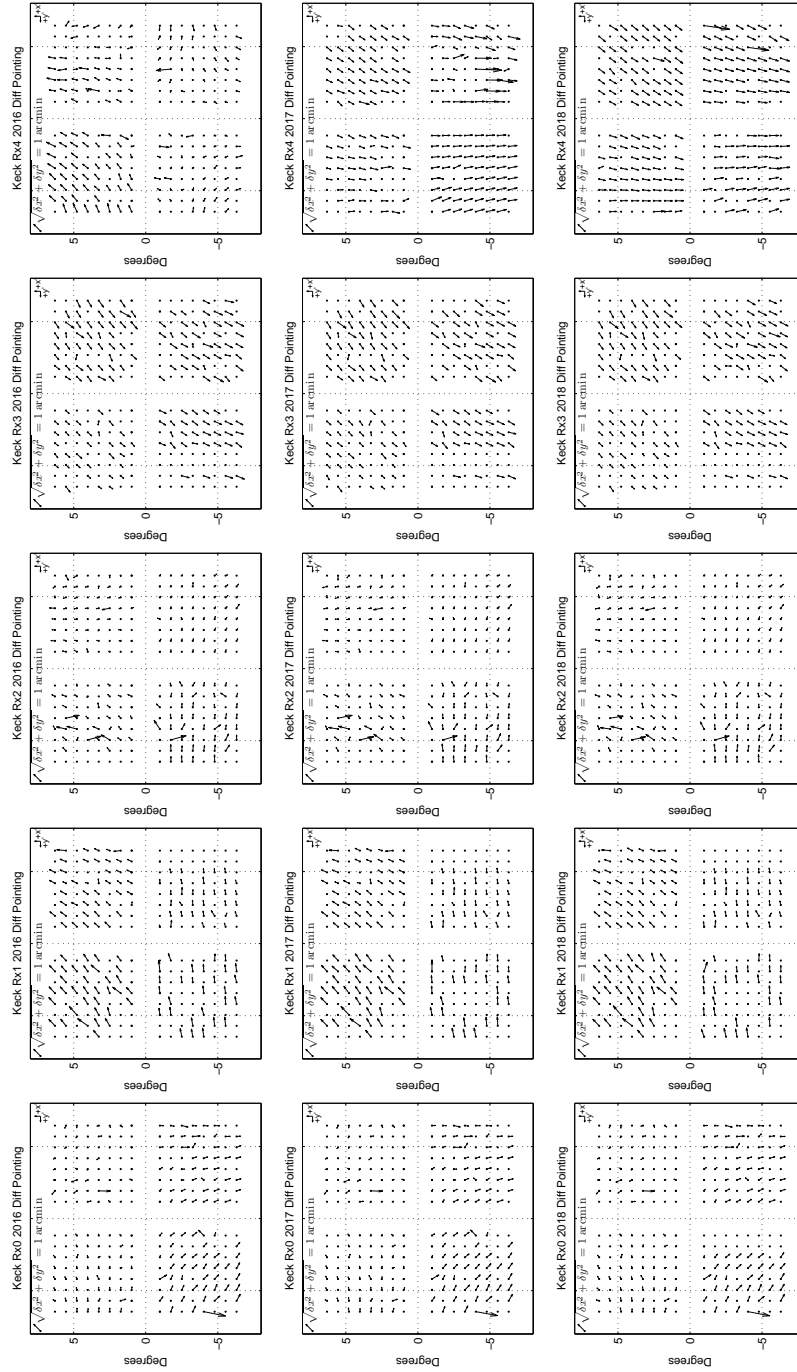


Figure 5.1: FPU quiver plots of the differential pointing taken from *Keck* FFBM observations in 2016 (top row), 2017 (middle row), and 2018 (bottom row). The arrows point from the A beam center to the B beam center, with the base of the arrow centered on the ideal beam centroid location. The arrows are scaled by a factor of 40 for clarity.

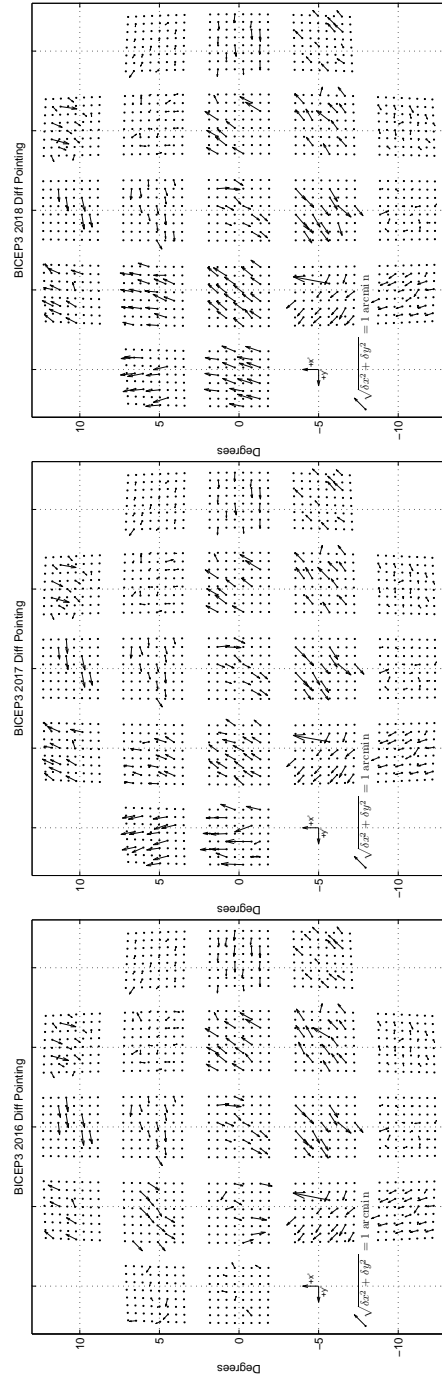


Figure 5.2: FPU quiver plots of the differential pointing taken from BICEP3 FFBM observations in 2016 (left), 2017 (middle), and 2018 (right). The arrows point from the A beam center to the B beam center, with the base of the arrow centered on the ideal beam centroid location. The arrows are scaled by a factor of 60 for clarity.

than from regression against deprojection templates, due to real TE correlations in the CMB (more on this in Section 6.1.2). This is one critical area where FFBM measurements directly intersect with results from primary CMB analysis. Differential ellipticity can arise a number of ways in the detectors and optical chain. As discussed earlier in this Section, the geometry of the slot antennas imparts an ellipticity associated with the antenna orientation, leading to an inherent differential ellipticity. An ellipticity in the phase or amplitude of the beam at the plane of the aperture can also give differential ellipticity in the far field.

FPU plots of differential ellipticity are shown for *Keck* in Fig. 5.3 and for BICEP3 in Fig. 5.4. For *Keck*, there is a clear negative differential ellipticity across each focal plane, likely associated with the slot antenna geometry. For BICEP3, the trend is similar, once one considers that fifteen tiles are rotated by 90° with respect to the other five. However, upon inspecting the $V - H$ differential ellipticity (flipping the signs of the $A - B$ values for those fifteen tiles), the sign of differential ellipticity is positive across the FPU, which is the opposite of *Keck*. This is unexpected if the dominant contribution is in fact antenna geometry, and is still being investigated.

5.3 ARRAY-AVERAGED BEAMS AND BEAM WINDOW FUNCTIONS

Using the per-detector composite far-field beam maps described in Section 4.2.5, we stack all the beams to form a single, high signal-to-noise beam for each frequency



Figure 5.3: FPU plots of the differential ellipticity taken from *Keck* FFBM observations in 2016 (top row), 2017 (middle row), and 2018 (bottom row). The ellipticities are scaled by a factor of 75 for clarity. The differences are evaluated as $A - B$.




Figure 5.4: FPU plots of the differential ellipticity taken from BICEP3 FFBM observations in 2016 (left), 2017 (middle), and 2018 (right). The ellipticities are scaled by a factor of 75 for clarity. The differences are evaluated as $A - B$.

band. After this final coadding, the noise floors can get as low as -60 dB from the peak, giving a signal-dominated measurement out to a radius of 10-15° from the main beam in the case of BICEP3.

The array-averaged composite beams for all BICEP/*Keck Array* frequency bands made with the newest beam map analysis pipeline are shown in Fig. 5.5. The beams are plotted in units of dBi (dB relative to an isotropic radiator). Along with the clear sidelobes in each beam, crosstalk beams are clearly visible in a “+” pattern surrounding the main beam. Crosstalk is an artifact of the time-domain SQUID-based multiplexed readout discussed in Section 2.2. These crosstalk beams are discussed in greater detail in Section 5.6. Different noise levels are also clear from comparing the different beams – BICEP3 benefits from having a redirecting mirror large enough to couple all detectors to the source in each FFBM observation schedule. This gives many more component maps in each per-detector composite map, allowing a greater signal-to-noise that is apparent when comparing to the high-frequency *Keck* beams which have a similar beamwidth to BICEP3.

The array-averaged beams are then azimuthally averaged and Fourier transformed in order to create 1D beam window functions, $B(\ell)$. These are of critical importance to the primary CMB analysis – any external maps observed in simulations (such as those from *Planck* or WMAP) must be convolved with the best representation of our beam in each frequency band before making simulated observations. We do this convolution in Fourier space, by multiplying the external map with the $B(\ell)$. The beam




Figure 5.5: Array-averaged beams for each *BICEP/Keck Array* frequency band. All are plotted in dBi (dB relative to an isotropic radiator). These are made by coadding all per-detector composite beams from each band. Multiple sidelobes are clearly visible, as well as crosstalk beams arising from the time-domain SQUID-based multiplexed readout.




Figure 5.6: Beam window functions $B(\ell)$ for each *BICEP/Keck Array* frequency band, made by taking the Fourier transform of the azimuthally averaged version of the beams shown in Fig. 5.5. The *Keck* 220 GHz band is omitted for clarity, due to its similarity to *Keck* 230 GHz. The light-colored bands represent the spread in $B(\ell)$ evaluated from the per-detector composite beams that comprise the array-averaged beams. All $B(\ell)$ have been corrected for the finite size of the thermal source aperture.

window functions for each frequency band are shown in Fig. 5.6. The same method to evaluate the array-averaged $B(\ell)$ can be applied to each of the per-detector composite beams. The light-colored bands plotted represent the spread in $B(\ell)$ values over these composite beams in each frequency band. All $B(\ell)$ have been corrected for the finite size of the thermal source aperture, using Eq. 4.4.

5.4 VALIDATION AGAINST CMB DATA

In this section we describe two different methods for validating the data derived from far-field beam maps against CMB observations. First, we can compare the 2D elliptical Gaussian fit parameters described in Section 5.2 with the deprojection coefficients derived from regression of CMB data against deprojection templates constructed from linear combinations of the *Planck* temperature map and its derivatives. Second, an effective beam window function can be extracted from the CMB itself, and compared to the $B(\ell)$ shown in Section 5.3.

5.4.1 DEPROJECTION COEFFICIENTS VS DIFFERENTIAL BEAM PARAMETERS

Deprojection is an analysis technique that filters out the leading-order modes of $T \rightarrow P$ leakage stemming from mismatched main beams within a detector pair (see Section 6.1.2 for a full description of deprojection). Deprojection templates are constructed from linear combinations of the *Planck* T map and its derivatives, and these templates are observed using real detector pointing timestreams. The real CMB timestreams are regressed against these observed templates to obtain deprojection coefficients, which quantify the amount of power in each difference mode. This allows a comparison of two independently-obtained measurements of differential beam parameters – those obtained from FFBM measurements, and those obtained from the deprojection coefficients.

2D histograms of the deprojection coefficients vs FFBM differential beam parameters for BICEP3 and *Keck* 220 and 230 GHz are shown in Fig. 5.7. For BICEP3, the differences are taken as $V - H$ as defined in Eq. 5.1 (for *Keck* this makes no difference). Note that differential gain is not extracted from the beam maps – the x-axis value instead comes from the “abscal” analysis done by taking correlations between *Planck* and per-detector temperature maps. The agreement is very good, although the deprojection coefficients are noisier than the FFBM parameters, especially for differential beamwidth and differential ellipticity. There is a clear bias in differential ellipticity plus – this is due to real, nonzero TE correlations in the CMB biasing the deprojection coefficient estimates. This is why we do not deproject differential ellipticity. Instead, we scale the deprojection templates by the value taken from FFBMs and subtract them from the detector timestreams.

5.4.2 BEAM WINDOW FUNCTIONS FROM THE CMB

To convert the raw CMB map pixel values into meaningful units (μK_{CMB}), we perform an absolute calibration (“abscal”) by cross-correlating the power spectra of our T maps against external temperature maps that are already in physical units, such as those from *Planck*. This calibration done for a *Keck* temperature map \hat{K} , for example, takes the form:

$$\alpha(\ell) = \frac{\hat{P}_{\text{ref}} \hat{P}_{\text{cal}}}{\hat{P}_{\text{ref}} \hat{K}} \quad (5.2)$$

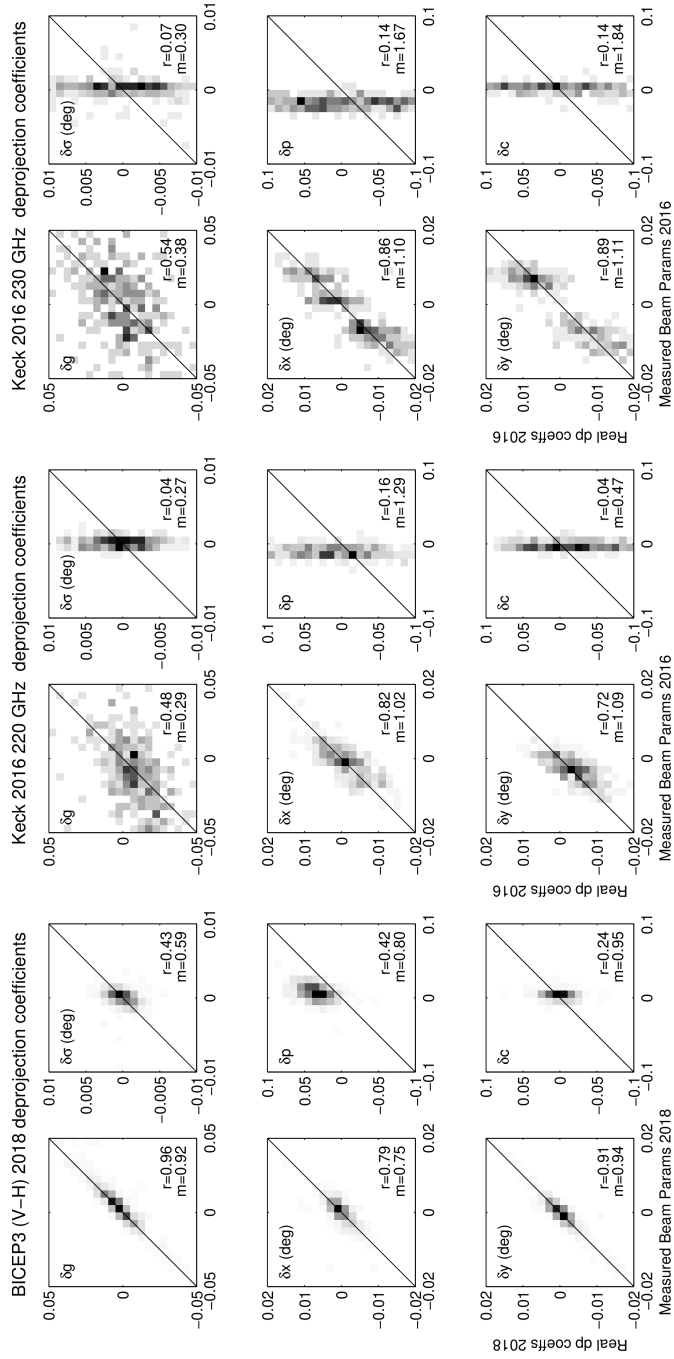


Figure 5.7: 2D histograms of deprojection coefficients vs. FFBM differential beam parameters for BICEP3 2018 (left), Keck 220 GHz 2016 (middle), and Keck 230 GHz 2016 (right). The differences are taken as $V - H$ for BICEP3 as defined in Eq. 5.1. The correlation coefficient r (where 1 is perfect correlation, 0 is no correlation, and -1 is perfect anticorrelation) and the slope m of the best fit line are shown. There is a notable bias in differential ellipticity plus ($\delta \hat{p}$) deprojection coefficients for all receivers, due to real TE correlations in the CMB. In general, differential beamwidth ($\delta \sigma$) and differential ellipticity ($\delta \hat{p}$, δc) values from deprojection are noisy compared to FFBM, leading to a wider distribution along the y-axis in each case.

where \hat{P}_{ref} is the external reference *Planck* T map at the nearest relevant frequency, and \hat{P}_{cal} is a second external T map at a different nearby frequency, which is needed to avoid self-correlating noise in the numerator. Note that each of these external maps are smoothed with the nominal *Keck* beam window function for that frequency, usually taken from the beam maps, as shown in Fig. 5.6. Any mismatch between the real *Keck* beam window function (inherently built into the map \hat{K}) and the $B(\ell)$ applied to the external maps will result in a non-uniform abscal $\alpha(\ell)$ in harmonic space. These slopes in $\alpha(\ell)$ vs ℓ were seen in the original BICEP3 and *Keck* 220 GHz abscal plots, indicating the $B(\ell)$ as derived from FFBM were uniformly over-estimating the true size of the beams. This prompted 1) a check that the redirecting mirrors using in FFBM observations were truly flat (which was found to be true, see Fig. 4.2 and Section 7.3), and 2) an audit of the FFBM analysis code, which led to the discovery of over-smoothing that was eliminated in the analysis pipeline upgrade detailed in Section 4.2.

While the FFBM analysis pipeline upgrade work was being implemented, we decided to use a beam profile derived from the abscal slope itself for BICEP3 and *Keck* 220 GHz. This “ $B(\ell)$ derived from the CMB” is constructed in such a way that the corrected abscal factors become flat. The procedure is shown for *Keck* 220 GHz in Fig. 5.8. The abscal as derived using the $B(\ell)$ from the original FFBM analysis pipeline is shown in blue in the right-hand plot, with error bars denoting an estimate of uncertainty due to experimental noise. There is a clear slope, indicating a mismatch be-

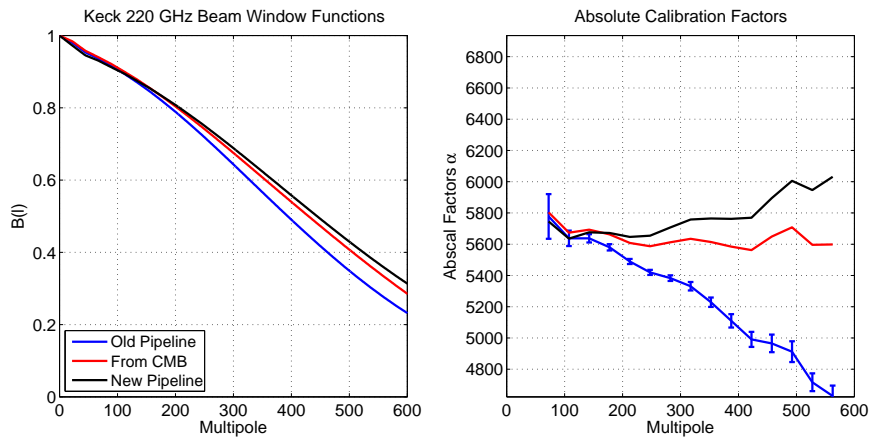


Figure 5.8: *Keck* 220 GHz beam window functions (left) and absolute calibration factors $\alpha(\ell)$ (right), as evaluated using the $B(\ell)$ derived from the old FFBM analysis pipeline (blue), the $B(\ell)$ as derived from the CMB (which leads to flat abscale vs ℓ , by construction; red), and the $B(\ell)$ derived from the updated FFBM analysis pipeline (black). Error bars in the abscale plot are an estimate of the uncertainty on $\alpha(\ell)$ due to experimental noise (the same uncertainties apply to the red and black lines, but are omitted for clarity). The “flatness” of the black line indicates that the $B(\ell)$ derived from the updated FFBM analysis pipeline closely matches the true beam shape.

tween the old pipeline $B(\ell)$ and the true beam. The red line shows the abscale after applying a correction that is deliberately constructed to make the abscales flat vs ℓ , which gives us the $B(\ell)$ from the CMB shown in the left plot. The black lines show the $B(\ell)$ and corresponding abscales when using the $B(\ell)$ derived from the updated FFBM analysis pipeline. The black line in the abscale is much more flat than the blue, indicating that the updated FFBM analysis pipeline results in a beam that is much closer to reality.

5.5 GHOST “BUDDY” BEAMS

Generally, a “ghost” beam is a secondary beam, which could be compact or diffuse, that arises due to internal reflections off of an optical element that eventually escapes the receiver and couples to the sky (in a time-reverse sense). We anti-reflection (AR) coat both surfaces of each optical element to minimize this effect, but some ghosting can still be seen in the beam maps. The most prevalent ghost beams, and the topic of discussion in this section, are “buddy” beams, which are seen in nearly every per-detector composite beam map.

Buddy beams are a particular form of ghost beam that results from a reflection off a 50 K filter, which then travels back to the focal plane, reflects again, and travels through the optics, emerging out of the window at a position on the sky that is exactly opposite the boresight from the main detector. This is illustrated in a simple diagram in Fig. 5.9, and a sample of one in a real map can be seen in Fig. 4.13. Since the reflection off the 50 K filter happens after traveling through both lenses, the buddy beam on sky is compact, with a shape very similar to the main beam. Buddy beams generally do not lead to significant $T \rightarrow P$ leakage, since these reflections happen exactly the same way for an A detector as it does for a B detector. However, it is still important to understand the total beam power residing in these buddy beams, because they can inform us about the performance and limitations of the optical elements, particularly the AR coating, in currently deployed telescopes.




Figure 5.9: Simple schematic demonstrating how buddy beams form. In reverse time, light coming from the main detector partially reflects off a 50 K filter, traveling back to the focal plane, which then travels back through the 4 K optics and emerges on sky at a location exactly opposite the boresight from the main detector beam. Because the reflection happens after both 4 K lenses, the buddy beam is compact with a similar beam width to the main beam. The optics are not to scale, and other filters have been omitted for clarity.

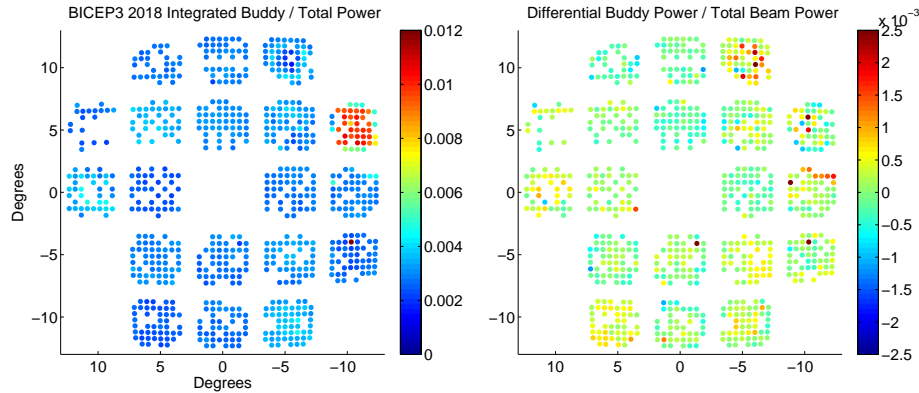


Figure 5.10: BICEP3 integrated buddy beam power as a fraction of total beam power. *Left:* average buddy beam power in each detector pair. *Right:* pair-difference buddy beam power. Tile 11 (center) is omitted due to the inability to isolate the buddy beam from the main beam. Tile 1 (top right) has abnormally high power in the buddy beams due to the lack of a module in the focal plane at the location opposite the boresight. Instead, there is reflective Mylar film which reflects more of the beam back to sky. In general, the buddy beams are not very polarized, with most detectors leaking $< 0.1\%$ of the total beam power to polarization.

An FPU plot of the integrated buddy power as a ratio of the total beam power for BICEP3 is shown in Fig. 5.10. The left plot shows the average buddy power within each pair, and the right plot shows the pair-difference buddy power. While most detectors show a uniform $\sim 0.3 - 0.5\%$ integrated buddy power, Tile 1 is an exception with 1% buddy power. This is because there is no module in the focal plane at the location opposite the boresight of Tile 1. Instead, there is bare copper covered in reflective Mylar polyester film, which accentuates the reflected power that goes back on sky. The right-hand plot shows that the buddy beams are not very polarized, leaking $< 0.1\%$ of the total beam power to polarization. Although the results are not plotted here for *Keck*, the buddy beam power is consistent and similar to that of BICEP3 (excluding Tile 1).

5.6 CROSSTALK BEAMS

Another feature commonly seen in beam maps is not an artifact of the optical chain, but instead stems from detector readout. As described in Section 2.2, BICEP/*Keck Array* uses a SQUID-based time-domain multiplexed readout to sample the detector signal. Each detector inductively couples to a SQ1, and 33 SQ1 are read sequentially within one SQ2 column. One form of crosstalk apparent in most detectors is inductive crosstalk, where each SQ1 unintentionally inductively couples to the nearest neighboring SQ1 in each direction within a column. This results in a small ($0.2 - 0.3\%$) crosstalk from the nearest “upstream” and “downstream” neighbors in each detector, which is seen in both the far-field beam maps and in an independent analysis of cosmic ray hits. Another form of crosstalk, which seems to only exist in *Keck* receivers (and mostly confined to one receiver), exists in the downstream direction only. “Settling time” or “row-switching” crosstalk occurs when the readout switches from SQ1 of detector N to the SQ1 of detector $N + 1$ before the previous signal has been allowed to completely decay, and varies in amplitude. We will revisit this mechanism momentarily.

