

Searching for Primordial Gravitational Waves at Degree Scales from the South Pole

Thesis by

Justus A. Brevik

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy



California Institute of Technology

Pasadena, California

2012

(Defended May 25, 2012)

© 2012

Justus A. Brevik

All Rights Reserved

To my parents

Abstract

We report on the preliminary performance of the BICEP2 mm-wave polarimeter, which was deployed in 2009 to the South Pole and will observe through 2012. BICEP2 is currently imaging the polarization of the cosmic microwave background at 150 GHz using an array of 512 antenna-coupled superconducting bolometers. It has been designed for high sensitivity and low systematics in order to pursue the primordial B-mode polarization signal. The instrument and its characterization are presented in this thesis, with particular attention to the detectors and readout system. The instrument sensitivity and mapping speed are discussed along with the expected constraint on the scalar-to-tensor ratio that should be made from one year of data.

Contents

Abstract	iv
List of Figures	ix
List of Tables	xii
1 Introduction	1
1.1 inflation	1
1.2 CMB polarization	3
1.3 polarized foregrounds	9
1.4 state of the field	11
1.5 thesis outline	12
2 Instrument	13
2.1 TES bolometers	15
2.2 cryostat and thermal design	20
2.2.1 cryostat	20
2.2.2 sub-Kelvin refrigerator	22
2.2.3 sub-Kelvin thermal architecture	23
2.3 optics	26
2.3.1 vacuum window	27
2.3.2 infrared blocking filters	29
2.3.3 low-pass filter	29
2.3.4 lenses	30
2.3.5 antireflection coatings	30
2.3.6 cold aperture stop	31
2.4 focal plane unit	31

2.4.1	detector plate	32
2.4.2	detector arrays	33
2.4.3	antireflection tiles	43
2.4.4	tile clips	43
2.4.5	printed circuit board	44
2.4.6	Nyquist and multiplexer chips	45
2.4.7	backshort	46
2.4.8	assembly	47
2.5	cold readout electronics	49
2.5.1	time-domain SQUID multiplexing	49
2.5.2	multiplexer chips	53
2.5.3	series arrays	55
2.5.4	Nyquist chips	56
2.5.5	electrical routing	58
2.6	shielding	59
2.7	mount	61
2.7.1	ground shield and forebaffle	63
2.8	optical star camera	63
2.9	warm electronics	64
2.9.1	control overview	64
2.9.2	motion control	65
2.9.3	housekeeping	65
2.9.4	GPS timing	66
2.9.5	multi-channel electronics	66
3	Observations	76
3.1	overview	76
3.2	observing site	76
3.3	scan strategy	78
3.4	observing fields	79
3.5	observing schedule	81
3.6	maintenance	83
3.6.1	cryogenics	83

3.6.2	star pointing	84
3.7	calibration measurements	85
3.7.1	relative gains	85
3.8	initial data reduction	87
3.9	data quality checks	88
3.10	observation efficiency	89
4	Detector and Readout Characterization	90
4.1	overview	90
4.1.1	device parameters	91
4.1.2	time constants and stability	93
4.1.3	loading	94
4.1.4	optical efficiency	95
4.1.5	crosstalk	96
4.1.6	noise characterization	103
4.1.7	aliased noise	109
4.1.8	optimization	110
4.1.9	sensitivity characterization	116
5	Instrument Characterization	122
5.1	overview	122
5.2	measurement techniques	124
5.3	beam characterization	125
5.4	polarization characterization	127
5.5	pointing	129
5.5.1	differential pointing	130
5.6	relative detector gains	132
5.6.1	temporal transfer functions	132
5.6.2	relative responsivities	134
5.6.3	spectral response	135
5.6.4	absolute calibration	136
5.6.5	leakage assessment	137
5.7	systematic impact	138