Only a small subset of detectors show row-switching crosstalk. One example is *Keck* 2013-2016 rx4, which is a 150 GHz receiver. Fig. 5.11 shows the downstream crosstalk levels seen in all B detectors in all five *Keck* receivers in 2015 beam maps. Most detectors show levels consistent with inductive crosstalk ($\sim 0.2 - 0.3\%$), except




Figure 5.11: Downstream crosstalk beam power relative to total beam power in *Keck* 2015 *B* detectors. Black dashed lines separate the five different receivers. Nearly all detectors in the first four receivers show crosstalk power that is consistent with inductive crosstalk (which happens both upstream and downstream). Rx4, the right-most receiver plotted, shows elevated downstream crosstalk levels as high as $\sim 2 - 3\%$ of the total beam power.

for rx4 which has elevated crosstalk as high as $\sim 2 - 3\%$ of the total beam power.

This elevated row-switching crosstalk persisted even after swapping the 150 GHz FPU in rx4 with a 270 GHz FPU prior to the 2017 season². This indicates that the source of the long time constant leading to this crosstalk could be in the SQUID Series Arrays (SSA), which are at the 4 K stage and are not changed during a focal plane swap. It could also indicate a problem in the tuning of the SQ2, if a similar tuning was used between the two FPUs.

Another feature of the row-switching crosstalk seen in *Keck* rx4 is a difference in

²In 2017, as a 270 GHz receiver, the crosstalk dipped to $\sim 1\%$, which is still quite elevated. It remained at this level after another FPU swap (to a different 270 GHz FPU) in 2018. This receiver is now observing the CMB as rx3 in BICEP Array, still with crosstalk levels of $\sim 1\%$ seen in FFBM.




Figure 5.12: Component signal beam map timestream (left) and corresponding phase beam map timestream (right) for a sample detector in *Keck* rx4 (150 GHz) taken in 2015. The row-switching crosstalk beam is the feature above the main beam, at roughly $(x', y') = (1.5, 0)$. The difference in phase between the crosstalk beam and the main beam is roughly 90° , for a chop rate of 16 Hz.

phase between the crosstalk beam and the main beam. This can be seen in “phase” beam maps, where the value of $\arctan(\text{sine component} / \text{cosine component})$ is plotted instead of the usual cosine component. An example of this phase difference is shown in Fig. 5.12. The phase beam map plotted on the right indicates a phase difference of roughly 90° between the main beam and crosstalk beam for a chop rate of 16 Hz. This initially seems impossible, because the row-switching rate from one SQ1 to the next is orders of magnitude faster than 16 Hz.

Although the mechanism of this phase delay is still under investigation, the current hypothesis is that a frequency dependence of the observed crosstalk arises from the finite servo bandwidth of the SQ1s. In a simplified sense, the output current from the TES detector is inductively coupled to the SQ1, whose voltage changes when the TES detector current changes from an incoming photon. A second current source is

coupled to the SQ1, which applies a feedback flux via the servo loop to maintain the SQ1 voltage at the chosen value. This feedback flux is what we use as the detector signal. With row-switching crosstalk, some of this output feedback signal will carry from row N to row $N+1$ as discussed earlier. The transfer function G from the SQ1 in row N to the SQ1 in row $N+1$ in this simplified model is:

$$G = A \frac{2 \pi i f \tau}{1 + 2 \pi i f \tau} \quad (5.3)$$

where A is the suppression from the exponential settling decay, f is the frequency of the signal (in this case the FFBM thermal chop rate, 16 Hz), and τ is the inverse of the servo bandwidth (3 ms for *Keck*³). For low frequencies this simplifies to $G \simeq 2 \pi i f \tau$, which has a 90° phase shift with an amplitude that depends on the source frequency f . This crosstalk amplitude would be maximized in the cosmic ray analysis (high frequency pulses) and would be significantly large in the FFBM analysis (16 Hz) compared to the low level of crosstalk expected from the low-frequency CMB ($\sim 0.1 - 1$ Hz). This difference in frequency makes it difficult to predict the amount of crosstalk expected in the CMB, given the amount of crosstalk seen in the beam maps. To explicitly calculate the expected CMB crosstalk, we need exact measurements of the servo bandwidth of the detectors as they were installed in *Keck*. Measurements

³The 3 ms servo bandwidth is quite slow compared to BICEP2 and BICEP3, and was set this way because of increased stability observed in this generation of SQUIDs with slower bandwidth. Other instruments with faster bandwidths do not see this form of crosstalk, which reinforces the hypothesis discussed in this section.

of these bandwidths are currently underway to test this hypothesis and predict the downstream, out-of-phase crosstalk seen in CMB data.

5.7 RESIDUAL BEAMS AND BORESIGHT ROTATION

Perhaps the most important quality of a pair of detector beams is the power remaining after *deprojecting* the leading order difference modes. To second order, the modes of a differential elliptical Gaussian couple to linear combinations of the CMB T sky and its first and second derivatives. Since the beam shapes are constant in time, we can construct leakage templates corresponding to these leading difference modes, sample them using each detector pair's real trajectory data, regress each detector pair's signal timestream against its leakage template, then subtract the fitted template from the signal. This analysis technique is called deprojection, and is fully discussed in Section 6.1.2. The power that remains after deprojection is called the *undeprojected residual*, or simply the residual beam. The residual beams are of the utmost importance, as they cause $T \rightarrow P$ leakage that is not removed in standard analysis. In this section, we briefly analyze the band-averaged difference beams and residual beams to look for patterns that may be correlated across many detectors. We also examine how this residual beam power may cancel under various boresight angle combinations.

Fig. 5.13 show the average V , H , difference, and residual beams for each of the frequency bands in BICEP/Keck Array. Starting with the difference beams, it is immedi-

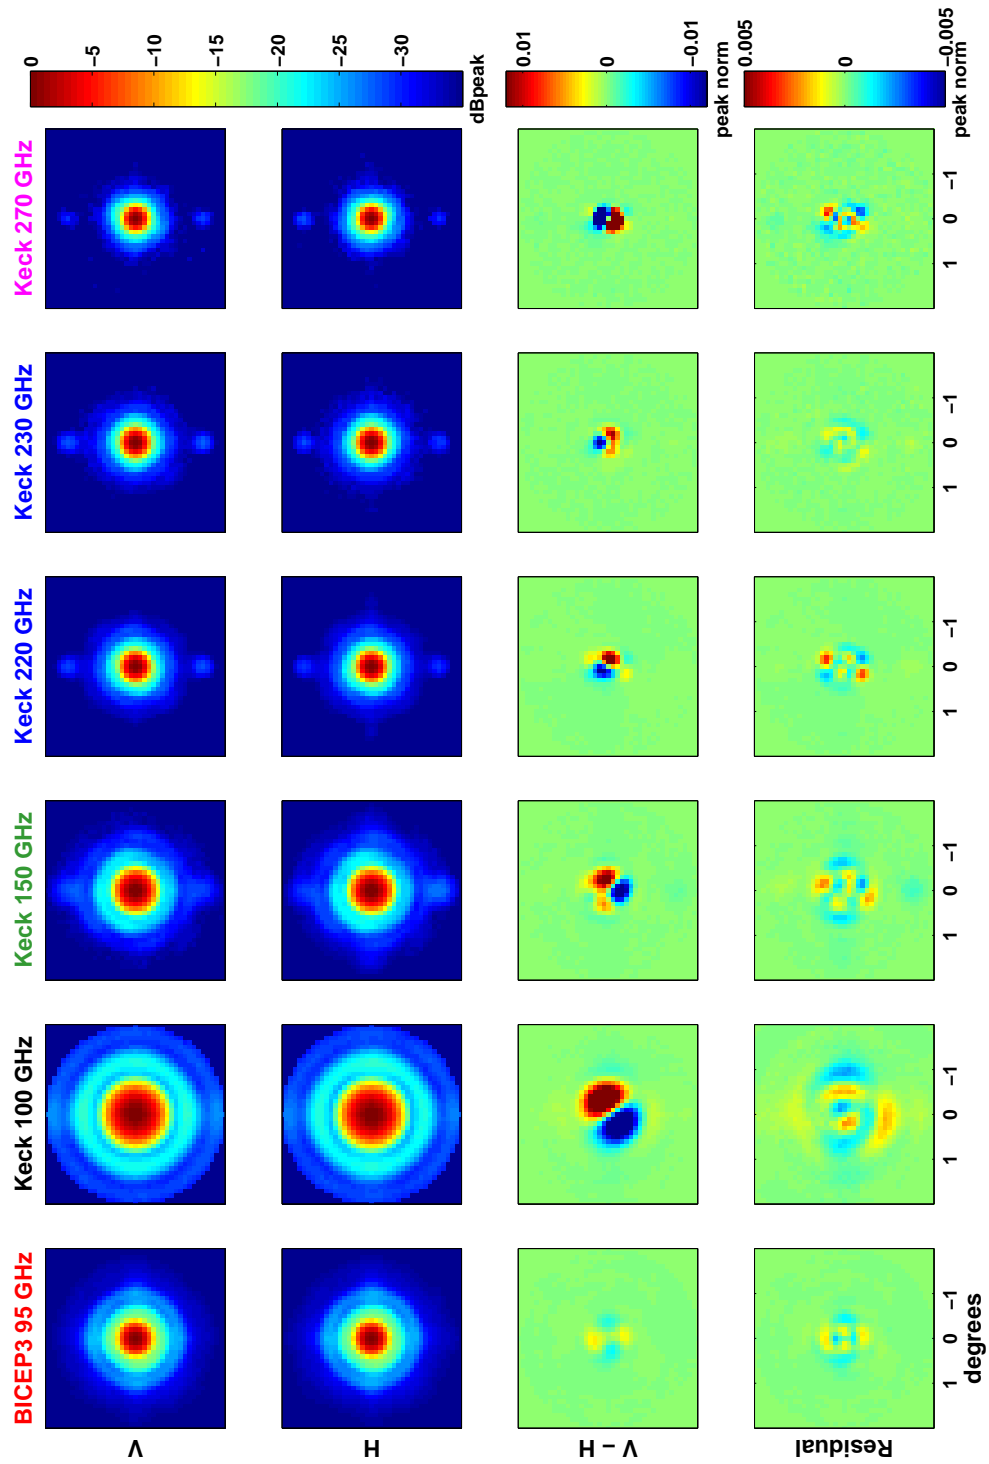


Figure 5.13: Band-averaged beams for each BICEP3/Keck/Array frequency band. The alternative polarization axis system (V, H) is used here instead of (A, B) – this change rotates 3/4 of the BICEP3 tiles so that all polarization axes are aligned. The first and second rows show the average V and H beams; the third row shows the average difference beam; the fourth row shows the average residual beam.

ately clear that the dominant difference mode is differential pointing⁴. BICEP3, on average, has remarkably lower amplitude difference beams compared to all other bands, as can also be seen in the Gaussian fit parameters (e.g. Table 5.3).

Looking at the residual beams, there are many interesting patterns seen in most bands. *Keck* 95 GHz has bright, relatively large patterns roughly a degree away from the beam center. This is a known effect from anomalous coupling between detectors near the edge of a tile and the tile corrugations. This was reduced in subsequent detector fabrication, as this pattern is not clearly present in the *Keck* 220, 230, and 270 GHz average residual beam. The BICEP3 residual beam may show a similar effect on a smaller scale (expected for a smaller beam) and at smaller amplitude. As discussed in Section 6.4.1, this feature, which aligns with the orientation of the polarization axes, is a suspected source of E leakage power at $\ell \sim 150 - 200$.

The *Keck* 220, 230, and 270 GHz residual beams show unique features that appear as well-defined dipoles with a Gaussian shape (most clear at 220 GHz). These are attributed to direct-island coupling – photons directly illuminating the TES detectors. The A and B detectors are offset from each other (forming a dipole), and they are located just below the antenna patch (causing the dipole to be offset from the center). There are two dipoles in this average because we are averaging over four tiles per *Keck* receiver – with two rotated 180° with respect to the other two. Since this

⁴Differential gain is generally quite large as well, but this is not captured in beam map measurements.

pattern aligns with a \times pattern (45° with respect to the polarization axes) in this coordinate system, we expect the $T \rightarrow P$ leakage to preferentially leak to B .

We expect natural cancellation of some components of the residual beam, either by the same detector pair hitting the same spot on the map at different boresight rotation angles, or by coadding over different detector pairs if the residual beam component is common-mode across the FPU. In general, the nature of these cancellations is very complex, and depends on the amount of weight each detector pair contributes at each boresight angle in each map pixel – for a full description and analysis with timestream simulations, see [Sheehy 2013](#) and BK III ([The BICEP2 Collaboration 2015](#)). Here, we consider a simplified analysis of this cancellation by examining how the band average residual beams shown in Fig. 5.13 average down when coadding over multiple boresight rotation angles.

Fig. 5.14 shows how the averaged residual beams for BICEP3 and *Keck* 150 and 230 GHz partially cancel out when coadding over various combinations of boresight rotation angles. When coadding over two angles separated by 180° , any “dipole-like” components will cancel. The average power remaining after this cancellation is plotted in the second row. When coadding over four angles separated by 90° , any “monopole-like” components will also cancel (because $Q \rightarrow -Q$ under 90° rotation). This is plotted in the third row. The bottom row of the figure shows a hypothetical beam after achieving the maximum possible cancellation by coadding over infinitely many boresight angles between 0° and 360° . For *Keck* (and now BICEP Array), we observe at




Figure 5.14: Band-averaged residual beams (first row), when coadded over two boresight rotation angles separated by 180° (second row), four angles separated by 90° (third row), and infinitely many angles (fourth row). The first column shows BICEP3, the second shows *Keck* 150 GHz, and the third shows *Keck* 230 GHz. Based on Fig. 5.13, the cancellation shown here for BICEP3 should be similar to the expected cancellation for *Keck* 95 GHz, and *Keck* 230 GHz should be analogous to *Keck* 220 and 270 GHz. All plots use the same color scale.

eight boresight angles each separated by 45° allowing cancellation of both monopole- and dipole-like terms. However for BICEP3, due to the geometry of the pulse tube cryocooler preventing full 360° rotation of the telescope within the mount, we only observe at $0^\circ, 45^\circ, 180^\circ$, and 225° . In this case the monopole-like terms do not cancel. Fig. 5.14 shows that, on average, we do not expect a heavy penalty from this lack of cancellation (the second row is very similar to the third row for BICEP3). This was confirmed upon running the $T \rightarrow P$ leakage simulations presented in Chapter 6.

*Everything was so new – the whole idea of going
into space was new and daring. There were no
textbooks, so we had to write them.*

Katherine Johnson

6

BK18 Temperature-to-Polarization

Leakage Simulations

PERHAPS THE MOST POWERFUL UTILITY of the library of composite beam maps is their ability to predict the temperature-to-polarization ($T \rightarrow P$) leakage in a real

CMB data set. The *BICEP/Keck Array* telescopes measure polarization by taking the difference in the signal between each co-located, orthogonally polarized pair of detectors (known as “pair differencing”). $T \rightarrow P$ leakage arises from any mismatch in the beams between two detectors in a pair:

$$d_{T \rightarrow P} = T(\vec{x}) * (B_A(\vec{x}) - B_B(\vec{x})) = T(\vec{x}) * B_\delta(\vec{x}) \quad (6.1)$$

where $T(\vec{x})$ is the temperature at some point \vec{x} on the sky, B_A (B_B) is the beam of the A (B) detector, and $*$ denotes a convolution. Because the CMB T anisotropies ($\sim 100 \mu\text{K}$) are much brighter than the polarized anisotropies ($\sim 1 \mu\text{K}$), any significant beam mismatch B_δ will noticeably bias the overall polarization measurement if left unchecked.

We use an analysis method, “deprojection,” to remove the leading order modes of the difference beam B_δ (see Section 6.1.2). The remaining residual beam, the “unde-projected residuals,” cause $T \rightarrow P$ leakage that is not removed by deprojection. We therefore require a method of quantifying the leakage from the residual beams in the real data. To do this, we run specialized “ $T \rightarrow P$ leakage simulations” (also called “beam map simulations” in previous publications), where each detector observes a simulated temperature (no Q/U) sky that is smoothed by its own composite beam map. Since there is no polarization in the input map, any polarization signal output from these simulations must be entirely due to the beam mismatch within each detector

pair. Using these simulations with the BK15 data set, the resulting bias on r due to $T \rightarrow P$ leakage is estimated to be $\Delta r = 0.0027 \pm 0.0019$ (The BICEP2 and Keck Array Collaborations 2019). In this chapter, we walk through the $T \rightarrow P$ leakage simulation pipeline, and show the results for the BK18 data set, which includes all observations through the 2018 observing season.

6.1 T-TO-P LEAKAGE SIMULATION PIPELINE

The $T \rightarrow P$ leakage simulations, like any other standard CMB simulation, are designed to mirror the structure and analysis of the real data as much as possible. The BICEP/Keck Array analysis pipeline utilizes a technique similar to the MASTER formalism (Hivon et al. 2002) – a computationally efficient Monte Carlo approach where an ensemble of simulated skies are observed and analyzed, with each realization drawn from a given astrophysical model. This has the advantage of including the non-trivial effects of varying data cuts, weights, and limited sky fraction in the simulations in the exact same way as they impact the real data. In this section we walk through each step of the analysis pipeline, specifically noting where $T \rightarrow P$ simulations differ from the mainline analysis methods.

6.1.1 INPUT MAPS AND GENERATING TIMESTREAMS

First, a series of pre-generated maps containing the desired astrophysical signal must be provided. In general, there are four flavors of signal maps used in primary *Br-CEP/Keck Array* simulations:

- Standard Λ CDM cosmology, with no gravitational lensing
- Lensed- Λ CDM
- $r = 0.2$ tensors (B modes only)
- Gaussian polarized dust (no T)

Each of these are in `Healpix` format with $N_{\text{side}} = 2048$ (Gorski et al. 2005). In standard simulations, each of these are smoothed by the array-averaged $B(\ell)$ for each frequency band used in the observations (such as those plotted in Fig. 5.6). In the case of Gaussian polarized dust, the amplitudes are scaled from 353 GHz (the assumed frequency used to generate the dust map realizations) to the integrated spectral response of the array-averaged bandpass (see Fig. 3.4). For the $T \rightarrow P$ leakage simulations, only the *Planck* PR1 T map at the nearest available frequency band is needed. No smoothing is applied to the T map, since the convolution is done using the per-detector composite beam maps when sampling.

Noise in the standard simulations is not generated through pre-constructed maps. Instead, we use a “sign-flip” technique using the real data itself. A noise realization is formed by randomly flipping the sign of half (by weight) of the scansets in a data

set, causing all common signal to cancel and leaving only noise. In the $T \rightarrow P$ leakage simulations, no noise is injected – the only noise contribution comes from the noise in the beam map measurement.

Detector pointing timestreams are taken from the real data and used to sample the input signal maps. The sampling for standard simulations is done via a second-order Taylor interpolation. For the $T \rightarrow P$ leakage simulations, the signal T map must first be projected onto a flat sky¹, then convolved with the flat-sky composite beam map of that detector, before sampling using linear interpolation. The composite beam map is a circular “thumbnail” of a chosen radius, usually between $2^\circ < r < 8^\circ$.

Simultaneously with the sampling of the signal maps, each detector also samples the deprojection templates used to filter out the leading order main beam difference modes. This is explored in full detail in Section 6.1.2.

At this point, the timestreams are pair-summed and pair-differenced, and filtering steps are applied. First, a third-order polynomial fit is subtracted from each pair-sum and pair-diff halfscan, which drastically reduces the atmospheric $1/f$ noise contribution. Next we do “ground subtraction,” where two scan-synchronous templates (one for each scan direction) are constructed and subtracted from the data. These templates are formed by taking the average of all samples at a constant azimuth – this effectively captures ground-fixed signal and not sky signal, due to the sky rotation

¹Care is taken such that the pixels are square along the scan direction of that scanset, to minimize distortion effects.

during a 50-minute scanset. Both the poly-filter and the ground subtraction preferentially reduces our sensitivity to low- ℓ modes along the scan direction. This step is identical between the real data, standard simulations, and $T \rightarrow P$ leakage simulations.

In the standard simulations and $T \rightarrow P$ leakage simulations, the per-pair, per-scanset weights are taken from the real data. For real data the weights are the inverse variance, taken over each scanset after the filtering is applied. Data cuts are also taken from the real data. We use a large suite of cuts, a small sample of which are listed below:

- Poor stability or goodness-of-fit of the elevation nods used for relative gain calibration.
- Abnormally large amounts of spikes and steps in the raw data, such as those due to cosmic ray hits or flux jumps in the SQUID readout.
- Large temperature or temperature instability at the focal plane.
- Abrupt changes in the scan-synchronous ground subtraction template, possibly arising from radio frequency interference due to local transmitters at the South Pole Station.

This is not a complete list; for more details on the cuts, see [Willmert 2019](#).

6.1.2 DEPROJECTION

In Eq. 6.1 we showed that the signal in a given detector pair from $T \rightarrow P$ leakage is equal to the convolution of the difference beam B_δ with the temperature sky at any given point. Here we discuss a timestream-based filtering method, *deprojection*, that

removes a majority of that leakage signal by constructing a set of template maps corresponding to each of the leading modes in a Taylor expansion of B_δ . The templates can then be observed with real detector pointings, and regressed against the real data to quantify the amount of leakage signal corresponding to each mode.

In Section 4.2.3, we parameterized the beam $B(x)$ of a single detector as a 2D elliptical Gaussian with beamwidth σ and plus and cross ellipticity p and c (Eq. 4.1). An ideal pair of beams would be circular ($p = c = 0$), have zero differential pointing ($\delta x = \delta y = 0$), and equal gain ($g_A = g_B = 0$). By perturbing one of the differential parameters δk , expanding in a Taylor series, and keeping only the leading terms of the expansion, we arrive at a set of analytic expressions for each of the six difference modes $B_{\delta k}$:

$$\begin{aligned}
B_{\delta g} &\simeq \delta g B(\vec{x}) \\
B_{\delta x} &\simeq \delta x \nabla_x B(\vec{x}) \\
B_{\delta y} &\simeq \delta y \nabla_y B(\vec{x}) \\
B_{\delta \sigma} &\simeq \sigma \delta \sigma (\nabla_x^2 + \nabla_y^2) B(\vec{x}) \\
B_{\delta p} &\simeq \frac{\sigma^2}{2} \delta p (\nabla_x^2 - \nabla_y^2) B(\vec{x}) \\
B_{\delta c} &\simeq \frac{\sigma^2}{2} \delta c (2 \nabla_x \nabla_y) B(\vec{x})
\end{aligned} \tag{6.2}$$

From these equations we can see which modes are sensitive to the zeroth derivative (δg), the first derivative ($\delta x, \delta y$), and the second derivative ($\delta \sigma, \delta p, \delta c$). For a more explicit calculation of the above expansions, see BK III ([The BICEP2 Collaboration](#)




Figure 6.1: The six differential beam modes corresponding to the expansion of B_δ shown in Eq. 6.2. Each of these correspond to a different deprojection template that is sampled and regressed against the real data, in order to remove the leading order $T \rightarrow P$ contamination. Figure from BK III (The BICEP2 Collaboration 2015).

2015) and Aikin 2013. Each of these differential modes are shown in Fig. 6.1. To relate this beam expansion to the temperature sky and ultimately the leakage signal to be deprojected, we revisit Eq. 6.1. Let's consider the leakage signal from differential beamwidth as an example:

$$\begin{aligned}
 d_{\delta\sigma} &= T(\vec{x}) * B_{\delta\sigma}(\vec{x}) \\
 &= T(\vec{x}) * \sigma \delta\sigma (\nabla_x^2 + \nabla_y^2) B(\vec{x}) \\
 &= \sigma \delta\sigma (\nabla_x^2 + \nabla_y^2) T(\vec{x}) * B(\vec{x})
 \end{aligned}
 \tag{6.3}$$

In the above equations, we substituted the expression for $B_{\delta\sigma}$ from Eq. 6.2, then

used the convolution theorem to move the derivative term from the beam component to the temperature component. We can therefore construct a deprojection template map for differential beamwidth: $(\nabla_x^2 + \nabla_y^2) T(\vec{x})$. The deprojection coefficient, the value obtained by regressing the real data against the sampled template, is equal to $\sigma \delta\sigma$. This process can be repeated for all six modes in Eq. 6.2 to generate six deprojection templates.

We pre-construct maps of the set of first and second derivatives of *Planck* T at the nearest frequency, smooth them with the appropriate $B(\ell)$ for that band, then accumulate the deprojection templates by observing these maps using the real detector pointing timestreams. These templates are then regressed against the real data to generate deprojection coefficients. We then can choose to 1) “deproject,” by removing the template scaled by the deprojection coefficient, or 2) “subtract,” by scaling the template by the coefficient taken from FFBM measurements, such as those shown in Section 5.2, before removing it from the data. The advantage of deprojection is that it guarantees a complete removal of leakage signal from that mode, but like any other filtering method, it also removes real cosmological signal that may exist in that mode. We usually elect to deproject the two most dominant modes – differential gain and differential leakage – and subtract differential ellipticity. We do not deproject differential ellipticity, because of real non-zero TE correlation that would be filtered out. We normally do not deproject or subtract differential beamwidth, since it is very small and therefore has negligible impact on the final data.

We use a shorthand notation to indicate the level of deprojection applied: “dpXXXX,” where $X = 0$ means nothing is done with that mode, $X = 1$ means it is deprojected, and $X = 2$ means it is subtracted. The ordering of the four digits is differential gain, differential pointing, differential beamwidth, and differential ellipticity. For example, the standard level of deprojection used in the final results is dp1102 – differential gain and pointing are deprojected, differential beamwidth is ignored, and differential ellipticity is subtracted.

It is important to note that, although the formalism shown above is in terms of Gaussians, the deprojection method is agnostic of the actual beam shape. Therefore it does not require external beam measurements to effectively remove $T \rightarrow P$ leakage. However, it is valuable to perform cross-checks between coefficients derived from deprojects vs. those derived from FFBM, as shown in Section 5.4.

6.1.3 BINNING AND COADDING MAPS

The next step is to bin the accumulated quantities into per-detector pair, per-scanset maps, called “pairmaps.” The grid is in R.A. and declination, with a pixel side length of 0.25° , square at a declination of -57.5° . The quantities that are accumulated are those needed for calculating Stokes Q and U parameters – the pair-difference, and the sine and cosine of the telescope boresight rotation angle, as demonstrated below – as well as the pair-sum, pair-sum/pair-difference weights, and the deprojection templates.

The response of an ideal pair of co-located, orthogonally polarized detectors at any point on the sky is:

$$\begin{aligned} A(\vec{x}) &= T(\vec{x}) + Q(\vec{x}) \cos(2\psi) + U(\vec{x}) \sin(2\psi) \\ B(\vec{x}) &= T(\vec{x}) + Q(\vec{x}) \cos(2(\psi + \pi/2)) + U(\vec{x}) \sin(2(\psi + \pi/2)), \end{aligned} \quad (6.4)$$

where ψ is the polarization angle of the A detector on sky at some boresight rotation angle. We have assumed for the moment that the gains are equivalent ($g_A = g_B = 1$), the beams are identical delta functions (“pencil beams”) for each detector, and the polarization angle of the B detector is offset from A by exactly $\pi/2$. The pair-sum S and pair-difference D results in a measurement of temperature and polarization, respectively:

$$\begin{aligned} S(\vec{x}) &\equiv \frac{A(\vec{x})+B(\vec{x})}{2} = T(\vec{x}) \\ D(\vec{x}) &\equiv \frac{A(\vec{x})-B(\vec{x})}{2} = Q(\vec{x}) \cos(2\psi) + U(\vec{x}) \sin(2\psi). \end{aligned} \quad (6.5)$$

The relationship between the pair-difference D and Stokes Q and U can be expressed as a matrix equation:

$$\begin{bmatrix} D \cos(2\psi) \\ D \sin(2\psi) \end{bmatrix} = \begin{bmatrix} \cos^2(2\psi) & \cos(2\psi) \sin(2\psi) \\ \cos(2\psi) \sin(2\psi) & \sin^2(2\psi) \end{bmatrix} \begin{bmatrix} Q \\ U \end{bmatrix}. \quad (6.6)$$

To extract Q and U values from the pair-difference observations, we need to accumu-

late the quantities D , $\cos(2\psi)$, and $\sin(2\psi)$. By inverting the orientation matrix on the right-hand side, we can obtain Q and U at any given map pixel. Note that the orientation matrix is singular for an observation at a single boresight rotation angle ψ . In principle this could be avoided by observing with detectors with varying polarization axes, or by utilizing natural sky rotation to observe the same sky patch at different angles. For *BICEP/Keck Array*, all focal planes have a fixed polarization angle for all detectors (up to a factor of 90°), and our patch as observed from the South Pole does not change orientation as the sky rotates. However, the freedom of boresight rotation allows us to accumulate observations at multiple angles ψ necessary to invert the matrix.

The pairmaps are saved to disk, then can be accumulated over any desired timescale, varying from one phase (a collection of ten elevation steps at a single boresight rotation angle, see Section 2.6), to multiple years of observations. It is important to note that there is freedom to choose the timescale at which deprojection is applied. The regression of the templates against the real data can in principle be done once every scanset, every phase, or on even longer time scales such as a full year. A shorter time scale allows us to capture time-varying systematics, whereas a longer time scale reduces the filtering of real signal, since more map pixels are being sampled. We find that deprojecting every phase to be a suitable compromise between these two effects. After deprojecting at every phase, the data is accumulated over all seasons to give full coadded maps in each observation band. At this step, the orientation matrix is

inverted to obtain Q and U maps, with the weights and cuts from each detector properly accounted.

6.1.4 ANGULAR POWER SPECTRA

Before transforming the Q/U maps into E/B , we apply an apodization mask to ensure each map smoothly rolls to zero. There are a couple motivations for this. First, the map pixels near the very edge are naturally noisier – due to the large field of view compared to the size of the sky patch we observe, there are more hits in the center of the map compared to the edge. An apodization mask downweights these noisy map pixels. The other benefit is to remove any discontinuities that would cause excessive ringing in the Fourier domain. We use one common apodization mask evaluated as the geometric mean of the smoothed inverse variance maps over all frequency bands. With the inclusion of BICEP3 in the BK18 data set, we have chosen a baseline analysis where BICEP3 uses an apodization mask from its own inverse variance map, which is larger than all other BICEP/*Keck Array* bands due to its larger field of view².

The simplest method to transform Stokes Q/U into the more useful E/B basis is to apply a rotation in Fourier space. Using $\tilde{Q}(\vec{\ell})$ to denote the Fourier transform of Q :

$$\begin{aligned} E(\vec{\ell}) &= \tilde{Q}(\vec{\ell}) \cos(2\phi_\ell) + \tilde{U}(\vec{\ell}) \sin(2\phi_\ell) \\ B(\vec{\ell}) &= -\tilde{Q}(\vec{\ell}) \sin(2\phi_\ell) + \tilde{U}(\vec{\ell}) \cos(2\phi_\ell) \end{aligned} \tag{6.7}$$

²In principle it is better to use a single common mask for all bands, to minimize loss of correlation in the auto- and cross-spectra. In practice this effect is negligible in this data set.

where $\phi_\ell = \arctan(\ell_y/\ell_x)$. Since this is a non-local transformation, any filtering of real signal (such as partial sky coverage or timestream filtering) will cause a mixing between E and B . Because of the relative brightness of the E modes compared to primordial B modes, this level of mixing is not acceptable. Other methods exist to mitigate the level of $E \rightarrow B$ mixing (Smith & Zaldarriaga 2007) due to partial sky coverage, but the remaining variance is still too large for the precision needed in BICEP/Keck Array measurements.

We have developed a matrix-based map purification step that is applied to the Q/U maps to isolate true B modes on the sky. This numerical approach, which follows the eigenvalue decomposition method in Bunn et al. 2003, utilizes a “matrix purification operator” that can be performed on any set of Q/U maps which identifies B modes that are orthogonal to both E modes and ambiguous modes. This operation includes the effect of both partial sky coverage and all real timestream filtering applied to a data set, reducing the $E \rightarrow B$ leakage to levels below $r = 10^{-4}$ (see BK VII, The BICEP2 and Keck Array Collaborations 2016a). While this is critical for real data and standard simulations, it is less important for $T \rightarrow P$ leakage simulations which generally have an equal amount of leakage E and B modes. However, we still apply this matrix purification step to ensure that the same B modes kept/rejected in the real data are treated identically in the $T \rightarrow P$ leakage simulations.

The next step is to correct for noise and the effect of suppression from filtering, the beam, and matrix purification. The real data and standard simulations apply noise




Figure 6.2: $B \rightarrow B$ bandpower window functions (BPWF) and total suppression factor for BICEP2. The suppression from the map window function, the filtering, and the beam are individually plotted. Filtering dominates the suppression at low ℓ , whereas the beam drives the suppression at high ℓ . Figure from BK VII ([The BICEP2 and Keck Array Collaborations 2016a](#)).

debiasing by subtracting the mean of 499 sign-flip noise realizations – for $T \rightarrow P$ leakage simulations, no noise is injected, so this step is skipped (there is a noise debiasing done to the auto spectra due to noise in the beam maps, as discussed in Section 6.4.1).

To calculate the bin-by-bin suppression from filtering, the beam, and purification, we calculate “band-power window functions” (BPWF) which contain the transfer functions that describe how a given bandpower sources power from a range of multipoles ℓ . These are calculated using an ensemble of simulations where the input sky has power at only a single ℓ . These are calculated for every combination of auto- and cross-spectra in a BICEP/Keck Array data set, then divided from the measured band powers. The BB suppression for BICEP2 is shown in Fig. 6.2.

6.2 MULTICOMPONENT LIKELIHOOD ANALYSIS

6.2.1 BAYESIAN APPROACH

Using the formalism described so far, we generate a series of auto- and cross-frequency power spectra using real BICEP/*Keck Array* maps, as well as external maps from *Planck* and WMAP. The same can be applied to maps that best estimate the noise in the data (using the sign-flip technique described in Section 6.1.1), and simulated maps corresponding to parametric models of Λ CDM cosmology (with or without gravitational lensing or tensor modes) and galactic foregrounds. Given this wealth of information in the form of power spectra, we wish to find the values of the cosmological model parameters that best describe the data.

In Bayesian inference, the probability of the model given the data (the posterior) is proportional to the likelihood of the data given the model and any prior information about that model. There are a number of methods to evaluate a general likelihood – one could calculate it analytically if it has a closed form, or computationally sample every point in the multidimensional parameter space. In our case, the likelihood is generally non-Gaussian due to partial sky coverage, filtering, and the presence of both correlated and uncorrelated noise, making a closed analytic solution difficult. The full computational approach is also infeasible due to the size of the parameter space. We therefore use the Hamimeche-Lewis (HL) approximation ([Hamimeche & Lewis 2008](#)),

which takes the following form:

$$-2 \log \mathcal{L} = \mathbf{X}^T \mathcal{M}_f^{-1} \mathbf{X} \quad (6.8)$$

where \mathcal{L} is the likelihood function, \mathcal{M}_f is a fiducial model bandpower covariance matrix (BPCM), and $\mathbf{X} = \{\mathbf{X}_\ell\}$ is a vector of modified bandpowers, where the modification to each bandpower is:

$$\mathbf{X}_\ell = \text{vecp}[\mathbf{C}_{\ell f}^{1/2} g(\mathbf{C}_\ell^{-1/2} \hat{\mathbf{C}}_\ell \mathbf{C}_\ell^{-1/2}) \mathbf{C}_{\ell f}^{1/2}] \quad (6.9)$$

Here the fiducial model bandpowers are $\mathbf{C}_{\ell f}$, the theory bandpower expectation values are \mathbf{C}_ℓ , and the real data bandpowers are $\hat{\mathbf{C}}_\ell$. The bandpower matrices are symmetric, with the auto-frequency bandpowers on the diagonal and the cross-frequency off-diagonal. The vecp operation converts the symmetric matrix inside the brackets into a vector, keeping only unique entries. Finally, g is a matrix operation that finds that values of the eigenvalues of the matrix, and applies $g(x) = \text{sign}(x-1)\sqrt{2(x - \ln(x) - 1)}$.

We evaluate the resulting likelihood for a data set using two different methods.

The first is a maximum likelihood (ML) search, which gives the single point that maximizes the likelihood function. The second method uses Markov Chain Monte Carlos (MCMC) to intelligently sample the posterior distribution. The software used to apply the MCMC technique is `CosmoMC` (Lewis & Bridle 2002).

6.2.2 PARAMETRIC MODEL

The multicomponent model used in our analysis framework includes four independent signal types: a cosmological component (Λ CDM + tensor modes), uncorrelated thermal dust, uncorrelated synchrotron emission, and a spatially correlated dust/synchrotron component. We also consider spectral decorrelation between dust and synchrotron emission, which manifests as a suppression in the cross-spectral response between the two components. Since these components are independent, the total model can be evaluated as the total sum of all individual terms.

The following are the parameters used to describe the component terms described above:

- A_d, A_s – amplitudes of the thermal dust and synchrotron emission at $\ell = 80$, with a pivot frequency of 353 GHz and 23 GHz, respectively,
- α_d, α_s – spatial spectral indices for dust and synchrotron,
- β_d, β_s – frequency spectral indices for dust and synchrotron,
- ϵ – dust-synchrotron spatial correlation,
- r – tensor-to-scalar ratio at pivot scale 0.05 Mpc^{-1} , with spectral tilt $n_t = 0$,
- A_L – scaling of the BB lensing amplitude, which is fixed to $A_L = 1$ in most baseline analyses, but can also be a free parameter,
- Δ'_d, Δ'_s – dust and synchrotron decorrelation between two frequency bands.

Physically-motivated priors are used for parameters where applicable, otherwise a generous uniform prior is generally used. Alternative analyses are usually employed in

parallel to a baseline scenario, where different parameters can be held fixed or allowed to vary, other prior distributions can be explored, and additional nuisance parameters corresponding to a known systematic are implemented and marginalized over.

When quantifying the bias on r due to $T \rightarrow P$ leakage as predicted by the leakage simulations, there are a number of different approaches. Similar to the approach described for bandpass uncertainties in Section 3.3.3, one can 1) inject a template leakage consistent with the simulation predictions into the power spectra, then attempt to fit and remove it using nuisance parameters; 2) inject leakage templates, then ignore it in the likelihood; 3) do not inject leakage, and attempt to fit and remove any leakage using nuisance parameters. Each of these approaches are considered and discussed in greater detail in Section 6.4.3.

6.3 BK18 BASELINE RESULTS

Here we present the results of the primary analysis of all BICEP/*Keck Array* data through the 2018 observing season (BK18). In total, this data set comprises 150 GHz data from BICEP2, multifrequency data (95, 150, 220 GHz) data from *Keck Array*, and 95 GHz data from BICEP3³. The analysis techniques used to generate simulated timestreams, bin into maps, and generate power spectra are described in detail in Section 6.1, which closely follows what was done in previous analyses such as those used

³At the maps \rightarrow power spectrum stage, the *Keck* 220 and 230 GHz bands are coadded into a single band, denoted as 220 GHz throughout this chapter.

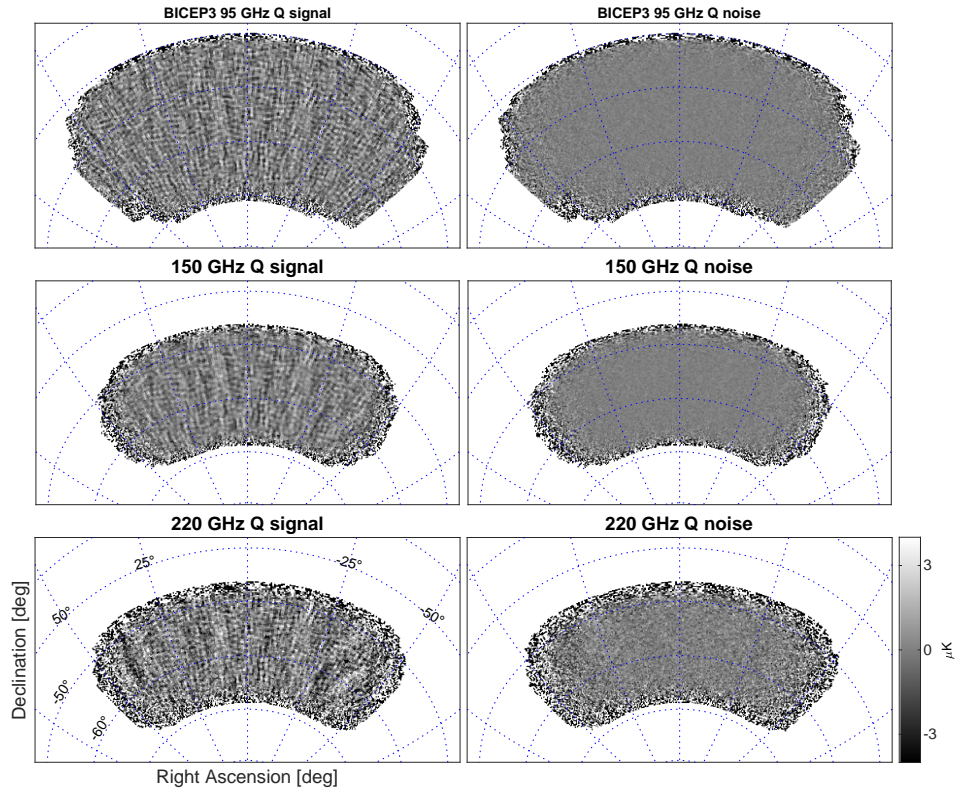


Figure 6.3: BK18 Q signal (left) and noise (right) maps, using data from BICEP3 at 95 GHz (top), BICEP2 + *Keck* at 150 GHz (middle), and *Keck* 220 GHz (bottom) observations. All plots use the same color scale. Note the larger field of BICEP3, which also has a noise level comparable to 150 GHz. The brightness of the E modes is clear in all signal maps, and some polarized dust emission is visible in the 220 GHz signal map between RA -50° and -25° .

in BKP, BK14, and BK15. For a full discussion on the results and interpretation, see BK18 (*The BICEP2 and Keck Array Collaborations 2021*). Note that due to the timing of the release of this dissertation and of BK18, results shown here may slightly vary compared to the final results in the BK18 publication.

Fig. 6.3 shows the Q signal and noise maps for all frequency bands in BK18. These are dp1102 maps – differential gain and pointing are deprojected and differential el-

lipticity is subtracted. Note BICEP3 has a larger field of view, leading to a larger effective sky area compared to *Keck* ($\approx 600^\circ$ vs 400°). E modes dominate the signal in all maps, though some polarized dust emission can be seen in the 220 GHz Q map between -50° and -25° . The map depth reached in 95, 150, and 220 GHz is 2.8, 2.8, and $8.8 \mu\text{K}_{\text{CMB-arcmin}}$, respectively. In three years of observations, BICEP3 has already reached the same map depth as 18 receiver-years at 150 GHz. The improvement at 220 GHz is also clear, comparing to the BK15 map depth of $26 \mu\text{K}_{\text{CMB-arcmin}}$.

After rotating from Q/U to E/B in harmonic space, we can apply an inverse Fourier transform to generate real-space maps of E and B , which are shown in Fig. 6.4. The matrix purification described in Section 6.1.4 has been applied. They are filtered to degree angular scales ($50 < \ell < 200$). ΛCDM signal dominates the E modes at all frequencies. For the B modes, at 95 GHz there is roughly equal contribution from lensed- ΛCDM and noise, and at 150 and 220 GHz it is dominated by thermal dust.

Along with the BICEP/*Keck Array* frequency bands, we also use data at 23 and 33 GHz from WMAP (Bennett et al. 2013), and *Planck* maps made using NPIPE (The *Planck Collaboration* 2020e). Auto- and cross-frequency EE and BB spectra between all BK18 bands and *Planck* 353 GHz (which is important for constraining polarized dust emission) are shown in Fig. 6.5. The spectra shows good agreement with the baseline lensed- ΛCDM + dust model (shown as red lines) used in the previous round of analysis, BK15, which had $r = 0$, $A_d = 4.7 \mu\text{K}^2$, $\beta_d = 1.6$, and $\alpha_d = -0.4$ (The BICEP2 and *Keck Array Collaborations* 2018). Note that the BK15 model did not

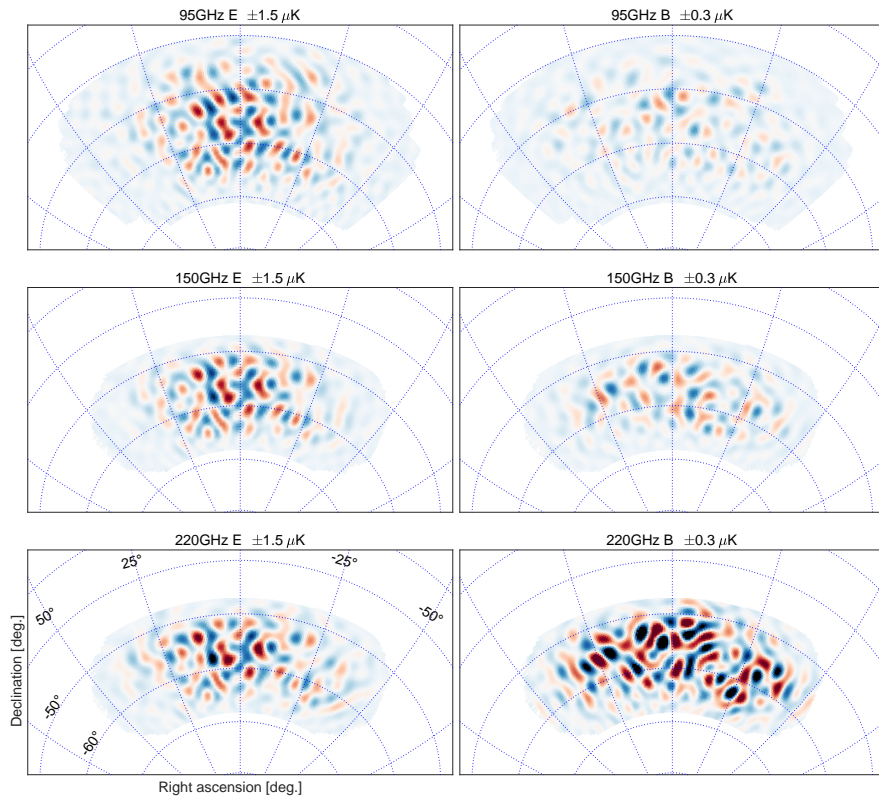


Figure 6.4: BK18 maps of E modes (left) and B modes (right) at 95 (top), 150 (middle), and 220 GHz (bottom), filtered to $50 < \ell < 200$. The color scale is given above each plot, in μK_{CMB} , with different color scales for E and B . The E modes are dominated by ΛCDM signal at all frequencies. The B modes at 95 GHz are roughly equal parts lensed- ΛCDM and noise, and the B modes at 150 and 220 GHz are dominated by thermal dust emission. Figure from BK18 ([The BICEP2 and Keck Array Collaborations 2021](#)).

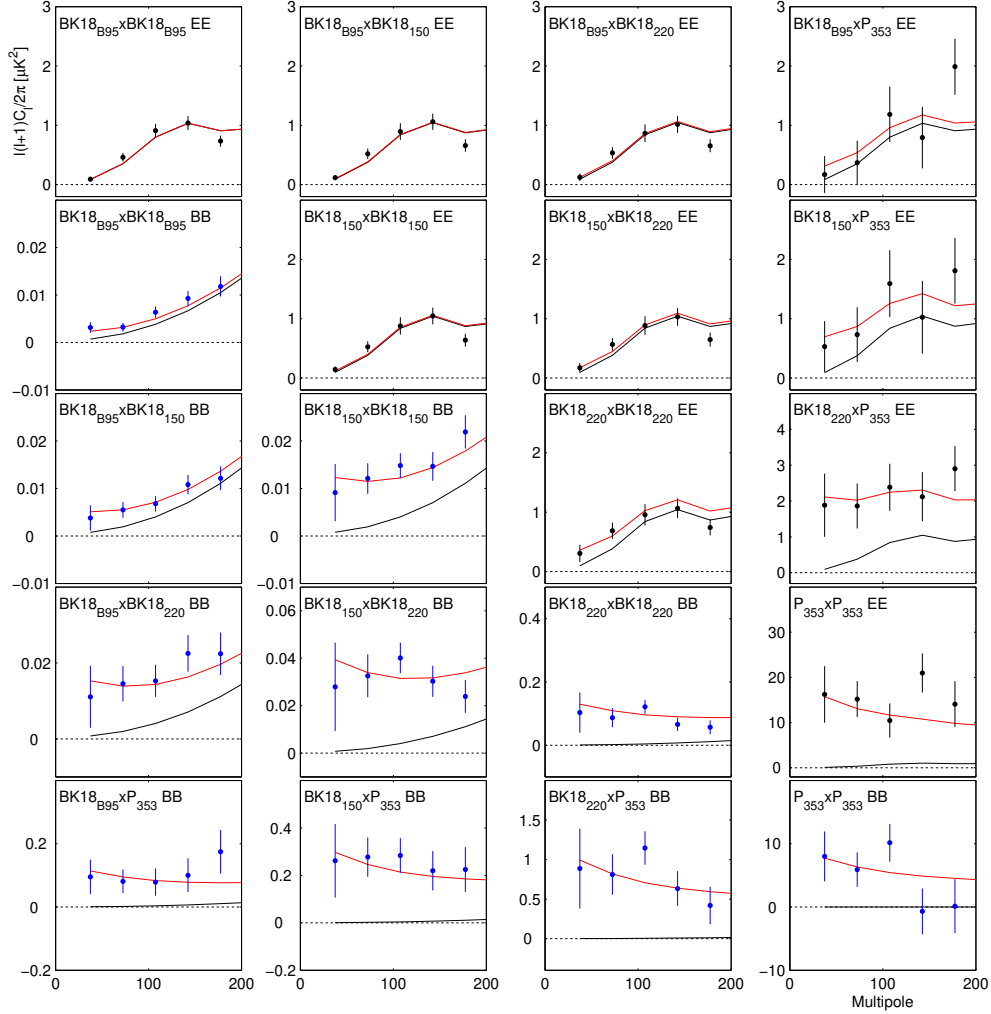


Figure 6.5: EE and BB auto-frequency and cross-frequency spectra using BK18 bands and *Planck* 353 GHz. The black lines show model expectation values for lensed- Λ CDM, and the red lines show the expectation values of the baseline lensed- Λ CDM + dust model used in the previous analysis, BK15, which was fit only to BB and did not include any BICEP3 data. The error bars are the standard deviation of the 499 lensed- Λ CDM + dust + noise simulations scaled to the BK15 model. Figure from BK18 (*The BICEP2 and Keck Array Collaborations 2021*).

have any BICEP3 data (though it did have four receiver-years of *Keck* 95 GHz), and was fit to only the *BB* data. The agreement between the new 95 GHz data dominated by BICEP3 and the previous baseline model therefore validates that model.