5.7.1	simulations	141
5.7.2	mitigation	142
6	Analysis	144
6.1	pipeline overview	144
6.2	cuts	145
6.3	mapmaking	146
6.4	power spectra estimation	148
6.4.1	pseudo- C_ℓ estimators	149
6.4.2	pure pseudo- C_ℓ estimators	150
6.4.3	band power window functions	151
6.5	simulations	152
6.5.1	noise-only simulations	154
6.5.2	signal-only simulations	159
6.5.3	filter and beam suppression factor	161
6.6	signal and noise simulations	162
6.7	consistency checks	164
6.7.1	jackknives	164
6.7.2	χ^2	165
6.8	constraint on r	167
6.8.1	band power selection	168
6.8.2	conclusions	169
	Bibliography	177

List of Figures

1.1	Pure E- and B-mode polarization signals for a single Fourier mode on a small (flat) patch of sky	4
1.2	Scalar and tensor quadrupole moments and local polarization produced from Thomson scattering	5
1.3	Plane wave modulation of the local quadrupole scattering polarization patterns	7
1.4	Modeled temperature and polarization power spectra for scalar and tensor perturbation	8
1.5	The polarized foregrounds that contribute to the B-mode polarization signal	10
2.1	The superconducting transition of a TES and its readout circuit	16
2.2	A cross-sectional view of the BICEP2 receiver	21
2.3	The three-stage He sorption fridge	23
2.4	The components of the sub-Kelvin structure and 4 K insert	24
2.5	A schematic of the optics chain	27
2.6	Layout of the BICEP2 beams with nominal locations, beam widths and polarization orientations	28
2.7	An image of the focal plane unit	32
2.8	The focal plane unit fully assembled and ready for installation	33
2.9	A close-up image of the illumination side of the focal plane unit	34
2.10	A schematic of a single detector tile with polarization orientation of the two bolometers	35
2.11	A micrograph of a portion of a focal plane pixel	35
2.12	A micrograph of a portion of a bolometer antenna	36
2.13	A micrograph of the inline microstrip filter	38
2.14	A micrograph of the free-standing bolometer island	39
2.15	The detector tile thermal components	42
2.16	An image of a pair of SQUID multiplexer and Nyquist interface chips	45
2.17	The superconducting niobium backshort	46

2.18	Micrographs of a multiplexer chip and close-up of the first-stage and dummy SQUID	53
2.19	An image of a Nyquist interface chip and the interface circuit for one TES	57
2.20	A schematic of the BICEP2 telescope and exterior view of the telescope looking out	62
3.1	An image of the Dark Sector Laboratory housing BICEP2	77
3.2	Optical depth and precipitable water vapor at leading millimeter-wave observing sites	78
3.3	The BICEP2 observing fields and the Galactic dust emission at 150 GHz	80
3.4	Scan pattern of the observing fields in rectangular and equatorial coordinates	82
3.5	Elevation nods and individual and pair-differenced timestreams calibrated with the relative gain	86
4.1	TES load curve analysis	92
4.2	The various types of crosstalk observed in bright source rasters	97
4.3	Timestreams of channels showing crosstalk when the source channels exceed a DAC output threshold	101
4.4	Noise spectra through the transition for a characteristic detector sampled at 400 kHz	106
4.5	Measured and modeled current noise for the characteristic detector at the observation bias	107
4.6	Histograms of the measured NEIs and NEPs for the operational detectors at the standard observation biases	109
4.7	Sensitivity versus bias profiles used to select the optimal detector biases	111
4.8	SQ1 $V - \Phi$ response before and after an overhaul of the SQUID biases	113
4.9	Per-detector sensitivity spectra averaged over all detectors and the distribution of per-detector NETs	117
4.10	Maps and histograms used to derived the map-based instrument sensitivity for the 2011 observing configuration	119
5.1	The setup and result of mapping a ground-fixed calibration source	124
5.2	Decomposition of differential beam systematics	125
5.3	Measured differential pointing	131
5.4	Measured spectral response of a pair of detectors	135
5.5	Simulated power spectra contamination from the measured differential pointing error	140
6.1	Band power window functions	151