The 66 auto- and cross-frequency *BB* spectra between BK18, WMAP, and *Planck* are fed into the multicomponent likelihood analysis framework described in Section 6.2. We use nine bandpowers spanning $20 < \ell < 330$ per spectrum, compared against a lensed- Λ CDM + tensors + dust + synchrotron + noise model, with baseline parameters and priors specified as follows.

Dust is included with amplitude A_d evaluated at 353 GHz at $\ell = 80$, with frequency spectral shape defined by a modified blackbody (greybody) spectrum with $T_d = 19.6$ K and frequency spectral index β_d . Unlike previous analyses such as BK15, we use no prior on β_d , due to the improved constraining ability of the deeper 220 GHz maps. The spatial spectral behavior is assumed to be a power law $\mathcal{D}_\ell \propto \ell^{\alpha_d}$ where $\mathcal{D}_\ell = \ell(\ell + 1) C_\ell / 2\pi$. A wide uniform prior is used for the spatial spectral index, $-1 < \alpha_d < 0$.

Synchrotron is similarly parameterized, with amplitude A_s evaluated at 23 GHz at $\ell = 80$, and a power law frequency spectrum with frequency spectral index β_s . A Gaussian prior of $\beta_s = -3.1 \pm 0.3$ is chosen based on WMAP 23 and 33 GHz data (Fuskeland et al. 2014). The spatial spectral shape is taken to be a power law $\mathcal{D}_\ell \propto \ell^{\alpha_s}$ with a uniform prior $-1 < \alpha_s < 0$.

The dust/synchrotron correlation parameter is used, with a uniform prior $-1 < \epsilon <$

1. This baseline scenario fixes the scaling of the lensing amplitude $A_L = 1$ as well as the decorrelation parameters $\Delta'_d = \Delta'_s = 1$ (corresponding to no decorrelation).

The 1D and 2D marginalized posteriors comparing the new results to the previous BK15 baseline scenario are shown in Fig. 6.6. The constraint on the tensor-to-scalar ratio has tightened to $r = 0.010^{+0.010}_{-0.009}$ (corresponding to $r < 0.034$ at 95% confidence), with a zero-to-peak likelihood ratio of 0.60. The value for the dust frequency spectral index is $\beta_d = 1.51^{+0.13}_{-0.12}$. In BK15, β_d was mostly constrained by the prior – here we see with no prior, the constraining power of the new receiver-years at 220 GHz has resulted in a distribution nearly as tight as the prior-dominated BK15 constraint. The amplitude constraints are $A_d = 4.4^{+0.8}_{-0.7} \mu\text{K}^2$, $A_s < 1.4 \mu\text{K}^2$ at 95% confidence.

We use a maximum-likelihood estimation as an alternative method for evaluating the peak of the global posterior. The parameter values that maximizes the likelihood include: $r = 0.009$, $A_d = 4.4 \mu\text{K}^2$, $A_s = 0.6 \mu\text{K}^2$, $\beta_d = 1.5$, $\beta_s = -3.1$, $\alpha_d = -0.68$, $\alpha_s = 0.00$, and $\epsilon = -0.13$. The probability to exceed (PTE) the observed χ^2 value is 0.89, demonstrating an acceptable model fit to the data.

6.4 BK18 T-TO-P LEAKAGE SIMULATION BASELINE RESULTS

In this section we present the primary results of the $T \rightarrow P$ leakage simulations run for the BK18 data set, using per-detector composite beam maps at the 95, 150, and

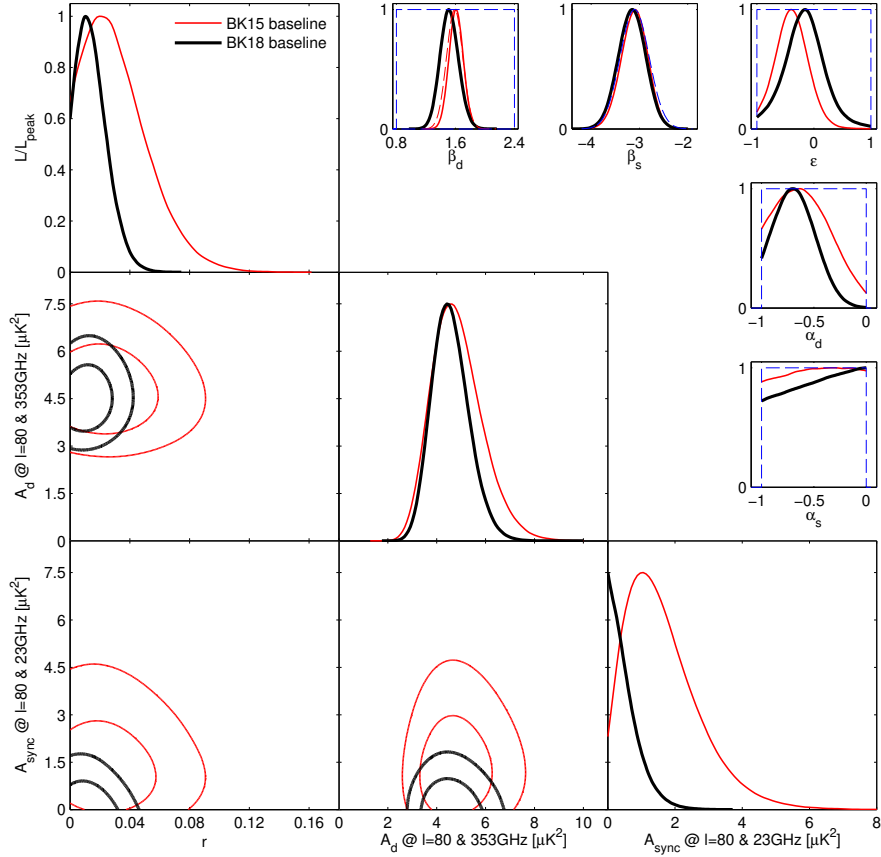


Figure 6.6: BK18 1D and 2D marginalized posteriors evaluated using MCMCs implemented by CosmoMC. The black lines show the new constraints and the red lines show the BK15 baseline results (*The BICEP2 and Keck Array Collaborations 2018*). The parameters plotted here include the dust and synchrotron amplitudes (A_d and A_s), frequency spectral indices (β_d and β_s), and spatial spectral indices (α_d and α_s), as well as dust/synchrotron correlation ϵ . The new upper limit on the tensor-to-scalar ratio shown here is $r < 0.034$ at 95% confidence. The 1D parameter posteriors in the top right also include the priors as blue dashed lines. This analysis uses no prior on the dust frequency spectral index β_d , while a Gaussian prior was used in BK15. Figure from BK18 (*The BICEP2 and Keck Array Collaborations 2021*).

220 GHz frequency bands. These results use composite beam maps made by stacking all component maps for each detector over all years where the receiver detector optics were unchanged. The beam maps for all years were re-analyzed using the updated beam map analysis pipeline described in Section 4.2. The baseline results here focus on beam maps out to a radius of 2° from the main beam (for an analysis of $T \rightarrow P$ leakage simulations at multiple beam radii, see [Karkare 2017](#)). After presenting the baseline $T \rightarrow P$ leakage results, we explore the impact of crosstalk on $T \rightarrow P$ leakage in Section 6.5.

6.4.1 BASELINE RESULTS: LEAKAGE MAPS AND POWER SPECTRA

The $T \rightarrow P$ leakage simulations are made by observing a noiseless *Planck* T sky, with no input Q/U signal. The simulated timestream signal is obtained by convolving that detector’s composite beam map with the T sky, then observing it using the real detector pointing timestreams. The resulting timestreams undergo the same operations as do the real data and primary CMB simulations – filtering, deprojection, cuts, and finally binning into Q/U maps and Fourier transformed to power spectra, after matrix purification. The cuts and weights are taken directly from the real data. The resulting polarized leakage spectra quantify the $T \rightarrow P$ leakage due to any mismatch between the beams of a detector pair. This process is repeated using per-detector split maps (see Section 4.2.5), to estimate the leakage contribution from noise in the beam map measurement.

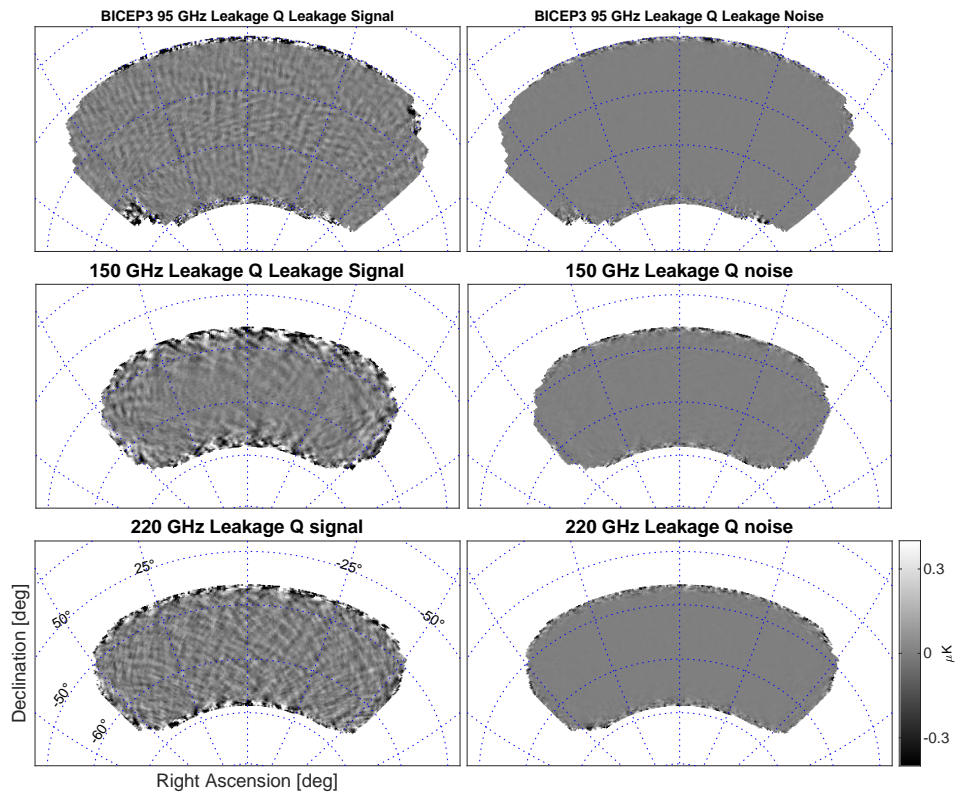


Figure 6.7: BK18 leakage Q maps, as predicted by $T \rightarrow P$ leakage simulations for BICEP3 95 GHz (top row) and *Keck* 150 (middle row) and 220 GHz (bottom row). The left column shows the leakage estimated using per-detector composite beam maps, and the right column shows the leakage associated with noise in the beam map measurement, estimated using split beam maps. Note the color scale is 10% of the BK18 real Q maps in Fig. 6.3.

The leakage simulation Q maps for the BK18 data set are shown in Fig. 6.7. These maps show predicted leakage for BICEP3 95 GHz and *Keck* 150 and 220 GHz. Although *Keck* 95 GHz is not explicitly shown here, it is still used in the final data set (no new *Keck* 95 GHz data was taken after the BK15 data set). These are dp1102 maps, chosen to match what is done with the real data. These maps are inspected before proceeding to power spectra, to check for any distinct patterns or cancellation trends. For example, most maps tend to have less predicted leakage at the center of the map, and higher leakage toward the edge. This is because the center of the map is more likely to be hit by more detectors at multiple boresight rotation angles, allowing partial cancellation of any unprojected residual beam patterns with certain symmetries, especially if that pattern is common-mode across detectors. Conversely, map pixels near the edge are more likely to be hit only by detectors near the edge of the focal plane, over fewer rotation angles per detector. For a deeper discussion on leakage cancellation, see Section 5.7 and [Sheehy 2013](#).

Comparing leakage maps at different levels of deprojection can highlight the amount of power in each differential mode. Fig. 6.8 shows Q signal maps for the BK18 frequency bands using varying levels of deprojection (dp0000, dp1100, dp1102, and dp1112). The dp0000 maps show a stripe of lower leakage in the middle of the maps, particularly at 150 and 220 GHz. This is likely from effective cancellation of differential pointing over multiple boresight rotation angles. Since the stripe is narrow and does not go to the edge of the map, this shows the cancelled mode is not common across

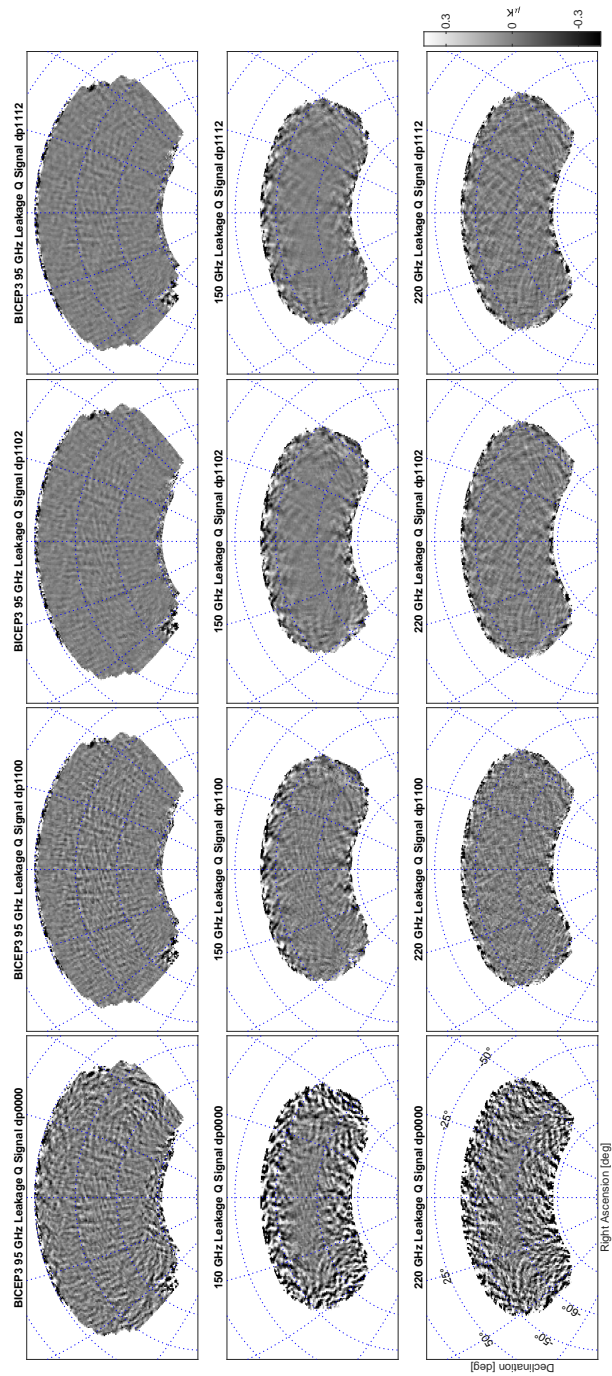


Figure 6.8: BK18 leakage Q signal maps, as predicted by $T \rightarrow P$ leakage simulations for BICEP3 95 GHz (top row) and $K_{\text{eff}}: 150$ (middle row) and 220 GHz (bottom row). Each column shows a different level of deprojection, incrementing from left to right: dp0000, dp100, dp1102, dp1112. The dp0000 maps show less leakage in the center of the map, demonstrating effective cancellation over boresight rotation. Most differential beam power is in differential gain and pointing, seen by the decrease in signal going from dp0000 to dp1100. Subtracting differential ellipticity (dp1102) slightly lowers leakage, mostly at smaller scales. Finally, deprojecting beamwidth (dp1112) only negligibly changes the predicted leakage.

all detectors, i.e. differential pointing is not the same in magnitude and direction across the FPU (this can also be seen in Fig. 5.1). It is clear that most leakage power is removed upon deprojecting differential gain and pointing (dp1100). The smaller difference between dp0000 and dp1100 for BICEP3 (as well as the lower overall dp0000 leakage compared to *Keck*) indicates that BICEP3 has very little differential beam power, even before deprojection. This agrees with the lower value differential beam parameters seen in BICEP3 compared to *Keck*, e.g. in Table 5.3. Subtracting differential ellipticity (dp1102) slightly lowers the leakage, especially on relatively finer scales. Including deprojection of differential beamwidth (dp1112) makes a negligible difference in the final results, which is why it is not deprojected or subtracted in the standard analysis.

The leakage maps are then apodized in the same way as the real data – using a common mask that is the geometric mean of the inverse variance maps over all *Keck* frequency bands; BICEP3 uses its own inverse variance map due to its larger field. We can then create auto-frequency and cross-frequency power spectra to examine the leakage in each bandpower. This is done using leakage simulations run with composite beam maps to estimate the signal, and with the split beam maps to quantify noise. The auto-frequency spectra for all BK18 bands in Fig. 6.9. The signal-to-noise is > 100 at $\ell < 200$. The leakage in *EE* and *BB* is roughly equal, especially near $\ell \sim 100$, which is expected for some general, incoherent undeprojected residual pattern.




Figure 6.9: BK18 leakage dp1102 EE (left) and BB (right) spectra estimated from $T \rightarrow P$ leakage simulations. The leakage signals from composite beam maps are shown in solid lines, and the noise signals from split beam maps are dashed lines. In general, the leakage noise levels are at least two orders of magnitude below the leakage signal, especially at low ℓ . The leakage in EE and BB is nearly equal in each band, with a few exceptions, such as the higher EE leakage at $\ell \sim 200$ for BICEP3 and *Keck* 95 GHz, and the higher BB in *Keck* 220 GHz at low ℓ . Each of these trends is discussed in the text.

Around $\ell \sim 150 - 250$, the BICEP3 and *Keck* 95 GHz leakage is slightly higher in EE than in BB . In the case of *Keck*, there is a known anomalous coupling between tile edge pixels and the tile corrugations, and this pattern is aligned with the detector polarization axes (aligned with a “+” pattern in the (x', y') coordinate system introduced in Section 4.2.2), which preferentially leaks into E . This corrugation coupling was reduced in future focal planes (the *Keck* 95 GHz FPUs were retired after 2015), though the EE leakage spectra of BICEP3 indicate they may still be present at a weaker level, and at higher ℓ (which would be expected due to the smaller BICEP3 beam size). It has not yet been confirmed that this shape in BICEP3 EE is strictly due to tile corrugation coupling – this is still being investigated, see Section 7.1.2.

At 220 GHz, there are interesting features in both EE and BB . In EE , there is a no-

ticeable bump near $\ell \sim 150$, which is likely due to crosstalk beams originally discussed in Section 5.6. The spacing between nearest readout neighbors, and therefore the separation between the main beam and the crosstalk beam for a detector, is 1.5° , which corresponds to $\ell \simeq 120$, very close to the peak seen in 220 GHz *EE*. In *BB*, there is also higher leakage, which could correspond to the differential beam pattern seen from direct-island coupling. This pattern, which is discussed further in Section 7.1.2, roughly aligns with an “ \times ” pattern in (x', y') and therefore leaks preferentially to *B*. This pattern is brighter in the new standalone 220 GHz band (which did not exist in BK15) than the standalone 230 GHz band, though identical in shape, which leads to a larger predicted leakage in the coadded BK18 220 GHz band compared to the previous BK15 220 GHz leakage. ([The BICEP2 and Keck Array Collaborations 2019](#)).

6.4.2 BASELINE RESULTS: CROSS SPECTRA WITH REAL MAPS

Next we take the leakage maps and cross them with the real maps shown in Section 6.3. Ideally, the cross spectra would completely overlap with the auto spectra, giving us confidence that the leakage predicted by the $T \rightarrow P$ leakage simulations truly exists in the real CMB data. Any difference between the cross and auto spectra should be able to be explained as physically-motivated differences between the measured far-field beam maps or simulation method and the real beams on the sky. The small amplitude of unprojected residuals, combined with their complex shape that generally varies across frequencies, makes this a difficult challenge. The ultimate end

goal is to place realistic upper limits on the expected bias on r , based on the leakage predicted by the simulations and seen in the real data. Again, these baseline results use beam maps that go out to a radius of 2° from the main beam, using beams analyzed with the updated FFBM analysis pipeline.

There are a few different possible noise sources when crossing $T \rightarrow P$ leakage simulations with real data: 1) noise in the beam map measurement, 2) noise in the CMB maps, and 3) coupling of the undeprojected residual beams to other astrophysical signal, such as dust and lensing. The noise in the beam map is subdominant, as shown in Fig. 6.9, so we assume the noise in the cross is primarily from the other two factors. We quantify the noise by crossing the fixed leakage simulations with 499 standard simulations of lensed- Λ CDM + dust + noise (so-called “type 8” simulations). The per-bin uncertainty in the leakage simulation cross with real data is taken to be the standard deviation of this cross with 499 type 8 simulations.

Here we introduce a couple metrics used to quantify the per-frequency band BB leakage, and the agreement between the cross spectra and the auto spectra. First, we define the ρ estimator, which is a quadratic estimator originally defined in [Aikin 2013](#), and also used in BK XI ([The BICEP2 and Keck Array Collaborations 2019](#)). This is a single number that best represents the scaling of a $r = 1$ primordial B -mode spectrum to a chosen spectrum, in this case a leakage BB spectrum. More formally, consider $\hat{\mathbf{C}}$ to be a set of bandpowers predicted by the leakage simulation, $\langle \mathbf{C} \rangle$ to be a set of expected bandpowers for an $r = 1$ theory spectrum (observed by our pipeline), and \mathbf{N}

as the bandpower covariance matrix from lensed- Λ CDM + dust + noise simulations.

The goal is to find a minimum variance estimate of ρ , where ρ is implicitly defined as:

$$\hat{\mathbf{C}} = \rho \langle \mathbf{C} \rangle + n, \quad (6.10)$$

where n is a noise term. The minimization of the quantity $\hat{\mathbf{C}} - \rho \langle \mathbf{C} \rangle$ leads to:

$$\hat{\rho} = \frac{\langle \mathbf{C} \rangle^T \mathbf{N}^{-1} \hat{\mathbf{C}}}{\langle \mathbf{C} \rangle^T \mathbf{N}^{-1} \langle \mathbf{C} \rangle}. \quad (6.11)$$

Throughout the rest of this dissertation, we use ρ to refer to $\hat{\rho}$. Since ρ is proportional to r , this estimator informs us the “amount of r ” seen in a BB spectrum. In general for BICEP/*Keck Array* receivers, a vast majority of the total weight comes from the first three bins ($\ell = 37.5, 72.5, 107.5$) (Barkats et al. 2014). Note that the ρ values discussed in this section are not themselves estimates of the final bias on r – this is calculated using maximum-likelihood searches in the next section.

The next metric is useful for determining the confidence that a given cross between leakage simulation and real data agrees with a model consistent with the full simulation auto spectrum, or with a null spectrum. Formally, consider the bandpowers of a cross spectrum to be \mathbf{d} and the bandpowers corresponding to a model (either the auto spectrum or a null spectrum) to be \mathbf{m} . We can then calculate a simple χ^2 and χ

statistic:

$$\chi^2 = (\mathbf{d} - \mathbf{m})^T \mathbf{N} (\mathbf{d} - \mathbf{m}) \quad (6.12)$$

$$\chi = \sum_i \frac{d_i - m_i}{\sigma_i} \quad (6.13)$$

where \mathbf{N} is the bandpower covariance matrix evaluated from the lensed- Λ CDM + dust + noise simulations. We use both χ and χ^2 to quantify the agreement between the cross spectra and the auto spectra or null. Since we know the number of degrees of freedom (the number of bandpowers), we can also calculate the probability to exceed (PTE) using the respective cumulative distribution functions of χ and χ^2 .

Fig. 6.10 shows the BK18 leakage BB spectra crossed with real data, including ρ and PTE metrics. The statistics are summarized in Table 6.1. Overall, the cross spectra are at an amplitude equal to about half the auto spectra; this can also be seen by comparing the auto spectra and cross spectra ρ estimates. They are generally positive and follow the auto spectra in shape. This is similar to what was seen in the BK15 $T \rightarrow P$ leakage analysis ([The BICEP2 and Keck Array Collaborations 2019](#)). The uncertainty in the cross spectra ρ come from applying Eq. 6.11 to the standard deviation of the cross between fixed leakage simulations and 499 realizations of lensed- Λ CDM + dust + noise.

Although the cross spectra ρ values and the PTE values do not indicate a confident detection of leakage in the real data, these cross spectra still provide valuable upper

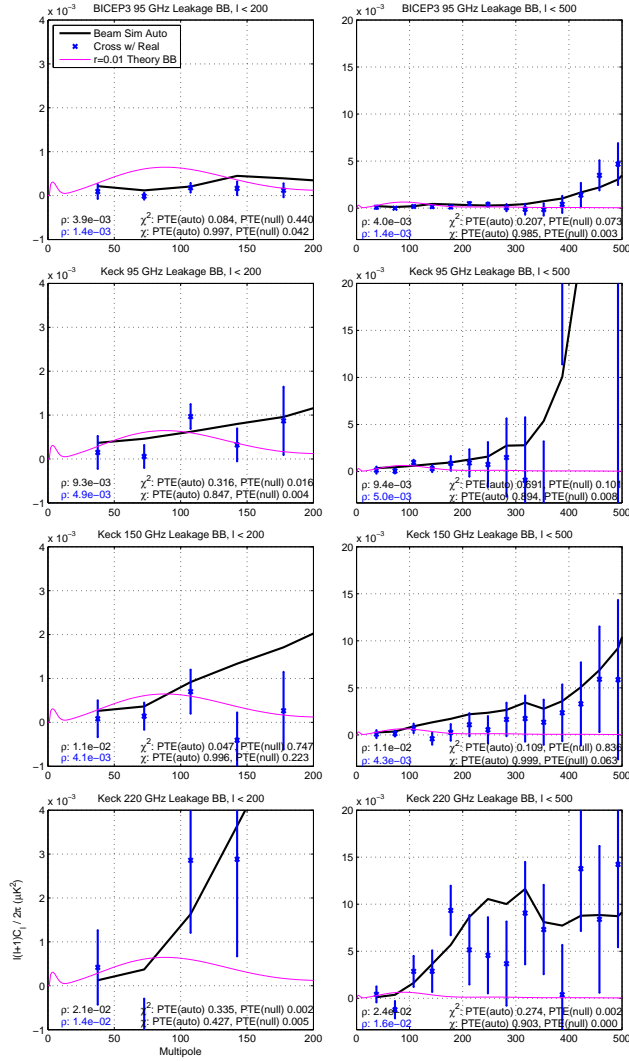


Figure 6.10: BK18 leakage dp_{1102} BB spectra estimated from $T \rightarrow P$ leakage simulations, crossed with real BK18 maps. This includes BICEP3 95 GHz (first row), *Keck* 95 GHz (second row), *Keck* 150 GHz (third row), and *Keck* 220 GHz (fourth row). Each plot shows the auto spectrum as a black line, and the cross with real data as blue crosses, with error bars taken as the standard deviation of the cross between the fixed leakage simulation and 499 realizations of lensed- Λ CDM + dust + noise. The left column ($\ell < 200$) and the right column ($\ell < 500$) show the same data, to different scales. The ρ and PTE metrics are shown as evaluated using only the bins included in that scale, i.e. the left column evaluates the metrics using only the first five bins, and the right column evaluates them using the first nine bins. In general, the cross spectra have a shape similar to the auto spectra, but with half the amplitude. The best agreement with the auto occurs in BICEP3 (with nine bins) and *Keck* 220 GHz.

Frequency Band	ρ (auto)	ρ (cross)	χ^2 PTE (auto)	χ^2 PTE (null)	χ PTE (auto)	χ PTE (null)
BICEP3 (95 GHz)	4.0×10^{-3}	$(1.4 \pm 2.6) \times 10^{-3}$	0.207	0.073	0.985	0.003
Keck 95 GHz	9.4×10^{-3}	$(5.0 \pm 6.4) \times 10^{-3}$	0.691	0.101	0.894	0.008
Keck 150 GHz	1.1×10^{-2}	$(4.3 \pm 8.3) \times 10^{-3}$	0.109	0.836	0.999	0.063
Keck 220 GHz	2.4×10^{-2}	$(1.6 \pm 2.5) \times 10^{-2}$	0.274	0.002	0.903	0.000

Table 6.1: Summary ρ and PTE statistics from the cross spectra between $T \rightarrow P$ leakage simulations and real data, corresponding to Fig. 6.10. In general, the cross spectra ρ values are about half that of the auto spectra ρ . The uncertainty in cross spectra ρ comes from calculating ρ on the $1\text{-}\sigma$ uncertainty on the cross between the fixed leakage simulation and 499 realizations of lensed- Λ CDM + dust + noise. The large uncertainty on ρ and the PTE statistics do not indicate a confident detection of leakage in the real data. However, the cross spectra can still provide valuable upper limits on the expected leakage.

limits on the expected leakage, which can be used with the maximum-likelihood algorithm to quantify a bias on r . Similar to the story with the auto spectra, the BICEP3 leakage cross spectra are lower than any *Keck* leakage cross spectra, by a factor of 3 – 10 in ρ . As we will see in the next section, three full years of observations with BICEP3, which has significantly less power in the difference beams (both before and after deprojection) drives the improvement in the estimate on the bias on r .

Given the lingering factor of ~ 2 between the cross spectra and the auto spectra, we are motivated to identify the root cause of this lack of confident detection of leakage in the real data. We have already seen that the effect of over-smoothing in the beam map analysis pipeline can significantly affect the overall beam shape (see Fig. 4.9) – these improvements alone do not drastically improve the cross with real data. Another possible difference between the beam maps as measured vs. the real beams on sky is crosstalk. As shown in Section 5.6, the current hypothesis for out-of-phase settling time crosstalk implies that the amount of crosstalk seen in the beam maps overestimates the crosstalk seen in the real CMB data by 1-2 orders of magni-

tude. The impact of crosstalk on $T \rightarrow P$ leakage is examined in Section 6.5.

6.4.3 BASELINE RESULTS: BIAS ON r FROM MULTICOMPONENT LIKELIHOOD ANALYSIS

The final step in this BK18 baseline $T \rightarrow P$ leakage analysis is to quantify a best estimate of the bias on r due to the presence of $T \rightarrow P$ leakage predicted from the simulations. The methodology used for the BK18 results here closely follows what was done for BK15 in [The BICEP2 and Keck Array Collaborations 2019](#). In general, there are a few options for how to account for the impact on r :

- If the specific amplitude and shape of the $T \rightarrow P$ morphology is known and confidently detected in the real data, it can be subtracted from the real data (similar to deprojection).
- If there were intermediate knowledge of the morphology, e.g. a well-characterized amplitude with uncertain shape, it could be debiased in the power spectra. The uncertainty of the leakage spectra would remain, inflating the overall uncertainty.
- If there is significant uncertainty on both the form and amplitude of the leakage, then the best option is to run simulations to estimate the possible amplitude of the leakage and ensure it is significantly below the experiment statistical uncertainty.

Given the large uncertainties seen in the cross between leakage simulations and real data presented in Section 6.4.2, we elect to use the third approach. We inject leakage into the set of 499 lensed- Λ CDM + dust + noise simulations in the power spectrum domain, using one of the leakage template options described below, then

find the shift in the value of r at the global maximum of the likelihood. For BK18, we consider three scenarios – two of which are “upper limit-driven” and one that is “CMB data-driven.” In each of these scenarios, the leakage is ignored in the likelihood.

The two upper limit scenarios involve adding the $T \rightarrow P$ leakage simulation auto spectra (black lines in Fig. 6.9) to the CMB simulations. In the first scenario, the full leakage auto spectrum is added to each CMB auto-frequency spectrum, and a spectrum equal to the geometric mean between leakage auto spectra at two frequencies is added to each CMB cross-frequency spectrum. Given the clear discrepancy between the leakage auto spectra and cross spectra with real, we expect this scenario to be an overestimate of the true bias. The median shift found in this scenario is $\Delta r = 0.0036$.

The second upper limit scenario also injects the leakage auto spectra in a similar way, except for each of the 499 realizations, the amplitude of the auto spectra templates added is scaled by a random variable drawn from a Gaussian distribution centered at zero with width $\sigma = 0.5$ of the nominal full amplitude, truncated at zero to allow only positive leakage injected. The value of 0.5 is chosen for the Gaussian distribution because the leakage cross spectra with real have amplitudes that are roughly half the amplitude of the leakage auto spectra. In this case, we find a median bias of $\Delta r = 0.0014$.

Scenario 3, the CMB data-driven scenario, uses the leakage cross spectra with real data (blue crosses in Fig. 6.9) as the injected templates. The leakage cross spectra are only added to the CMB auto-frequency spectra, since the cross-frequency leakage

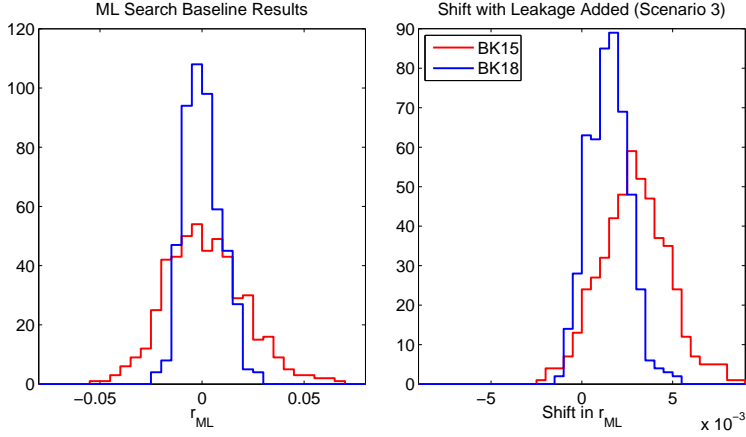


Figure 6.11: Histograms of the recovered value of r_{ML} from an ML search over 499 realizations of the baseline scenario (left) and the shift in r_{ML} after re-running the ML search with $T \rightarrow P$ leakage injected as described in the “CMB data-driven” scenario. The red histogram shows the results from the previous analysis, BK15 ([The BICEP2 and Keck Array Collaborations 2019](#)), and the blue histogram shows the new BK18 results ([The BICEP2 and Keck Array Collaborations 2021](#)). Note the different x-axis scale between both plots.

spectra cross with real data are more consistent with zero. Each of the 499 realizations is additionally biased by a random contribution drawn from a Gaussian centered at zero with standard deviation equal to the standard deviation of the cross between the leakage simulations and the lensed- Λ CDM + dust + noise simulations. This scenario was quoted as the baseline result of the BK15 $T \rightarrow P$ leakage analysis ([The BICEP2 and Keck Array Collaborations 2019](#)), and we continue to use it as the baseline in this BK18 leakage analysis. Using the median and 14th/86th percentiles of the distribution of shifts in r , we find a bias of $\Delta r = 0.0015 \pm 0.0011$. Histograms showing the realization-by-realization shifts in the maximum-likelihood value of r for this scenario are shown in Fig. 6.11.

In the BK15 $T \rightarrow P$ analysis we also considered three other “recovery” scenar-

ios, where nuisance parameters were introduced into the multicomponent likelihood framework in an attempt to fit and remove the $T \rightarrow P$ leakage in each band. In that analysis, we found that any attempt to recover the leakage (whether it was artificially injected similar to scenarios 1-3, or if it was only the leakage in the real data) led to a possible negative bias in r . For this reason, we did not use any of these as the baseline scenario in the BK15 analysis, and we did not run these scenarios for BK18.

Going from BK15 to BK18, the experiment statistical uncertainty on r decreased from $\sigma = 0.020 \rightarrow 0.009$. The decrease in the reported bias on r from $T \rightarrow P$ leakage ($\Delta r = 0.0027 \pm 0.0019 \rightarrow 0.0015 \pm 0.0011$) is fractionally very similar. It is reassuring that the bias (and uncertainty thereon) remains subdominant to the experiment statistical uncertainty – however, the bias is only a factor of $\sim 8 - 10$ smaller than $\sigma(r)$. This margin may be comfortable for now, but with the rapid improvement of current and future generation CMB experiments, there is still improvement to be desired in order to keep this systematic controlled going forward. This provides ample motivation to reduce the $T \rightarrow P$ leakage from beam mismatch on many fronts – detector fabrication issues leading to beam mismatch at the focal plane; imperfections in the optical elements, especially filters, lenses, and AR coats; systematics in the FFBM measurement and analysis procedure leading to a misrepresentation of the true far-field differential beam shape. All of these are heavy areas of focus in future experiments such as BICEP Array and CMB-S4, and motivate the use of optical modeling software, in conjunction with achieved past performance, to optimize the telescope

detector and optical design for these experiments. We visit this topic in Chapter 7.

6.5 IMPACT OF CROSSTALK ON T-TO-P LEAKAGE

Due to crosstalk in our time-domain SQUID-based multiplexed readout, most detectors see an additive signal equal to $\sim 0.3\%$ of the next-nearest neighbors’ (“upstream” and “downstream”) signal in the readout. This is due to inductive crosstalk, where the suspected mechanism is mutual inductance between neighboring SQ1 that are coupled to each TES detector. A small subset of detectors see a high level of downstream crosstalk at a $1 - 3\%$ level which also appears out of phase by up to 90° from the main detector signal. This mechanism is due to residual signal from row switching (sampling one SQ1 after the other sequentially in time) before the previous SQ1 signal has been allowed to fully decay. For a full discussion, see Section 5.6. Here, we specifically consider the impact of crosstalk, particularly the elevated “row-switching” crosstalk, on the $T \rightarrow P$ leakage simulations.

Along with the six deprojection templates discussed in Section 6.1.2, we also have the option to observe deprojection templates corresponding to crosstalk. These templates simply correspond to the temperature signal sampled using the real trajectories of the nearest upstream and downstream neighbor detectors. Like any other deprojection mode, these templates can be regressed against the real detector timestream and removed. Because previous studies have shown crosstalk contamination in the

real maps to be small ($r \approx 3 \times 10^{-3}$ in BK III, [The BICEP2 Collaboration 2015](#)), this deprojection mode is not included in standard CMB analysis. Here, we enable these crosstalk deprojection modes (one for upstream, one for downstream) in both the real data and in the $T \rightarrow P$ leakage simulations, and compare how the real BB and predicted leakage BB bandpowers change. In this section, we use the fifth and sixth deprojection shorthand digits (e.g. dpXXXX00 or dpXXXX11) to refer deprojection of the upstream and downstream crosstalk templates, respectively.

We consider the *Keck* 2016 CMB observing year, which had two receivers at 220 GHz, two at 230 GHz, and one at 150 GHz. This 150 GHz receiver (*Keck* rx4) demonstrated high-amplitude, out-of-phase downstream crosstalk in the far-field beam maps (see Fig 5.11 and Fig. 5.12), while the other four receivers only have $\sim 0.3\%$ inductive crosstalk. The beam maps for the 150 GHz receiver show out-of-phase crosstalk, but because we are not confident that the crosstalk that may be seen in the real CMB data is out of phase (due to the hypothesis that the amplitude and phase varies with signal frequency, see Section 5.6), we regenerate the beam maps from all receivers in 2016 using the quadrature sum of the cosine and sine demodulation components⁴. By comparing real data and $T \rightarrow P$ leakage simulations with crosstalk deprojection enabled and disabled, we hope to answer three questions. 1) Does enabling crosstalk deprojection significantly shift the real BB bandpowers? 2) How do the leakage ρ esti-

⁴While making beam maps this way ensures the full amplitude of all beam signal is captured, it also rectifies noise, which is why this method is not done in the standard $T \rightarrow P$ simulations.

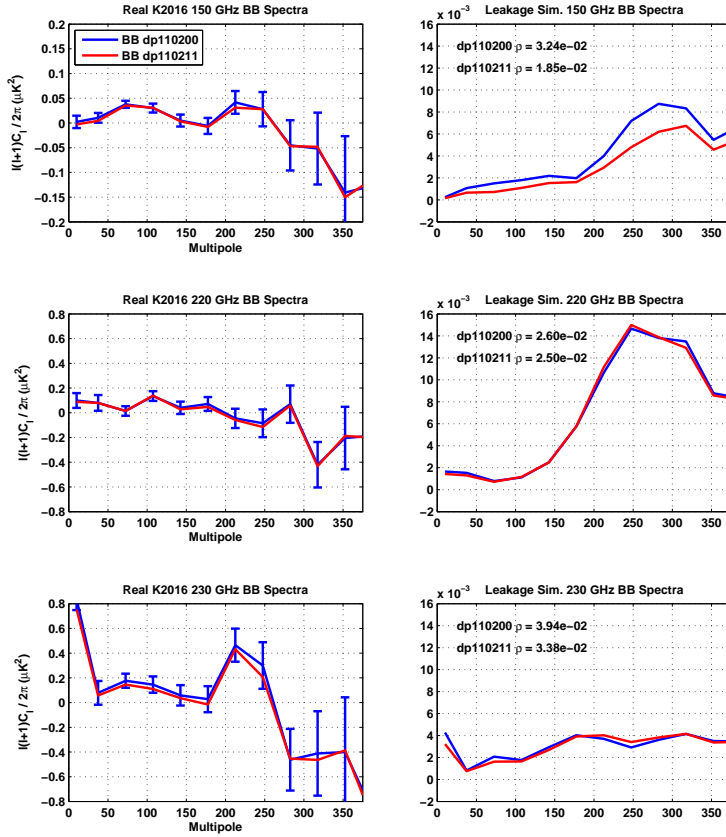


Figure 6.12: *Keck* 2016 real *BB* bandpowers (left column) and leakage *BB* predicted by $T \rightarrow P$ leakage simulations (right column), for 150 GHz (top), standalone 220 GHz (middle), and standalone 230 GHz (bottom). The blue lines are spectra using the standard deprojection dp110200 (without crosstalk), and the red lines additionally include crosstalk deprojection dp110211. The error bars on the real bandpowers are the standard deviation of 499 realizations of lensed- Λ CDM + dust + noise from 2016 alone. The real data spectra are noise debiased and have had the suppression correction factor applied. The same noise debiasing and suppression correction is done to the dp110211 data, which technically has slightly different filtering and therefore different suppression, but this is a small effect. The leakage spectra have had the same suppression correction factor applied, but are not debiased from noise in the beam maps. Including crosstalk deprojection leads to minimal change in the real bandpowers in any band, but reduces the leakage ρ estimate for the 150 GHz band roughly by a factor of 2. This could indicate that the high level of crosstalk seen in the beam maps may not be present in CMB observations.

mates change when enabling crosstalk deprojection in the $T \rightarrow P$ leakage simulations?

3) Do $T \rightarrow P$ leakage simulation cross spectra with real data significantly change?

The shift in the real and $T \rightarrow P$ leakage BB bandpowers for the *Keck* 2016 season are shown in Fig. 6.12. The uncertainty in the real BB bandpowers comes from the standard deviation of 499 realizations of lensed- Λ CDM + dust + noise from 2016 season alone. The shift in the real bandpowers after deprojecting crosstalk is smaller than 0.5σ in any bin, and the 150 GHz band does not seem to shift any more significantly than the other bands. Looking at the $T \rightarrow P$ leakage simulation bandpowers, however, we see a notable decrease in leakage BB at 150 GHz, resulting in a decrease in ρ by a factor of ~ 2 . This agrees with the general expectation from looking at the beam maps themselves – in many 150 GHz detectors (for this receiver specifically), the downstream crosstalk beam is as bright or brighter than the main beam unprojected residuals themselves. The fact that no such shift happens in the real BB bandpowers could suggest that this leakage is simply below the real data noise level, or that the crosstalk is not as bright in the CMB data, which would agree with the hypothesis that the amplitude (and phase) of the crosstalk signal is proportional to the frequency of the signal itself (the CMB is 1 – 2 orders of magnitude lower than the FFBM chop rate in frequency).

Any change in the $T \rightarrow P$ leakage cross spectra with real are of critical importance, since the cross spectra are used as the templates injected into real simulations when quantifying the bias on r (see Section 6.4.3). Although this study considers only the




Figure 6.13: *Keck* 2016 $T \rightarrow P$ leakage simulation cross with real data (BB), with crosstalk deprojection excluded (dp110200) and included (dp110211). The left plot shows 150 GHz, the middle shows standalone 220 GHz, and the right shows standalone 230 GHz. The solid and dashed black lines show the BB auto-spectra for dp110200 and dp110211, respectively. The blue crosses are the standard dp110200 cross with real data (2016 season only), with the error bars coming from the standard deviation of the cross between the fixed leakage simulation and 499 realizations of lensed- Λ CDM + dust + noise. The red crosses show the dp110211 cross with real data. The change in the cross spectra after enabling deprojection is small in all frequencies, especially compared to the change in the auto spectrum for 150 GHz, which again suggests that the bright FFBM crosstalk may not exist (as bright) in the CMB data.

Keck 2016 season and not all of BK18, any significant change in the cross with real data after deprojecting crosstalk could indicate that the bias Δr found for BK18 may be an underestimate over overestimate of the true leakage. Given that the auto spectra in Fig. 6.12 all decreased after deprojecting crosstalk (which is expected, since deprojection is removing power from the simulated maps due to bright crosstalk beams), it is more likely that the cross with real data will also decrease, meaning that the BK18 value of $\Delta r = 0.0015 \pm 0.0011$ is likely a slight overestimate, rather than an underestimate.

Fig. 6.13 shows the $T \rightarrow P$ leakage simulation cross with real data for the *Keck* 2016 observing season only, with crosstalk deprojection excluded and included. As usual, the uncertainties shown on the cross come from the standard deviation of the

Frequency Band	ρ (auto) dp110200	ρ (cross) dp110200	ρ (auto) dp110211	ρ (cross) dp110211
150 GHz	3.2×10^{-2}	$(1.3 \pm 1.8) \times 10^{-2}$	1.9×10^{-2}	$(1.3 \pm 1.8) \times 10^{-2}$
220 GHz	2.6×10^{-2}	$(5.1 \pm 5.0) \times 10^{-2}$	2.5×10^{-2}	$(4.5 \pm 5.0) \times 10^{-2}$
230 GHz	3.9×10^{-2}	$(-5.0 \pm 3.4) \times 10^{-2}$	3.4×10^{-2}	$(-3.2 \pm 3.4) \times 10^{-2}$

Table 6.2: Keck 2016 $T \rightarrow P$ leakage simulation ρ values, with crosstalk deprojection excluded (dp110200) and included (110211). These correspond to the spectra shown in Fig. 6.13. The uncertainty in the ρ for the cross with real data comes from evaluating ρ on the standard deviation of the fixed leakage simulation crossed with 499 realizations of lensed- Λ CDM + dust + noise (the error bars in Fig. 6.13). There is minor improvement between the leakage auto and cross with real data after including crosstalk deprojection for all bands, particularly 150 GHz. However, the uncertainties are relatively large, partially due to using only one season of data for this crosstalk study.

fixed leakage simulation crossed with 499 realizations of lensed- Λ CDM + dust + noise (using 2016 simulations only). The red X's relative to the blue X's indicate the shift in the cross with real data upon including crosstalk deprojection. The corresponding ρ values for all spectra are listed in Table 6.2. The shifts in the ρ values indicate slightly better agreement between cross with real data and the auto spectra, especially for 150 GHz, but the shifts relative to the uncertainties are quite small, just as we see with the real BB bandpowers in the left column of Fig. 6.12 (note the uncertainties are naturally larger in this study due to only using one observing season). In fact, even though the auto spectrum for 150 GHz is reduced by nearly a factor of two, the ρ value for the cross with real data remains unchanged upon deprojecting crosstalk. This seems to favor the hypothesis that the bright crosstalk seen in the far-field beam maps may not be present at such an amplitude in the real CMB data, though again, it is difficult to make a confident conclusion at this level of uncertainty.

The results here show that more could be learned by expanding this crosstalk deprojection analysis to include more years of real data, to beat down the noise and

obtain a more precise estimate on the shift in the $T \rightarrow P$ leakage due to crosstalk seen in the CMB data. However for now, the results from *Keck* 2016 indicate that the leakage contribution from crosstalk is likely an overestimate. It is important to note that $T \rightarrow P$ leakage simulations presented in this section and in Section 6.4.3 *naturally include* all leakage contributions from crosstalk. The BK18 bias on r of $\Delta r = 0.0015 \pm 0.0011$ includes the effect of crosstalk in all receiver-years, and now we have reason to believe that the crosstalk contribution encapsulated in that number is an overestimate. This number is an improvement over the systematic uncertainty of $r \approx 3 \times 10^{-3}$ due to crosstalk quoted in BK III ([The BICEP2 Collaboration 2015](#)), which assures us that the leakage due to crosstalk is reduced in *Keck* and BICEP3 compared to BICEP2.

*The limits of the possible can only be defined by
going beyond them into the impossible.*

Arthur C. Clarke

7

Beam Metrics and Model Comparisons

The bias on the measurement of r from $T \rightarrow P$ leakage predicted from the leakage simulations in Chapter 6 is subdominant to the overall statistical uncertainty for the BK18 analysis, but this statistical uncertainty is expected to decrease with future CMB experiments. CMB-S4 – the next generation ground-based CMB experiment – is expected to constrain $r < 0.001$ at 95% confidence in the absence of a detection

([The CMB-S4 Collaboration 2020](#)). This means we will need better control of this systematic than the bias of $\Delta r = 0.0015 \pm 0.0011$ found for BK18.

There are many experimental design aspects to consider in order to build a series of telescopes with exquisite beam control. How do the detectors couple to free space from the focal plane? Phased array antennas, lenslet arrays, and feed horn arrays each have their own advantages and disadvantages regarding ease of fabrication, proven performance, and possible systematics associated with the beam shaping elements. What should the configuration and constituent material be for the optical elements? The optical, thermal, and mechanical properties must be considered in order to guarantee well-focused beams for all detectors in a receiver. Given the optical configuration, what amount of ghosting and sidelobe levels should be expected? This just a sample of many questions that must be considered, simulated, and verified using past performance before implementing these designs in a real telescope (see [Abitbol et al. 2017](#) for details on each of these different choices).

The modeling and review of past performance needed to make all these decisions is beyond the scope of this paper. However, in this chapter, we review the development of a small-aperture telescope optical model using the GRASP physical optics modeling software, with emphasis on comparisons and feedback to real near-field and far-field beam maps discussed in Chapters 4 and 5. We also discuss metrics used to evaluate characteristics of the beam pattern, including edge taper at the aperture, and estimates of $T \rightarrow P$ leakage without the need of full timestream simulations. These

metrics can be used to facilitate the comparisons between optical models and data, and to identify correlations or trends between two optical properties. For instance, is a detector pair with high edge taper at the aperture more likely to have high levels of unprojected residuals in the far field? The utility of such metrics can prove to be of great value when validating baseline designs or exploring alternative designs for small aperture telescopes.

We start with an introduction to simple beam metrics, evaluated on real near-field and far-field beam maps. We explore comparisons to other experimental quantities such as optical loading from coupling to warm baffling structure. We emphasize the beam map-derived ρ metric, which is analogous to the full-timestream metric developed in Section 6.4.2. We then move to the GRASP model efforts, comparing outputs in the near field and far field. Finally, we discuss a second, simpler modeling code that can be used to derive a beam shape at any distance from the receiver, given a beam distribution at the aperture. BICEP3 will be featured in most of the discussion and examples presented in this chapter, since it was the focus of the GRASP modeling work.

7.1 SIMPLE BEAM METRICS

7.1.1 NFBM METRICS

We start with near-field beam maps (NFBM), which are originally discussed in Section 4.3. These beam measurements are a proxy for the illumination at the aperture – in practice, the measurement takes place ~ 1 m above the aperture. They are primarily sensitive to the phase of the beam distribution at the focal plane, as opposed to the far-field beam maps which are more sensitive to the amplitude. Historically in the BICEP/*Keck Array* program, they have been very useful for diagnosing detector fabrication issues, and non-idealities in the full optical chain.

We consider two separate NFBM metrics: *beam steer* and *edge taper*. Beam steer arises when variation in the thickness of the index of refraction in the inner-layer dielectric “pushes” a beam in the aperture plane off-center, sometimes causing severe truncation at the aperture stop. Since beam steer arises from dielectric thickness variation or refractive index variation, this is more useful for receivers that use phased array antennas (such as BICEP/*Keck Array*) or lenslet arrays. Edge taper is more general; it refers to the amount of truncation a beam undergoes at the aperture stop. This quantity is of great importance regardless of detector coupling – the harder the beam is stopped, the higher the response in the sidelobes. Some detectors may experience higher-than-intended edge taper, whether due to vignetting, beam steer or

otherwise, which could cause distorted beams in the far field, or abnormally extended power that may couple to warm baffling or the ground. Much of the development of both the beam steer and edge-taper metrics were done in collaboration with James Cornelison of the Harvard University CMB Group.

The beam steer is simply the spatial offset between the center of the beam at the aperture plane and the center of the aperture. This work follows similar analysis done for BICEP2 and *Keck* 2012 receivers (Wong 2014). For a completely ideal system, the detector beam should be centered in the middle of the aperture, regardless of the location of the detector in the focal plane. In order to calculate this from NFBM data taken for BICEP/*Keck Array* receivers, we must first account for the natural divergence of the beams between the aperture and the plane of the NFBM instrument. In principle these numbers can be explicitly measured, using knowledge of the receiver and NFBM apparatus dimensions. To account for possible variation in standard mounting/dismounting of the apparatus to the receiver, we elect to fit a four-parameter model to the data, where the parameters describe describe the position and orientation of the NFBM instrument, relative to the aperture plane. Given the ideal detector pointings r and θ (see Section 4.2.2), the model is:

$$\begin{aligned} x_0 &= d \tan(r) \sin(\theta + \phi) + \Delta x \\ y_0 &= d \tan(r) \cos(\theta + \phi) + \Delta y \end{aligned} \tag{7.1}$$

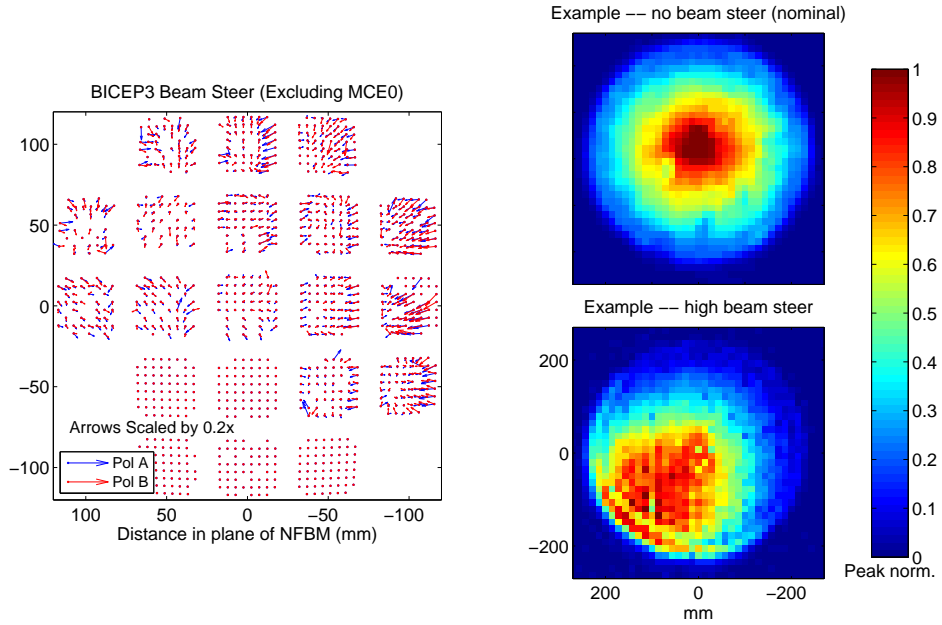


Figure 7.1: *Left:* quiver plot of the beam steer measured in BICEP3 near-field beam maps taken in January 2017. The arrows point from the aperture center (as seen in that detector’s beam map) to the fit beam center. The arrows are in units of mm, but scaled to a factor 0.35 of the axis scaling. *Right:* sample NFBMs of a nominal BICEP3 detector (top) and a detector with very high beam steer (bottom). The nominal beam steer is small, but there are a small handful of detectors with exceptionally high beam steer (for reference, the average beamwidth in the NFBM plane is ~ 170 mm). These detectors are removed from CMB analysis via auxiliary “channel” cuts, but still provide valuable points of comparison for this and other metrics discussed in this chapter. MCEO was disabled during this data set, for unrelated testing.

where (x_0, y_0) is the center of the aperture in the raw NFBM for a given detector, d is the vertical separation between the aperture plane and the NFBM plane, ϕ is the rotation of the NFBM instrument with respect to the instrument at $\theta = 0$, and $(\Delta x, \Delta y)$ is the offset between the center of the NFBM and the center of the aperture.

Once the aperture center has been located, a 2D Gaussian can be fit to the near-field beam to extract the beam center, and the beam steer metric can be evaluated

as the 2D distance between the fit beam center and the aperture center. The beam steer for BICEP3 detectors evaluated from this procedure is shown in Fig. 7.1, as well as sample near-field beams of nominal and high-beam steer detectors. Most detectors have small beam steer ($\lesssim 40$ mm) compared to the average beamwidth (~ 170 mm) in the NFBM plane. There is a small subset of detectors with abnormally high beam steer, mostly near the tile edge and near the focal plane edge. Most, but not all, of the detectors in this subset are excluded from CMB analysis, via auxiliary cuts (“channel flags”)¹. This subset will be revisited often throughout this chapter, when comparing the metric values of nominal detectors vs outlier detectors.

The next, perhaps more widely applicable NFBM metric is edge taper. To minimize sidelobes and extended response in the far-field beams, most receivers are designed such that the edge taper at the aperture is very far down from the main beam peak. For example, all BICEP/*Keck Array* receivers have target edge taper values between -8 and -15 dB_{peak}. To evaluate the edge taper from the near-field maps, we apply the same algorithm to identify the center of the aperture in a given beam map, and then define an annulus with an outer diameter equal to the physical edge of the aperture and an inner diameter equal to 90% the physical aperture diameter. The width of this annulus is somewhat arbitrary, but was chosen to allow enough map pixels to give a confident estimate of the edge taper over the noise, without making the

¹The channel flag that removes some of these outlier detectors is `pos_err`, which cuts on instability in the beam center analysis done by cross-correlating *T* maps with *Planck*.




Figure 7.2: Edge-taper metrics evaluated on real BICEP3 near-field beam maps. From left to right: the minimum value (5th percentile), median value, and maximum value (95th percentile) of the beam map within an annulus with outer radius equal to the aperture radius and inner radius equal to 90% the aperture radius. Each point is the average of the edge taper between the two detectors in that pair. MCE0 is not included in this NFBM data set due to unrelated testing. A subset of detectors in the middle tile of the right-most column (Tile 2) show very high maximum edge taper, and low median edge taper, due to the high amount of beam steer also seen in Fig. 7.1.

annulus thick enough to substantially over estimate the true edge taper. Given that most beams are not perfectly centered and circular, we use three separate numbers to fully quantify the extracted edge taper – the minimum, median, and maximum values within the annulus. To avoid outliers, we actually use the 5th percentile and 95th percentile as proxies for the minimum and maximum, respectively.

Tile plots of the three edge-taper metrics for BICEP3 are shown in Fig. 7.2. Most detectors show edge tapers varying from -12 dB to -5 dB, with an average edge taper of roughly -8 dB. Trends across the FPU are most clear in the maximum edge taper plot. There appears to be an overall trend with focal plane radius – the farther out on the focal plane, the higher the edge taper. This may indicate vignetting as far-edge detector beams clip the edge of, say, the field lens (we will see this effect occurring

in GRASP simulations presented in Section 7.2). There is also a slight variation in edge taper with tile radius, which may be due to beam steer caused by variation in the dielectric layer across the tile. Note the small subset of detectors in Tile 2 (right-most column, center row) which has very high max edge taper but low median edge taper. Those detector suffer from very high beam steer, such that the center of the beam is at or even surpassing the edge of the aperture stop. This causes a very large maximum taper (the beam peak is within the annulus) and low median taper (most map pixels do not have significant beam power). To see example BICEP3 NFBMs of detectors with varying levels of edge taper, see Fig. 4.11.

The left-hand plot of Fig. 7.3 shows how maximum edge taper and total beam steer vary with detector distance from boresight. As expected, detectors with very high beam steer lead to high edge taper – the beams are being steered into the edge of the aperture stop. However, the same plot indicates that high edge taper can also be generated without beam steer (as seen in the example in Fig. 4.11). At the far edge of the focal plane, almost all detectors see very high edge taper, likely due to vignetting within the optics. The small rise at a distance of $\sim 2^\circ$ from boresight is a bit more mysterious. Those points are all at the edge of Tile 11, which is at the center of the focal plane (all points within $< 2^\circ$ are Tile 11). Typically in detector fabrication, dielectric thickness across a tile can vary 1) linearly in a single direction across the tile and 2) in a circularly symmetric pattern, due to the symmetry of the “sputtering” procedure in building the detector film layers. This uptick at $\sim 2^\circ$ in the

left of Fig. 7.3 is likely sampling such a circular variation in dielectric thickness in the central tile.

During CMB observations, every receiver is coupled to an ambient temperature, absorptive, comoving forebaffle designed to intercept any portion of the beam $\gtrsim 15^\circ$ from the main beam. Although this prevents the extended beam from coupling to an unknown portion of the ground or sky, the forebaffle itself adds thermal loading to the detectors, depending on the fraction of the beam coupled. Every austral summer, we estimate this forebaffle loading by performing “noise stares” (collecting detector response with the telescope fixed at zenith) and load curves with and without the forebaffle attached to the receiver. With knowledge of the optical efficiency of the detectors, this allows us to extract the forebaffle loading onto the detectors in pW or K_{CMB} .

We may expect detectors that demonstrate high beam steer or edge taper to correlate with greater forebaffle loading – as the truncated portion of the beam exits the aperture, it is more likely to diffract more power into the far sidelobes. There are two trends to be aware of before looking for such a correlation: 1) We generally expect the forebaffle loading to increase with detector distance from boresight, due to edge detector beams exiting the aperture at a steeper angle and naturally coupling more power to the forebaffle. 2) We also know from the left-hand plot of Fig. 4.11 that detectors with very high edge taper are mostly at the edge of the focal plane, due to a different physical effect (possibly vignetting). Therefore, it may be difficult to identify

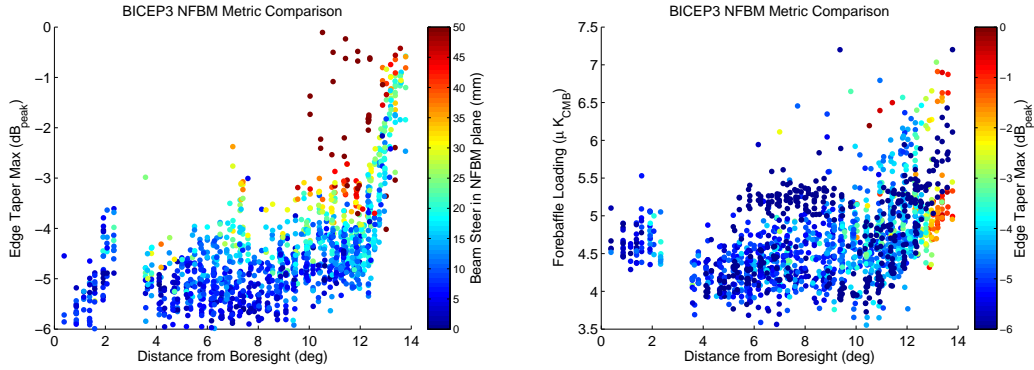


Figure 7.3: BICEP3 NFBM metric comparison plots. *Left:* NFBM maximum edge taper vs distance from boresight, with color scaled by beam steer magnitude. It is clear from this plot that high beam steer generally leads to high edge taper. However, the sharp rise beyond $\gtrsim 12^\circ$ from the boresight indicates a separate effect generating edge taper, which is likely vignetting of the optics. *Right:* forebaffle loading vs distance from boresight, with color scaled by maximum edge taper. Although both forebaffle loading and maximum edge taper generally increase with focal plane radius, there is no clear correlation between edge taper and forebaffle loading.

a true correlation between edge taper and forebaffle loading. The right-hand plot of Fig. 4.11 looks for possible correlation by plotting both forebaffle loading and edge taper against distance from boresight. There is no obvious trend – the detectors with higher edge taper do not seem more likely to have higher forebaffle loading.

This is a physically interesting result derived from this NFBM metric analysis – using two independent data sets in near-field beam maps and forebaffle loading measurements, it is not necessarily likely that a detector sharply truncated at the aperture stop will suffer excess loading from warm baffling. It is of course difficult to make a strong generalizing statement without considering more data sets, and without properly considering systematic uncertainties in the forebaffle loading measurement. Further work could be done to test these NFBM metrics on other existing data, e.g. from

Keck Array, and develop more confident estimates in the forebaffle loading by comparing multiple measurements on the same receiver, or by cross-checking against far-sidelobe maps taken with and without the forebaffle.

7.1.2 BEAM-MAP DERIVED LEAKAGE ESTIMATOR

The $T \rightarrow P$ leakage simulation technique described in detail in Chapter 6 is the proper method for quantifying the amount of leakage from main beam mismatch in a given data set. The full timestream simulations naturally include the complex averaging down and cancellation of leakage contributions from multiple detector pairs over multiple boresight rotation angles in each map pixel, after factoring in real data cuts and weights. However, there is also value in quantifying the intrinsic $T \rightarrow P$ leakage associated with a single unprojected residual beam. One can use a beam map-derived metric to get quick, simple leakage estimates on real or modeled beams where full timestream simulations are unavailable or inconvenient. [Karkare 2017](#) used this method with *Keck* 95, 150, and 220 GHz beams to examine scaling of residual beam leakage with beam size, and explore how the leakage averages down with detector count and with coaddition over many boresight rotation angles. In this section, we introduce the methodology, compare the leakage predictions against ρ evaluated from the full $T \rightarrow P$ leakage simulations, and explore correlations with other metrics described earlier in this chapter.

The procedure for evaluating the leakage metric closely follows [Karkare 2017](#). A

simulated realization of a Λ CDM sky is generated, and convolved with a given beam map, simulated or real, at its chosen level of deprojection. Since there are no timestreams in this case, deprojection is done in beam map space – a pair of beams is integral normalized and pair differenced, then regressed against the templates shown in Fig 6.1. As long as the beam templates are derived from beams with the window function $B(\ell)$ used to smooth the timestream deprojection template maps, beam map space deprojection and timestream deprojection are equivalent. The convolution between the simulated T sky and the difference beam is repeated with the beam rotated at the desired number of boresight rotation angles (at least two, ideally separated by 45° , are needed to separate the convolved maps into Q/U). The Q/U maps are coadded over all rotation angles, then rotated into E/B using Eq. 6.7. Power spectra are then binned into the same multipole bins used in standard BICEP/*Keck Array* analysis. From the BB power spectra obtained using this method, one can evaluate ρ in a very similar manner to what is described in Section 6.4.2. To give distinction between the two methods, we denote the leakage metric from this beam map-derived method as ρ_{bm} .

Due to the simplified nature of this procedure, it is agnostic of any real cuts and relative weights used in real data. When coadding over multiple boresight rotation angles, it assumes a perfectly even distribution of weight in each “observation.” In other words, when coadding over two boresight angles separated by 180° , leakage from the dipole-like component of the difference beam will *perfectly* cancel, which is unrealistic




Figure 7.4: Comparison of the leakage in the standalone 220 and 230 GHz bands as estimated from full $T \rightarrow P$ leakage simulations and from the beam map-derived procedure discussed in this section. *Left:* leakage BB from full simulations for both bands over the 2016 observing season. These results fully consider the effects of cuts, weights, and filtering. *Middle:* leakage BB for the standalone 220 GHz band derived from the undeprojected residual beams, coadded over eight boresight rotation angles. The faint lines are per-detector pair leakage spectra, and the bold line is the leakage derived from a single beam coadded over all detectors. *Right:* Same data as the middle plot, for the standalone 230 GHz band. The fractional 220/230 difference in ρ shown in the left plot is very similar to the fractional difference in ρ_{bm} in the middle and right plot.

in real data. Therefore the value of the ρ_{bm} metric is in relative comparisons, and not necessarily in its absolute value.

To give confidence in the trends and comparisons we can extract from the ρ_{bm} metric, we can validate relative predictions made on real *BICEP/Keck Array* beams against ρ values obtained from full $T \rightarrow P$ leakage simulations. One prediction made with the leakage simulations is that the BB leakage in the standalone *Keck* 220 GHz band is similar in shape but larger in amplitude to the standalone *Keck* 230 GHz BB spectrum². This is shown in the left-hand plot of Fig 7.4. From a full simulation of the 2016 observing season, the ratio of is $\rho(220)/\rho(230) = 2.3$. The middle and right

²Here the “standalone” 220/230 bands refer to these two bands separately, before being coadded. The results shown as 220 GHz in Chapter 6 are coadded over these two bands.

plots of Fig 7.4 show the expected leakage derived from the beam maps for the standalone 220 (middle) and 230 (right) GHz bands, coadded over eight boresight rotation angles to mimic the true *Keck* observing strategy. These include per-detector unde-projected residual leakage spectra, and the leakage predicted from the band-averaged residual beam. The amplitude of the beam map-derived leakage spectra is lower than the simulation-derived spectra at low ℓ , likely due to the beam map method assuming perfect cancellation of relevant modes over boresight rotation angles, which does not happen in the full timestream simulations. However, the ratio of leakage between the band-averaged residual beams is $\rho_{\text{bm}}(220)/\rho_{\text{bm}}(230) = 2.5$, which is very similar to the timestream simulation predicted ratio, which lends confidence to this method.

Next we evaluate the leakage metric ρ_{bm} on BICEP3 residual beam maps, and make comparisons to the edge-taper metric and forebaffle loading introduced in Section 7.1.1. The leakage metric is computed after coadding over four boresight rotation angles ($0^\circ, 45^\circ, 180^\circ, 225^\circ$), matching the real BICEP3 observation strategy. The results and comparisons are shown in Fig. 7.5. Starting with the tile plot on the left, there is no clear trend across the FPU or within a tile – pair-to-pair variation is the dominant effect. We isolate a subset of four neighboring detector pairs which all show elevated values of ρ_{bm} , and which have demonstrated high edge taper as seen in Fig. 7.2. These same pairs are marked in red in the scatter plots against edge taper and forebaffle loading³. Although there is no clear general correlation between ρ_{bm}

³As a reminder, these pairs are cut from the main CMB analysis due to external “channel”




Figure 7.5: Beam map-derived BB leakage ρ_{bm} for BICEP3 shown in tile plot format (left), and compared against NFBM maximum edge taper (top right) and forebaffle loading (bottom right). We isolate four neighboring detector pairs in the top-right corner of Tile 2 (highlighted in gray in the left figure, red dots in the top right and bottom right figures), which have demonstrated high leakage, edge taper, and forebaffle loading. Outside of this anomalous subset, there is no clear trend between ρ_{bm} and edge taper or forebaffle loading.




Figure 7.6: Beam map-derived EE (left) and BB (right) leakage for BICEP3 and *Keck* 95 GHz, averaged over all detectors on a tile perimeter (dashed lines) and near the tile center (solid lines). The average residual beam leakage is coadded over four boresight rotation angles for BICEP3 (0° , 45° , 180° , 225°) and eight evenly-spaced rotation angles for *Keck*, matching the real observation strategy. In both BICEP3 and *Keck* 95 GHz, the leakage is higher in EE than BB . The tile perimeter leakage dominates in both cases, more so for *Keck* 95 GHz. The shape and location of the “bump” in tile perimeter EE matches the same bump seen in the timestream simulation EE auto spectra shown in Fig 6.9.

and edge taper/forebaffle loading, this subset makes an interesting case study. The metric comparisons made throughout this chapter give slight evidence that detectors with largely anomalous aperture illumination – detectors with maximum edge taper or beam steer well above the nominal or design levels – could lead to critical levels of $T \rightarrow P$ leakage from main beam mismatch. Greater confidence could be made in such a claim if this correlation were seen in other receivers – future work may involve performing this analysis on BICEP Array rx0 (30/40 GHz), which, like BICEP3, shows a small subset of detectors with non-ideal aperture illumination.

Another utility of the beam map-derived leakage metric is the ability to check flags and therefore do not contribute to the bias Δr shown in Sec. 6.4.3.

for trends across specified detector subsets. We’ve already seen this method predict higher leakage in the standalone 220 GHz band compared to the 230 GHz band, agreeing with the full timestream simulations (Fig. 7.4). Another observation from the timestream simulations is that the BICEP3 leakage EE spectra is larger than the BB , especially near $\ell \sim 200$, similar to the behavior in *Keck* 95 GHz caused by anomalous coupling between tile edge detectors and tile corrugations (see Fig 6.9). We can evaluate the expected leakage in the average tile edge pixels and compare to the tile inner pixels, to verify if the high EE leakage in BICEP3 is also due to tile corrugations. This is shown in Fig. 7.6. The left-hand plot shows the EE leakage for tile perimeter and inner-tile detector subsets. The tile perimeter leakage dominates in this case for both bands, though the difference between inner and outer tile is smaller for BICEP3. The shape of the “bump” in tile perimeter EE in both bands matches the timestream simulation EE auto spectra in Fig 6.9, possibly indicating a small level of coupling to tile corrugations in BICEP3. The right-hand plot shows the BB leakage spectra, which is smaller than EE in all cases, as expected from the “+” shape of the residual beam patterns. Even in BB , the tile perimeter leakage is significantly higher than the inner tile leakage for *Keck* 95 GHz, which is consistent with previous work (The BICEP2 and *Keck Array Collaborations* 2019, Karkare 2017). The difference in tile perimeter and inner tile BB leakage for BICEP3 is relatively small.

Another application of the beam map-derived leakage method is to implement and test other modes of deprojection that we may consider adding to the timestream de-




Figure 7.7: Deprojection template corresponding to direct-island coupling. Unlike the standard deprojection templates, this one requires knowledge of the orientation of the tile with respect to the standard $\theta = 0$ axis. The spacing of the dipole corresponds to the relative spacing of the A and B bolometer; the offset from the center of the map corresponds to the physical distance between the center of antenna patch and the bolometer pair. The beamwidth of each dipole component is calculated analytically and is slightly less than the main beam beamwidth, since it assumes a uniform aperture illumination.

projection templates. The general philosophy is to use deprojection to remove physically motivated difference modes that we know are repeatable and exist in the real data. For instance, deprojection of crosstalk has been implemented and studied previously (Section 6.5). Another such physically motivated difference mode not removed by the standard suite of deprojection templates is direct-island coupling, which is especially prominent in the standalone 220, 230, and 270 GHz bands (see the bottom row of Fig. 5.13). We can easily implement deprojection of this mode in beam map space, using knowledge of the TES island placement relative to the pixel antenna patch. The beamwidth of these patterns can also be analytically calculated assuming a uniformly illuminated aperture at the band center frequency (this results in a smaller beamwidth than the main beam, which forms from an under-illuminated aper-




Figure 7.8: Beam map-derived EE (left) and BB (middle) leakage for the standalone *Keck* 220 GHz band with standard beam map-space deprojection (differential pointing, beamwidth, and ellipticity) and with direct-island coupling deprojection additionally included. The beams in the top right and bottom right are the band-averaged residual beams with standard and standard + direct-island coupling deprojection, respectively. The leakage spectra were made by coadding over eight boresight rotation angles uniformly spaced by 45° , matching the real *Keck* observation strategy. Including direct-island coupling deprojection reduces the predicted BB leakage by an order of magnitude in ρ_{bm} , with negligible change in EE (the blue and magenta lines overlap).

ture). The beam map-space template is therefore a dipole with fixed separation, at a fixed distance from the center of the beam map, as shown in Fig. 7.7. The only free parameter is the amplitude of the dipole. Note that this mode also requires knowledge of the orientation of the tile which houses a given detector, which determines which side of the main beam to expect the dipole. We include this extra deprojection mode, and compare it to the standard beam map space deprojection (differential pointing, beamwidth, and ellipticity) by deprojecting each beam, then coadding over the entire standalone 220 GHz band (which showed the brightest direct-island coupling features) over eight boresight rotation angles, and generate leakage EE and BB spectra.

The results of this new beam map-space deprojection mode are shown in Fig. 7.8. The band-averaged residual maps in the right-most column show a significant reduc-

tion in leakage power when direct-island coupling deprojection is included. This is reflected in the predicted leakage BB spectra in the middle plot – the ρ_{bm} metric decreases by an order of magnitude. The change in EE is negligible, which is expected for a band with roughly equal contributions from tiles relatively rotated at 0° and 180° (in other words, the “bottom” dipole and the “top” dipole in the top right of Fig. 7.8 have roughly equal power; removing such power leads to no significant change along the “+” direction corresponding to EE leakage). This drastic reduction in BB is exciting – since direct-island coupling likely exists in the real CMB data at the same amplitude as seen in the beam maps, implementing this deprojection mode in the real analysis may show real improvement in the BB leakage predicted in the full $T \rightarrow P$ leakage simulations, giving a straightforward path to reducing the bias on r .

7.2 GRASP SIMULATIONS

In this section, we discuss simulations made in GRASP, a physical optics simulation software made by TICRA⁴. These simulations were run by Dr. Paul Grimes, as part of a collaboration at the Center for Astrophysics | Harvard & Smithsonian between the Receiver Lab and the CMB Group. They are focused primarily on the BICEP3 instrument in order to validate the models against a currently-deployed receiver with a wealth of optical data. Here we briefly describe the details of the model itself, and focus more on the results and how they compare against near-field beam maps (NFBM)

⁴www.ticra.com/software/grasp

and far-field beam maps (FFBM) taken on BICEP3.

7.2.1 THE BICEP3 GRASP MODEL

An electric field distribution, representing the antenna feed pattern of a single detector at some specified location in the focal plane, emits radiation through an optical chain specified by the user. The field emitted from the detector pattern is used to calculate surface currents on the surface of the dielectric elements (lenses, filters, etc.) in the Physical Optics (PO) approximation⁵, which assumes that the surface current induced on a curved surface is the same as what would be induced on flat, infinite surface tangent to that point. These surface currents then induce a field which is propagated to the next optical surface in the chain, and this process repeats up to the final surface of the optical chain. The complex field can be evaluated in any desired plane within optical chain, or outside it. To make comparisons to real BICEP3 NFBM data, the simulated beam is evaluated in a plane just above the window of the cryostat, at the same location as the NFBM apparatus (no forebaffle is included in the simulated optical chain for NFBMs). To evaluate the far-field beam, the field contribution from the surface currents at the cryostat window and from those at the forebaffle are summed together.

This method of sequentially calculating induced surface currents does not naturally

⁵GRASP is generally used to simulate reflections off of reflectors, but modifications to the basic PO method allow for complex conductivities needed to simulated dielectric lenses.

include standing waves or multiple reflections within the optical chain. However, we can simulate the ghost “buddy” beam described in Section 5.5 by including a second series of PO calculations, where the surface currents induced on the 50 K filter are reflected and propagated backward through the optical chain, reflected off the focal plane, then forward through the optical chain again. The field of the buddy beam is then added incoherently⁶ with the field of the main beam. All GRASP simulations shown in this chapter include both the main beam and the buddy beam in this way.

GRASP implements Physical Theory of Diffraction (PTD), which modifies the standard PO assumption for surface currents near the edge of an optical surface. This PTD method has shown to be very accurate for optics set into metal surfaces (surfaces with perfect conductivity). However, PTD cannot be implemented on surfaces with complex conductivities, which means diffracted fields from absorptive surfaces and aperture stops in this model generally have overestimated amplitudes. This is a key point to keep in mind when comparing to real data from BICEP3, which uses both an absorptive aperture stop and an absorptive forebaffle.

GRASP evaluates the fields for a single frequency at a time, so in order to simulate a beam from a wideband incoherent bolometer, we must run simulations on a number of frequencies throughout the band and sum them incoherently. In the models presented here, we sum over five single frequencies uniformly spanning the BICEP3 band,

⁶The coherence length $L = cn/\Delta f \sim 10$ mm for BICEP3, which is much less than the difference in path length between the buddy beam and main beam. This is why they are summed incoherently.

which seems to be sufficient for predicting the level and shape of the near sidelobes. However, some features we will see below, such as deep nulls between the main beam and first sidelobe not seen in the real data, may indicate more frequencies and a more sophisticated treatment of the multi-frequency summation while accounting for coherence lengths may be necessary.

7.2.2 MODEL COMPARISON TO REAL DATA

We simulate linearly-polarized detector feed patterns at many locations in the focal plane. As a function of the total physical FPU radius $R_0 = 230$ mm, the simulated feeds discussed in this section are roughly located at $R = (0, 0.2R_0, 0.5R_0, 0.8R_0)$. Note that $0.8R_0$ corresponds to a detector at the edge of the focal plane. First, we compare the near-field beam patterns for feeds at $0.2R_0$ and $0.8R_0$ in Fig. 7.9. The GRASP beams are obtained by evaluating the field in an (x,y) plane at a distance above the aperture that matches the location of the NFBM apparatus during real NFBM measurements. For each feed position, the GRASP beam is compared to an average of ~ 10 real beams corresponding to working detectors near the simulated feed position. The agreement is very good, especially for feeds closer to the center of the focal plane. The detectors used in the comparison at the edge of the focal plane have significant beam steer in the real beam maps due to phase gradients in the beam pattern at the focal plane, leading to an illumination that is not centered on the aperture. The far edge GRASP simulation does not include beam steer by design, but

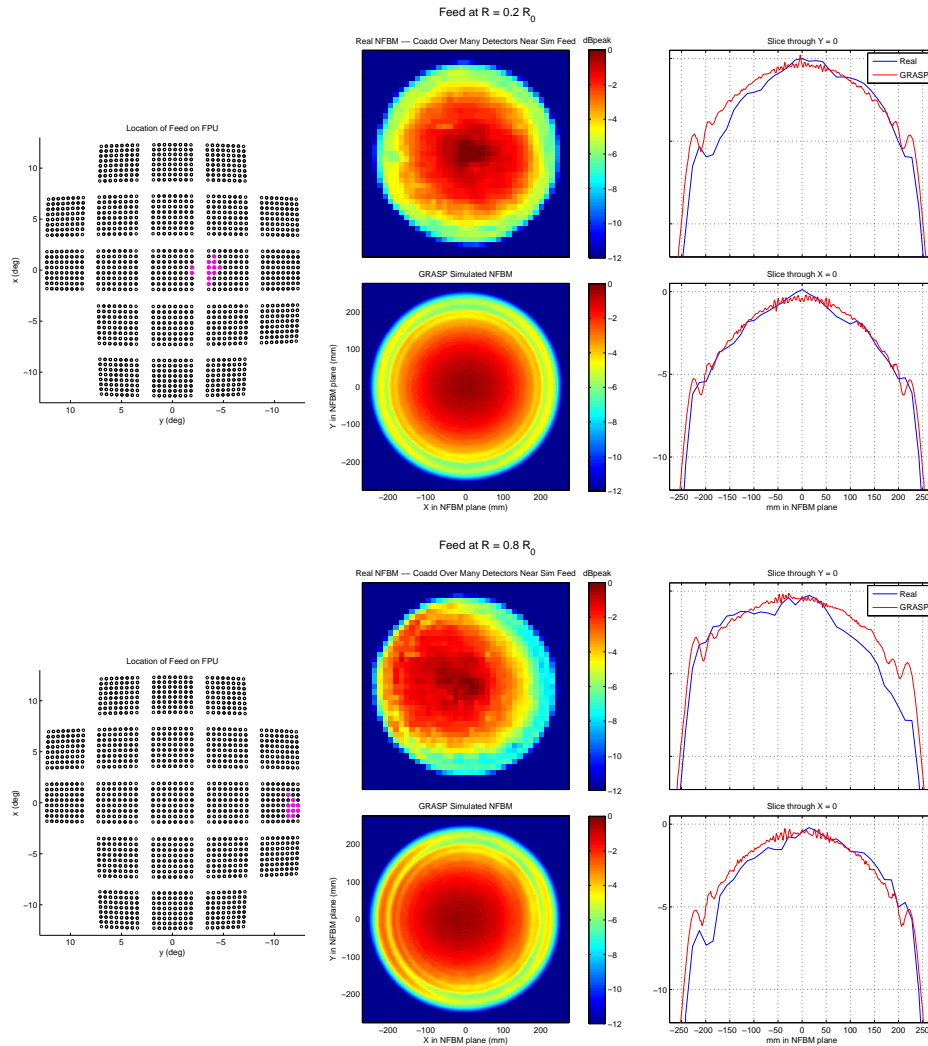


Figure 7.9: Near-field beam comparison between GRASP and real data, for two feed locations in BICEP3. The top (bottom) figure shows the comparison for a feed at radius $0.2R_0$ ($0.8R_0$), where R_0 is the physical FPU radius. Each GRASP model is compared to an average over ~ 10 real detector maps. In each figure, the FPU plot on the left highlights the location of the real detectors averaged over to create the real NFBM, chosen to be near the simulated feed location. The middle column shows the real and simulated NFBMs in dB relative to peak, and the right column shows cuts through $X = 0$ and $Y = 0$. The agreement is generally good, although beam steer in the real beams in the $R = 0.8R_0$ plot skew the $Y = 0$ response pattern more than what's predicted in GRASP. Separate GRASP simulations that include beam steer (by inducing a phase gradient in the feed pattern) have led to better agreement in that case.




Figure 7.10: GRASP simulated beam for a detector near the edge of the focal plane ($R = 0.8R_0$) evaluated at the plane of the field lens. The magenta circle indicates the outer edge of the field lens. One edge of the beam is truncated at about -12 dB from the peak, leading to $\sim 3\%$ of the integrated beam power missing the lens. No beam steer is included in this simulation.

still shows higher edge taper, possibly due to vignetting at the field lens. Specialized GRASP simulations with phase gradients in the detector feed pattern show improved agreement with real detectors that exhibit beam steer.

To see the extent of the truncation happening at the field lens to an edge detector beam, we plot the GRASP simulated beam evaluated at the plane just before the field lens in Fig. 7.10. In this simulation, beam steer is not included, so the truncation seen in the figure is representative of a nominal working detector at that radius on the focal plane. The main beam is truncated at about -12 dB from peak value, and the total integrated power missing the lens altogether is about 3%. This is a possible source of the high maximum edge taper seen in the GRASP NFBM at the bottom of Fig. 7.9. This could also explain some of the very high edge taper seen in the corresponding real NFBM at the bottom of Fig. 7.9, although in that case the edge taper

is also increased due to beam steer. An example of a NFBM with high edge taper but nearly zero beam steer can be seen in Fig. 4.11.

Next, we compare the GRASP simulations and real beams in the far field, again using a detector located at $R = 0.2 R_0$ and one at $R = 0.8 R_0$. Fig 7.11 shows the far-field beam comparison. The model matches the real data at the main beam peak, but below -10 dB from the peak, we see a discrepancy in the extended beam response. Although the sidelobes are at the right location, the lack of scattering in the optics and the inability to simulate diffraction off of an absorptive aperture stop (we are forced to use a conductive aperture stop in this model) leads to underestimated power beyond a radius of $\sim 0.5^\circ$. As alluded to earlier, the deep null between the main beam and first sidelobe in the $R = 0.2 R_0$ detector indicates that more frequencies, or a more complete treatment of the coherence length when adding these frequencies, may be needed. The far-field beam from the edge detector correctly predicts a radial ellipticity which is seen in the real data, but the magnitude of the ellipticity is overestimated, due to the same assumption of a conductive aperture stop. The model also predicts a small extended radial lobe a few degrees away from the main beam, which is seen at a similar level in some edge detectors.

What are the shape and amplitude of the differential beams predicted by this GRASP model, and what level of $T \rightarrow P$ leakage can we expect from the unprojected residuals? All results shown so far have been for a single linear polarization orientation, but simulations were also run with the orthogonal polarization at each detector location.

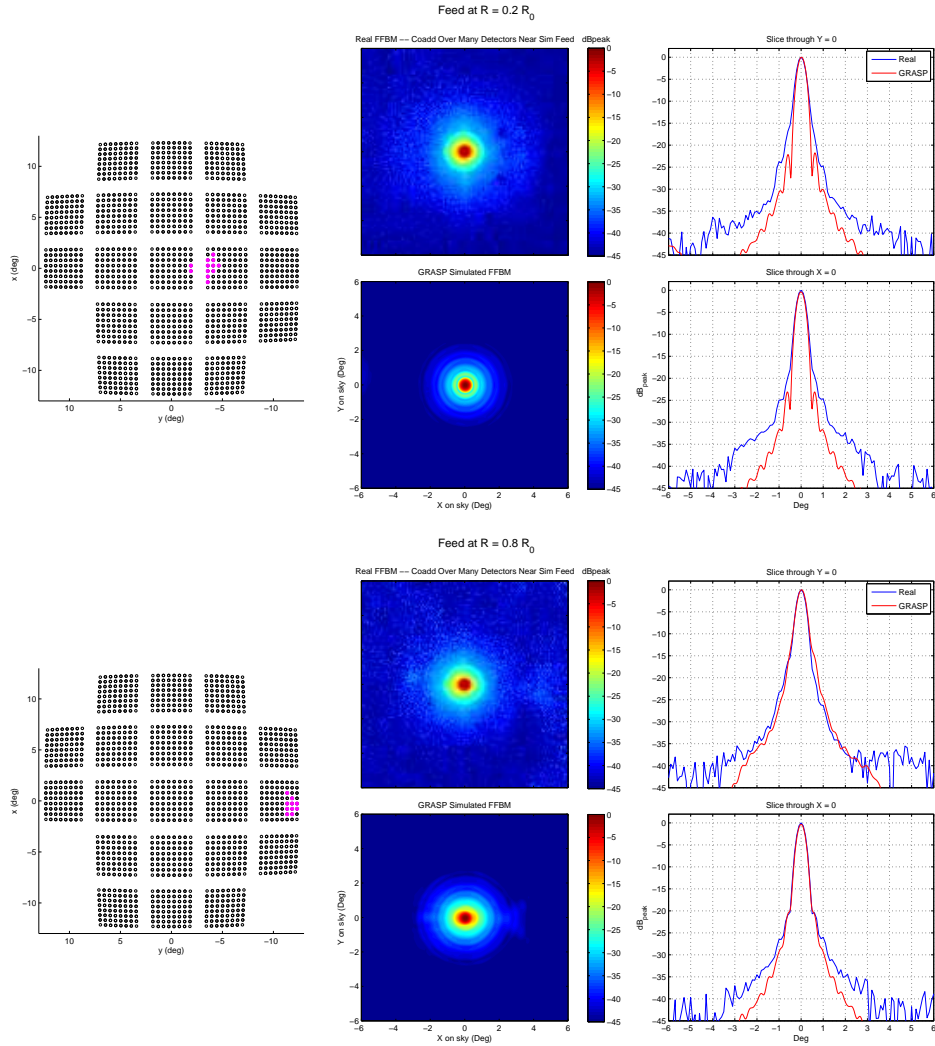


Figure 7.11: Far-field beam comparison between GRASP and real data, for two feed locations in BICEP3. The top (bottom) figure shows the comparison for a feed at radius $0.2R_0$ ($0.8R_0$), where R_0 is the physical FPU radius. Each GRASP model is compared to an average over ~ 10 real detector maps. In each figure, the FPU plot on the left highlights the location of the real detectors used to create the real averaged FFBM, chosen to be near the simulated feed location. The middle column shows the real and simulated FFBMs in dB relative to peak, and the right column shows cuts through $X = 0$ and $Y = 0$. The agreement in the main beam peak is good for the center ($R = 0.2R_0$) detector, but below -10 dB there is a clear discrepancy in the extended beam, primarily due to the lack of scattering and PTD off of absorbing aperture stops in the GRASP model (the model here uses a conductive aperture stop). For the edge ($R = 0.8R_0$) detector, the model overestimates the ellipticity seen in the far field, again due to the inability to implement PTD for an absorbing aperture stop. The extended lobe at $X \simeq +3^\circ$ in the model may arise from the truncation at the field lens or aperture stop – the same feature may be seen in the real map, though it is difficult to distinguish from noise.

We can take these pairs of modeled beams and difference them to look for the dominant difference modes, then deproject them (in beam map space), just as we do with the real beams. Finally, we can derive ρ_{bm} estimates using the same methodology described in Section 7.1.2. Although the model does not perfectly match real BICEP3 beams due to limitations in scattering and PTD, given the good agreement near the main beam peak, $T \rightarrow P$ leakage estimates from this model can still be quite informative. Given that the only inputs to the model are the antenna beam pattern at the focal plane and the elements and configuration of the optical chain, any residual beam pattern seen in the far field could in principle be isolated to a certain element by probing the beam pattern at various stages throughout the optical chain. Ultimately, if the model can predict undeprojected residual beam patterns seen in the real data, even to a rough extent, it would give great confidence in the modeling efforts currently underway for the development and refinement of future small-aperture telescope designs.

The difference beams, residual beams, and beam map-predicted BB leakage for GRASP modeled BICEP3 beams are shown in Fig. 7.12. The leakage spectra are evaluated after coadding over boresight rotation angles $0^\circ, 45^\circ, 180^\circ$, and 225° . All four detector pair locations are included in this comparison, ranging from the center of the focal plane to the far edge. The difference beam patterns are dominated by differential ellipticity, likely originating from the input antenna feed pattern – differential ellipticity is expected due to the asymmetric shape of the antennas themselves.

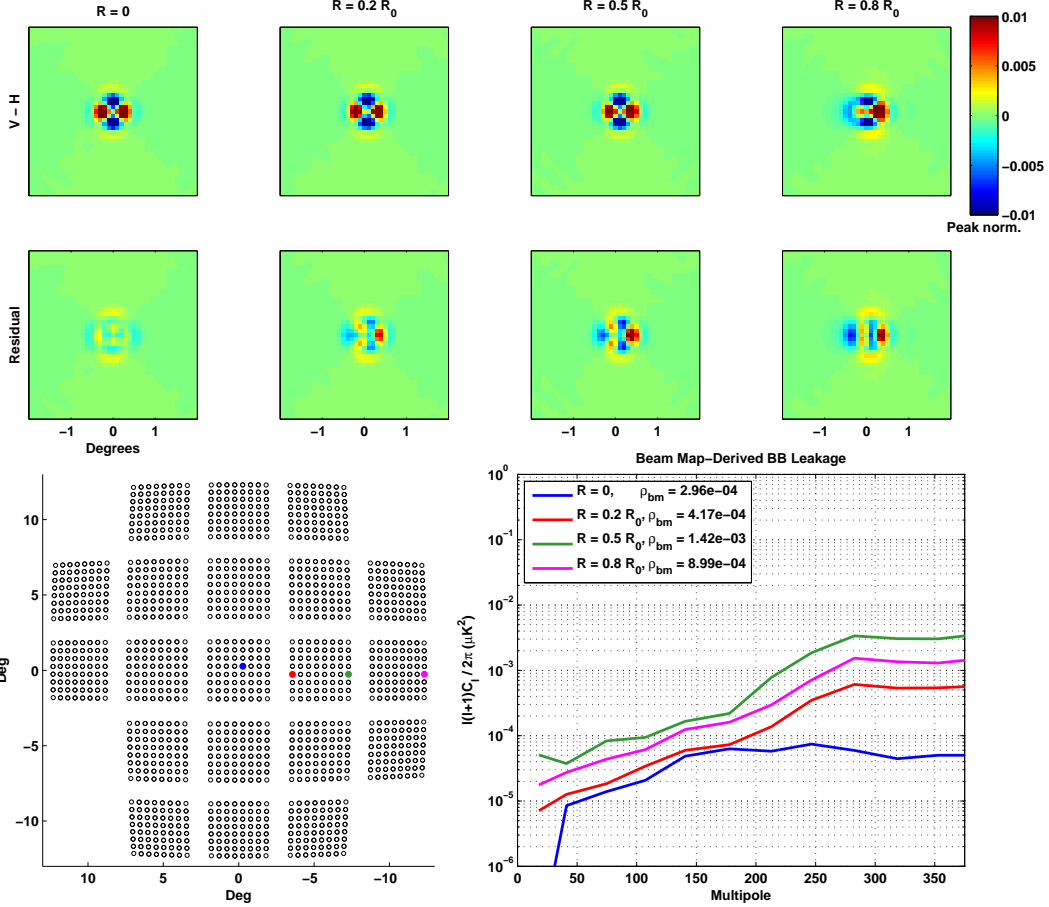


Figure 7.12: Difference beams (first row), undeprojected residual beams (second row), and beam map-predicted BB leakage (bottom right) for GRASP modeled BICEP3 beams. The four detector pair locations included in this comparison (0 , $0.2R_0$, $0.5R_0$, $0.8R_0$) are highlighted in the FPU projection plot in the bottom left. All difference beams are dominated by differential ellipticity, with increasing power in the undeprojected residuals as the detector pair location is farther from the center of the focal plane. The BB spectra plotted were estimated by coadding over rotation angles (0° , 45° , 180° , 225°), matching the real BICEP3 observation strategy. The BB spectrum for the farthest edge detector pair ($R = 0.8 R_0$) is slightly lower than that of the $R = 0.5 R_0$ pair, because the undeprojected residual incidentally has more of a “dipole-like” pattern that cancels with rotation over 180° .




Figure 7.13: GRASP-modeled residual beam for the center detector pair (left) and BICEP3 band-averaged real residual beam (middle). The beam map-derived leakage BB spectra are plotted on the right, when coadded over the minimum two boresight rotation angles (0° , 45°) and when coadded over four angles to match the real observation strategy (0° , 45° , 180° , 225°). The pattern is strikingly similar - note the quadrupolar structure at roughly 0.5° from the main beam in both maps. The GRASP modeled residual is more symmetric than the real averaged residual, leading to greater cancellation when averaging over more boresight rotation angles (going from the solid to the dashed lines in the BB spectra plot). The leakage metric for the GRASP and real beam is $\rho_{\text{bm}} = 0.0008$ and 0.003 respectively for two rotation angles, and $\rho_{\text{bm}} = 0.0003$ and 0.002 for four angles.

No differential pointing or gain is injected at the antenna feed pattern, and is not expected to arise from ideal optics. The undeprojected residual pattern gains more power for detector pairs farther from the center of the focal plane. Interestingly, the detector pair farthest from the center ($R = 0.8 R_0$) has a slightly lower BB spectrum than the $R = 0.5 R_0$ detector pair, due to a more “dipole-like” pattern which cancels more effectively under 180° rotation.

To facilitate a direct comparison with real BICEP3 residual beams, Fig. 7.13 shows the simulated center feed compared against the band-averaged BICEP3 residual beam, with associated beam map-derived leakage BB spectra. The leakage spectra shown are evaluated after coadding over the minimum two boresight rotations angles (0°

and 45°) and over four boresight rotation angles ($0^\circ, 45^\circ, 180^\circ, 225^\circ$). The overall shape of the residual beam is quite similar, particularly the “quadrupole-like” pattern 0.5° away from the beam center. The pattern is more symmetric in the model than in the average data, which may be expected since the model beam is evaluated at a single location at the center of the focal plane, and the average real beam has varying contribution from each location on the focal plane. It is also worth noting that the model beam is not *exactly* at the center – BICEP3 has no detector precisely at $(0,0)$, so the modeled detector is one of the four inner-most detectors in the real focal plane (see the blue point in the bottom left of Fig. 7.12). This slight offset from the exact center likely explains the small asymmetry in the modeled residual beam.

The ability of the GRASP model to predict a residual beam pattern that is repeatedly seen in real BICEP3 residual beams not only boosts confidence in the model itself, but also encourages us to consider adding a deprojection template (or two, as needed) corresponding to this quadrupole-like pattern, which is near the first null in the far-field beam. It is worth pointing out that despite the similarity in shape, the GRASP ρ_{bm} values are an order of magnitude lower than those predicted by the real data. Because of this difference (which is at least partially due to the higher degree of symmetry predicted in an ideal optical model like GRASP), we would always deproject this mode, not subtract it, to allow for fitting of the proper amplitude seen in real data. This is another path forward to reducing the systematic bias from $T \rightarrow P$ leakage in future BICEP/*Keck Array* data sets.

Looking forward to CMB-S4, the baseline detector coupling mechanism for the small-aperture telescopes (SATs) are feed horns, not the planar antennas modeled here for BICEP3. Substituting the antenna feed patterns with realistic horn illumination patterns will allow for modeling of SAT beams with different optical configurations (e.g. different lens materials). Comparing the predicted residual $T \rightarrow P$ leakage (among other optical constraints) from these configurations will be critical as these SAT designs are iterated and finalized.

7.3 HUYGENS-FRESNEL SIMULATIONS

Although optical simulation software such as GRASP or Zemax are well-suited for sophisticated optical analyses, such as those required to simulate the full optical chain of a CMB small-aperture telescope, the licenses are expensive and the simulations may require prohibitive computation times. In this section, we briefly introduce a simple optical simulation framework that takes in an aperture illumination pattern (e.g. from real near-field beam maps) and predicts the distribution in any plane in the near or far field by applying the Huygens-Fresnel principle.

The Huygen-Fresnel (HF) principle postulates that every point on a wavefront can be treated as a source of spherical waves, and the superposition of these spherical waves forms the wavefront. This idea was initially proposed by Christiaan Huygens ([Huygens 1690](#)), and was able to explain linear and spherical wave propagation

as well as simple reflection. However, it could not explain diffraction effects until Augustin-Jean Fresnel incorporated the principle of interference in the early 19th century (Crew 2009). A field characterized by wavelength λ has amplitude at some point in space $U(x, y, z)$ given by:

$$U(x, y, z) = \frac{1}{2i\lambda} \int_S U(x', y', 0) \frac{e^{i2\pi r/\lambda}}{r} (1 + \cos \theta) dx' dy' \quad (7.2)$$

where the illuminated surface (the circular telescope aperture in this case) is S , which is located in the plane $z = 0$, the 3D distance between the field point and the point on the surface is r , and $\cos \theta = z/r$. Fresnel needed to make empirical changes, which were physically unmotivated at the time, in order to obtain agreement between theory and experimental results. Similar results were eventually derived from first principles by Gustav Kirchhoff, using scalar diffraction theory. Kirchhoff's method uses Green's theorem to express a scalar field (such as a field amplitude) outside the surface S in terms of the values of the field and its first derivatives on S , assuming the field satisfies the Helmholtz wave equation. The resulting Kirchhoff integral is a bit more generalized than the HF principle, and both the Kirchhoff's method and the HF method fail when $\lambda \sim r$ or $\lambda > r$ (Jackson 1999). However, for the purpose of simulating diffraction of microwave wavelengths (~ 1 mm) from a $\gtrsim 26$ cm aperture, Eq. 7.2 is perfectly adequate.

The HF simulation framework takes some real near-field beam map as the input,

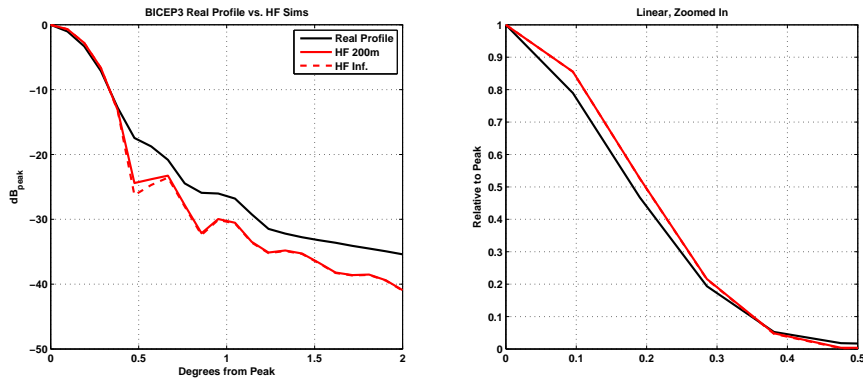


Figure 7.14: BICEP3 real beam profile (black) compared against HF simulations evaluated at a distance of 200 m (the true distance of the FFBM source; solid red line) and at a distance of infinity (dashed red line). The modeled beams are convolved with a uniform disk to emulate observing a 24" chopper (the chopper is not deconvolved from the real beam). The left plot shows the beam at a radius $< 2^\circ$ on a dB scale relative to peak, and the plot on the right shows a linear scale for the inner 0.5° . The agreement is generally good, though the model slightly overestimates the inner beam response and underestimates the extended response. The effect of observing a source at 200 m instead of infinity is very small.

assuming it to be the real illumination at the spectral band center at the receiver aperture. For a single frequency, the size of the input illumination is scaled from the nominal band center to the desired frequency, then Eq. 7.2 is evaluated at some specified distance z from the aperture. This is repeated for a set of discretized frequencies within the specified frequency band, then the resulting field distributions are coadded, weighted by the real measured spectral bandpass (such as those in Fig. 3.4). The results were verified against analytic far-field solutions of uniform and perfectly Gaussian aperture illuminations, as derived in, for example, [Goldsmith 1998](#).

We compare the predicted beam profile from the HF simulation to the real measured profile from FFBM measurements for BICEP3 in Fig. 7.14. The agreement in




Figure 7.15: HF simulated beam window functions $B(\ell)$ for BICEP3 (red) and BICEP Array 220 (blue) and 270 GHz (magenta). *Left:* the solid lines show the $B(\ell)$ simulated at a distance of 200 m, and the dashed lines show the $B(\ell)$ at infinity. *Right:* the $B(\ell)$ fractional difference (200 m - infinity) / infinity, which can be interpreted as an error in the measured $B(\ell)$ from not observing a source in the true far field. The fractional errors in the current BICEP/Keck Array science bins do not exceed 3%, even for future BICEP Array receivers.

the main beam is very good, though the size of the main beam is slightly overestimated (see the right-hand plot). Like the GRASP model, the HF model also underestimates the power in the sidelobes and extended beam response. The sidelobes occur at the expected location, and are smoothed due to the integration over a frequency spectrum with finite bandwidth.

The application we consider here is the quantitative impact on the overall beam shape from measuring our beams using a source at a distance that is not quite in the far field. The far-field criterion $2D^2/\lambda$ varies from 44-171 m for past or currently-deployed BICEP/Keck Array receivers (see Table 4.1), which is all within the true distance of the FFBM thermal sources (200 m). The difference in predicted beam shape between 200 m and infinity are very small for BICEP3, as seen in Fig. 7.14. As Br-

CEP Array expands to 150, 95, and 220/270 GHz, the far-field criterion increases to as high as 545 m, which may cause concern for our standard FFBM observation procedure (though it is important to remember that $2D^2/\lambda$ is not a hard limit). Therefore we simulate the BICEP Array 220 and 270 GHz beams at 200 m and at infinity, and compare the fractional difference to quantify the error in the real $B(\ell)$ that would be measured with these receivers. In lieu of real near-field beam maps for input aperture illuminations, we use Zemax simulations of the beam at the aperture. The bandpasses used are the nominal *Keck* 220 and 270 GHz bandpasses. The results are shown in Fig. 7.15. In the current BICEP/*Keck Array* science bandpowers ($\ell \lesssim 350$), the fractional error in the $B(\ell)$ is less than 3%, though at much higher multipoles, the error gets as high as 13%.

A phase gradient can be applied to the aperture illumination to simulate the effect of, for example, curvature in the mirror used to redirect the beams to the thermal source. The photogrammetry measurements shown in Fig. 4.2 are used to obtain the phase gradient applied across a beam at the spot at which the beam intersects the mirror. HF simulations showed the fractional change in the $B(\ell)$ to be less than 1% at $\ell = 200$ for all BICEP/*Keck Array* bands.

There is room for future improvement of this simulation framework. Although the results were verified against analytic solutions and real data in the far field, comparison with existing data *not* in the far field could give more confidence in the results shown in Fig. 7.15. One such comparison could be with far-sidelobe maps (Sec-

tion 4.4) which are measured using a source only ~ 10 m away. The code as written also azimuthally symmetrizes the input aperture illumination pattern, in order to drastically reduce the computation time. One could optimize this code to work with a general asymmetric 2D aperture illumination, such as a real near-field beam map with high edge taper or beam steer, and compare the results against GRASP simulations or real FFBMs.

*May it be a light to you in dark places, when all
other lights go out.*

J.R.R. Tolkien, *The Fellowship of the Ring*

8

Conclusions

IN THE QUEST TO CONSTRAIN THE TENSOR-TO-SCALAR RATIO r and other cosmological parameters, CMB telescopes continue to improve in detector count and mapping speed with each generation. As the statistical uncertainties continue to plummet, the systematic uncertainties associated with the measurement must follow suit. All sys-

tematics, both those known to be significant and those previously considered to be subdominant, must be brought under scrutiny if we hope to sustain the impressive progress seen in the field the past two decades. The BICEP/*Keck Array* program has just published a sensitivity to r of $\sigma(r) = 0.009$, which will only continue to improve as the next three high-throughput BICEP Array receivers at 95, 150, and 220/270 GHz are developed and deployed. Looking even further ahead, CMB-S4, which will deploy $\sim 500,000$ detectors with frequency bands spanning from 20 GHz to 270 GHz, is forecasted to achieve a sensitivity of $\sigma(r) = 5 \times 10^{-4}$ (The CMB-S4 Collaboration 2020). Needless to say, systematic control must be a top priority for these experiments.

One such systematic considered in detail in this dissertation is uncertainty in the spectral bandpass measurements made with a field-deployable Fourier Transform Spectrometer. Accurate estimates of each detector’s spectral response is critical in foreground component separation, and systematics in the FTS measurement may lead to distortions or shifts in the measured spectra. After detailed analysis of the real data and simulations of frequency-dependent truncation effects, it was shown that the natural detector-to-detector scatter dominates all other systematic uncertainties. Using a conservative 2% band center uncertainty with the multicomponent likelihood framework leads to a maximum bias on r of $\Delta r = 8.4 \pm 5.0 \times 10^{-4}$ in the BK18 data set. In future experiments, improved uniformity in the optics and detector fabrication may be required to keep this uncertainty comfortably under the experiment statistical uncertainty. We have also seen that failure of some optical elements after a receiver

cooldown may cause degraded performance, so FTS measurements should be taken after cooldown (before observations) and before reopening the receiver (after observations) whenever possible.

The other systematic we extensively explored is temperature-to-polarization leakage due to main beam mismatch between co-located, orthogonally polarized detectors within a pair. Using per-detector composite beam maps made from *in-situ* far-field beam measurements in conjunction with specialized $T \rightarrow P$ leakage simulations, the best estimate of the bias on r in the BK18 data set was found to be $\Delta r = 0.0015 \pm 0.0011$. These full timestream simulations are the most reliable method for estimating the leakage in any real data set, because it accounts for real beam shapes and non-idealities along with the same weights, cuts, and filtering applied to the real data. Perhaps the most important message one should take from this dissertation is that direct, high-fidelity measurements of the far-field beams are crucial. The source should be bright and as stable as possible – the 24 inch thermal chopper used by BICEP/Keck Array runs nearly continuously every season for 4+ weeks, allowing us to build a suite of composite far-field beam maps necessary to constrain the expected $T \rightarrow P$ leakage.

While ongoing and future experiments such as BICEP Array and CMB-S4 are continuing to refine their baseline optical designs, the development of beam models and simple, beam map-derived metrics has helped facilitate comparisons between the optical designs and previous real, achieved performance. We have shown that beam mod-

els of small-aperture telescopes made in GRASP can reliably predict the near-field beam shape, far-field main beam shape, and most importantly, the shape of the unde-projected residual beam pattern. The success of this model lends great confidence in the optical design parameters of future receivers that are informed by GRASP simulations. It is critical to note, however, that beam models such as these are not suitable replacements for real beams. The beam map-derived $T \rightarrow P$ leakage metric has showed us that 1) although the GRASP model accurately predicts the residual beam shape, it significantly underestimates the amplitude and overall complexity, and 2) real beam measurements encapsulate other imperfections that may lead to $T \rightarrow P$ leakage that are difficult to model or do not arise from optics, such as crosstalk beams and direct-island coupling. While iterating on GRASP, Zemax, and other optical models, it is critical to benchmark their performance against real beam measurements.

In the short-term future, we consider ways to reduce the predicted $T \rightarrow P$ leakage in currently-deployed and future BICEP/*Keck Array* receivers, and improve the agreement between predicted leakage BB power and real data BB power. The $T \rightarrow P$ leakage analysis shown in this dissertation has already demonstrated a few paths forward. The first is implementation of deprojection of direct-island coupling, which is relatively bright in the unde-projected residuals of the high-frequency bands in particular. The second is deprojection of crosstalk beams, particularly in the $T \rightarrow P$ leakage simulations themselves, since the beam maps likely overestimate the actual level

of crosstalk seen in the real data. The third is the deprojection of the “quadrupole-like” feature seen in BICEP3 GRASP models and real data. Further investigation into the root cause of this pattern would lend more confidence in the choice to implement deprojection in other receivers as well. There are other far-field beam data sets available (far-sidelobe maps, mid-sidelobe maps made with the high-powered noise source) that can be used in conjunction with the standard thermal chopper far-field beam maps in order to develop complete per-detector beam maps that have high signal-to-noise out to $\sim 80^\circ$ from the main beam. This would give us a complete picture of the beam shape, amplitude and polarization level of the extended sidelobes, and expected $T \rightarrow P$ leakage in every detector. Overall, BICEP/*Keck Array* has a great advantage in its exhaustive library of beam measurements, and it is imperative to extract as much information from this library as possible.

References

- Abitbol, M. H., Ahmed, Z., & et. al. (2017). CMB-S4 Technology Book, First Edition. *arXiv e-prints*, (pp. arXiv:1706.02464).
- Ade, P. A. R., Pisano, G., Tucker, C., & Weaver, S. (2006). A review of metal mesh filters. In J. Zmuidzinas, W. S. Holland, S. Withington, & W. D. Duncan (Eds.), *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III*, volume 6275 (pp. 248 – 262).: International Society for Optics and Photonics SPIE.
- Aikin, R. (2013). *Testing inflationary cosmology with the BICEP1 and BICEP2 experiments*. PhD thesis.
- BICEP/Keck Array and SPIDER Collaborations (2015). Antenna-coupled TES bolometers used in BICEP2, Keck Array, and SPIDER. *Astrophys. J.*, 812(2).
- Barkats, D., Aikin, R., Bischoff, C., & et. al. (2014). Degree-scale Cosmic Microwave Background Polarization Measurements from Three Years of BICEP1 Data. *Astrophys. J.*, 783, 67.
- Baumann, D. (2009). TASI Lectures on Inflation. *arXiv e-prints*, (pp. arXiv:0907.5424).
- Bennett, C. L., Larson, D., & et. al. (2013). Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. 208(2), 20.
- BICEP/Keck Collaboration & SPTpol Collaboration (2021). A demonstration of improved constraints on primordial gravitational waves with delensing. *Phys. Rev. D*, 103(2), 022004.
- Bradford, J. K. (2012). *A Rotating Polarized Source: Instrument Development for Precision Calibration of Cosmic Microwave Background Polarimeters at the South Pole*. PhD thesis, Harvard University.

- Bunn, E. F., Zaldarriaga, M., & et. al. (2003). E/B decomposition of finite pixelized CMB maps. *Phys. Rev. D*, 67(2).
- Buza, V. (2019). *Constraining Primordial Gravitational Waves Using Present and Future CMB Experiments*. PhD thesis, Harvard University.
- Carlstrom, J., Abazajian, K., & et. al. (2019). CMB-S4. In *Bulletin of the American Astronomical Society*, volume 51 (pp. 209).
- Cornelison, J., Ade, P. A. R., Ahmed, Z., & et. al. (2020). Polarization calibration of the BICEP3 cmb polarimeter at the south pole. In *Proc.SPIE*, volume 11453.
- Crew, H. (2009). *The Wave Theory of Light - Memoirs by Huygens, Young and Fresnel*. Chapman Press.
- Cukierman, A., Ahmed, Z., & et. al. (2020). Microwave Multiplexing on the Keck Array. *Journal of Low Temperature Physics*, 199(3-4), 858–866.
- de Korte, P. A. J., Beyer, J., & et. al. (2003). Time-division superconducting quantum interference device multiplexer for transition-edge sensors. *Review of Scientific Instruments*, 74(8), 3807–3815.
- Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. (1965). Cosmic black-body radiation. *Astrophys. J.*, 142, 414–419.
- Duband, L. & Collaudin, B. (1999). Sorption coolers development at cea-sbt. *Cryogenics*, 39(8), 659–663.
- Einstein, A. (1916). The foundation of the general theory of relativity. *Annalen der Physik*, 354(7), 769–822.
- Fixsen, D. J. (2009). *The Temperature of the Cosmic Microwave Background*, volume 707. American Astronomical Society.
- Fuskeland, U., Wehus, I. K., Eriksen, H. K., & Næss, S. K. (2014). Spatial variations in the spectral index of polarized synchrotron emission in the 9 yr wmap sky maps. 790(2), 104.
- Goldsmith, P. F. (1998). *Quasioptical Systems: Gaussian Beam Quasioptical Propagation and Applications*. Wiley-IEEE Press.

- Gorski, K. M., Hivon, E., & et. al. (2005). Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *Astrophys. J.*, 622(2), 759–771.
- Grayson, J. (2016). *The BICEP3 Millimeter-Wave Polarimeter: Measuring the Cosmic Microwave Background from the South Pole*. PhD thesis, Stanford University.
- Grayson, J. A., Ade, P. A. R., & et. al. (2016). BICEP3 performance overview and planned *Keck Array* upgrade. In W. S. Holland & J. Zmuidzinas (Eds.), *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VIII*, volume 9914 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (pp. 99140S).
- Guth, A. H. (1981). Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23, 347–356.
- Hamimeche, S. & Lewis, A. (2008). Likelihood analysis of cmb temperature and polarization power spectra. *Physical Review D*, 77(10), 103013–.
- Hanson, D., Hoover, S., & et. al. (2013). Detection of *b*-mode polarization in the cosmic microwave background with data from the south pole telescope. *Phys. Rev. Lett.*, 111, 141301.
- Hivon, E., Gorski, K. M., & et. al. (2002). The american astronomical society, find out more the institute of physics, find out more master of the cosmic microwave background anisotropy power spectrum: A fast method for statistical analysis of large and complex cosmic microwave background data sets. *Astrophys. J.*, 567(1), 2–17.
- Hu, W. & Dodelson, S. (2002). Cosmic Microwave Background Anisotropies. *Ann. Rev. Astron. Astroph.*, 40, 171–216.
- Hubble, E. (1929). A relation between distance and radial velocity among extragalactic nebulae. *Proceedings of the National Academy of Sciences*, 15(3), 168.
- Huygens, C. (1690). *Treatise on Light*. chez Pierre Vander Aa marchand libraire.
- Irwin, K. D. (1995). An application of electrothermal feedback for high resolution cryogenic particle detection. *Applied Physics Letters*, 66(15), 1998–2000.

- Irwin, K. D., Cho, H. M., & et. al. (2012). Advanced code-division multiplexers for superconducting detector arrays. *Journal of Low Temperature Physics*, 167(5), 588–594.
- Jackson, J. D. (1999). *Classical Electrodynamics*. Wiley, third edition.
- Kang, J. H., Ade, P. A. R., & et. al. (2018). 2017 upgrade and performance of BICEP3: a 95GHz refracting telescope for degree-scale CMB polarization. In J. Zmuidzinas & J.-R. Gao (Eds.), *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, volume 10708 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (pp. 107082N).
- Karkare, K. S. (2017). *Multifrequency Beam Systematics and Measurements of B-Mode Polarization with the BICEP/Keck Array Cosmic Microwave Background Experiments*. PhD thesis, Harvard University.
- Kernasovskiy, S. A. (2014). *Measuring the Polarization of the Cosmic Microwave Background with the Keck Array and BICEP2*. PhD thesis, Stanford University.
- Kovac, J. M. (2003). *Detection of Polarization in the Cosmic Microwave Background using DASI*. PhD thesis, University of Chicago.
- Kovac, J. M. & Barkats, D. (2007). CMB from the South Pole: Past, Present, and Future. *arXiv e-prints*, (pp. arXiv:0707.1075).
- Kovac, J. M., Leitch, E. M., & et. al. (2002). Detection of polarization in the cosmic microwave background using DASI. *Nature*, 420(6917), 772–787.
- Kuo, C. L., Bock, J. J., & et. al. (2008). Antenna-coupled TES bolometer arrays for CMB polarimetry. In W. D. Duncan, W. S. Holland, S. Withington, & J. Zmuidzinas (Eds.), *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy IV*, volume 7020 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (pp. 70201I).
- Lesurf, J. C. G. (1990). *Millimetre-wave Optics, Devices, and Systems*. Bristol, England: A. Hilger.
- Lewis, A. & Bridle, S. (2002). Cosmological parameters from cmb and other data: A monte carlo approach. *Physical Review D*, 66(10), 103511–.

- Lewis, A., Challinor, A., & Lasenby, A. (2000). Efficient computation of cosmic microwave background anisotropies in closed friedmann-robertson-walker models. *The Astrophysical Journal*, 538(2), 473–476.
- Liddle, A. R. & Lyth, D. H. (2000). *Cosmological Inflation and Large-Scale Structure*. Cambridge University Press.
- Linde, A. (1982). A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6), 389–393.
- Martin, D. H. & Puplett, E. (1970). Polarised interferometric spectrometry for the millimetre and submillimetre spectrum. *Infrared Physics*, 10(2), 105–109.
- Moncelsi, L., Ade, P. A. R., Ahmed, Z., & et. al. (2020). Receiver development for BICEP Array a next-generation cmb polarimeter at the south pole. In *Proc.SPIE*, volume 11453.
- O’Brient, R., Ade, P. A. R., & et. al. (2012). Antenna-coupled tes bolometers for the *Keck Array* spider, and polar-1. In *Proc.SPIE*, volume 8452.
- Paine, S. (2019). The am atmospheric model.
- Penzias, A. A. & Wilson, R. W. (1965). A measurement of excess antenna temperature at 4080 mc/s. *The Astrophysical Journal*, 142, 419–421.
- Polarbear Collaboration (2014). A measurement of the cosmic microwave background *b*-mode polarization power spectrum at sub-degree scales with polarbear. *Astrophys. J.*, 794(2), 171.
- Rees, M. J. (1968). Polarization and spectrum of the primeval radiation in an anisotropic universe. *Astrophys. J.*
- Sheehy, C. (2013). *Progress toward detection of inflationary B-modes with the BICEP2 and Keck Array polarimeters*. PhD thesis.
- Smith, K. M. & Zaldarriaga, M. (2007). General solution to the *e*–*b* mixing problem. *Phys. Rev. D*, 76(4).

St Germaine, T., Ade, P. A. R., Ahmed, Z., & et. al. (2020). Optical characterization of the *Keck Array* and BICEP3 cmb polarimeters from 2016 to 2019. *Journal of Low Temperature Physics*, 199(3), 824–832.

Starobinsky, A. (1980). A new type of isotropic cosmological models without singularity. *Physics Letters B*, 91(1), 99–102.

Stiehl, G. M., Cho, H. M., & et. al. (2011). Time-division squid multiplexers with reduced sensitivity to external magnetic fields. *IEEE Transactions on Applied Superconductivity*, 21(3), 298–301.

The BICEP2 and *Keck Array* Collaborations (2015a). BICEP/*Keck Array* IV: Optical Characterization and Performance of the BICEP2 and *Keck Array* Experiments. *Astrophys. J.*, 806, 206.

The BICEP2 and *Keck Array* Collaborations (2015b). BICEP/*Keck Array* V: Measurements of B-mode Polarization at Degree Angular Scales and 150 GHz by the *Keck Array*. *Astrophys. J.*, (pp. 126).

The BICEP2 and *Keck Array* Collaborations (2016a). BICEP/*Keck Array* VII: Matrix based *E/B* Separation applied to BICEP2 and the *Keck Array*. *ArXiv e-prints*.

The BICEP2 and *Keck Array* Collaborations (2016b). BICEP/*Keck Array* VI: Improved Constraints on Cosmology and Foregrounds from BICEP2 and *Keck Array* Cosmic Microwave Background Data with Inclusion of 95 GHz Band. *Physical Review Letters*, 116(3), 031302.

The BICEP2 and *Keck Array* Collaborations (2017). BICEP/*Keck Array* IX: New bounds on anisotropies of CMB polarization rotation and implications for axionlike particles and primordial magnetic fields. *Phys. Rev. D*, 96(10), 102003.

The BICEP2 and *Keck Array* Collaborations (2018). BICEP/*Keck Array* x: Constraints on primordial gravitational waves using *Planck* wmap, and new BICEP/*Keck Array* observations through the 2015 season. *Physical Review Letters*, 121(22), 221301.

The BICEP2 and *Keck Array* Collaborations (2019). BICEP/*Keck Array* XI: Beam Characterization and Temperature-to-Polarization Leakage in the BK15 Data Set. *Astrophys. J.*, 884(2), 114.

- The BICEP2 and *Keck Array* Collaborations (2021). BICEP/*Keck Array* XIII: Constraints on Primordial Gravitational Waves using *Planck*, WMAP, and BICEP/*Keck* Observations through the 2018 Observing Season. In Preparation.
- The BICEP2 Collaboration (2014a). BICEP2 I: Detection of B-Mode Polarization at Degree Angular Scales by BICEP2. *Phys. Rev. Lett.*, 112(24), 241101.
- The BICEP2 Collaboration (2014b). BICEP2 II: Experiment and three-year Data Set. *Astrophys. J.*, 792, 62.
- The BICEP2 Collaboration (2015). BICEP2 III: Instrumental Systematics.
- The CMB-S4 Collaboration (2020). CMB-S4: Forecasting Constraints on Primordial Gravitational Waves. *arXiv e-prints*, (pp. arXiv:2008.12619).
- The *Planck* Collaboration (2014). Planck 2013 results. XXII. Constraints on inflation. *Astr. & Astroph.*, 571, A22.
- The *Planck* Collaboration (2020a). Planck 2018 results. I. Overview and the cosmological legacy of Planck. *Astr. & Astroph.*, 641, A1.
- The *Planck* Collaboration (2020b). Planck 2018 results. IX. Constraints on primordial non-Gaussianity. *Astr. & Astroph.*, 641, A9.
- The *Planck* Collaboration (2020c). Planck 2018 results. VI. Cosmological parameters. *A&A*, 641.
- The *Planck* Collaboration (2020d). Planck 2018 results. X. Constraints on inflation. *Astr. & Astroph.*, 641, A10.
- The *Planck* Collaboration (2020e). Planck intermediate results. LVII. Joint Planck LFI and HFI data processing. *A&A*, 643.
- Verde, L., Treu, T., & Riess, A. G. (2019). Tensions between the early and late universe. *Nature Astronomy*, 3(10), 891–895.
- Willmert, J. (2019). *Constraining Inflationary B-Modes with the BICEP/Keck Array Telescopes*. PhD thesis, University of Minnesota.
- Wong, C. L. (2014). *Beam Characterization and Systematics of the BICEP2 and Keck Array Cosmic Microwave Background Polarization Experiments*. PhD thesis.

Yoon, K. W., Ade, P. A. R., Barkats, D., & et. al. (2006). The Robinson Gravitational Wave Background Telescope (BICEP): a bolometric large angular scale CMB polarimeter. In *Proc.SPIE*, volume 6275 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*.

Zaldarriaga, M. & Seljak, U. (1997). All-sky analysis of polarization in the microwave background. *Phys. Rev. D*, 55(4), 1830–1840.