6.2	Real and simulated T, Q, and U maps	153
6.3	Angular power spectra for noise-only simulations	155
6.4	Real and simulated timestreams	156
6.5	Real and simulated PSDs	157
6.6	Angular power spectra for signal-only simulations	160
6.7	Filter and beam suppression factors	161
6.8	Angular power spectra for signal-plus-noise simulations	163
6.9	95% confidence limit on r from systematic-free simulations	168

List of Tables

2.1	BICEP2 instrument design summary	14
2.2	BICEP2 bolometer design parameters for the titanium and aluminum TESs	19
2.3	MCE electrical output	67
3.1	Phases of the BICEP2 observing schedule	81
4.1	Designed and measured detector performance	91
4.2	Measured and predicted crosstalk mechanisms	96
4.3	Noise budget for the BICEP2 detectors from the noise model	103
4.4	Sensitivity and mapping speed comparison for BICEP1 and BICEP2 from the map-based and timestream-based methods	118
4.5	Designed and measured sensitivity	120
5.1	Systematic errors that could produce false B-mode polarization	138
5.2	Calibration uncertainties affecting the power spectrum amplitudes only	138
6.1	Comparison of BICEP2 2010 and BICEP1 2006–2007 data quality	148

Chapter 1

Introduction

1.1 inflation

Cosmic inflation was proposed to solve issues that the hot big-bang paradigm alone isn't able to explain [22]. The first issue — the horizon problem — is that the universe is homogenous to a part in 10^5 over scales on which the universe could have had no causal contact. The angular scale subtended by the horizon at the last scattering surface is $\sim 1.6^\circ$. Thus the angular scales we observe should have no causal contact and no isotropy on scales larger than a few degrees. Yet the CMB is isotropic to one part in 10^5 in completely opposite directions. The second issue — the flatness problem — is that in order to explain the level of flatness measured for the universe the initial value must have been incredibly finely tuned. The third issue — the monopole problem — is the relative absence of relics, such as magnetic monopoles, that are predicted by the standard model for an expanding universe. If the early universe underwent a period of accelerated expansion, increasing in scale by at least 26 orders of magnitude, then these issues are resolved. During inflation the size of the horizon stays fixed while the universe rapidly expands. All observable scales are stretched beyond the horizon leaving a homogenous and incredibly flat early universe.

To generate this accelerated expansion a negative vacuum energy is needed. This is similar to, but not the same as, the vacuum energy that dominates the recent universe. In order to allow the universe enough time to expand and solve the problems listed above, inflation must have proceeded relatively slowly. The scalar field rolled slowly down a potential from its original value to lower energy and then decayed into a hot plasma of matter and radiation, the so-called reheating.

In addition to solving these problems inflation has experimentally verifiable predictions, notably in the properties of the CMB anisotropy. Foremost, inflation predicts a nearly scale-invariant background of perturbations to the metric. The perturbations are generated by quantum mechanical

fluctuations that get stretched to observable scales by inflation. There are three distinct perturbations: scalar (compressional), vector (vortical) and tensor (gravitational wave). These perturbations can be represented by independent oscillatory modes with wave vector \vec{k} . Inflation stretches the physical scale of these modes until they are larger than the horizon and, due to causality, no longer evolve. After the end of inflation the horizon begins expanding again, the modes re-enter the horizon and begin to evolve. Vector perturbations decay away and are not expected to give rise to observable anisotropies.

Scalar perturbations are oscillations in the gravitational (Newtonian) potential and spatial curvature. Scalar perturbations give rise to density fluctuations in the hot plasma of the universe when they re-enter the horizon. These develop into oscillations in the photon-baryon fluid and later form large-scale structure through gravitational instability. The frequency of the oscillations is set by the speed of sound in the photon-baryon fluid. Scalar perturbations are responsible for the observed temperature and polarization anisotropy in the CMB. The spectrum of scalar perturbations are characterized by a spectral index n_s and amplitude A_s . Consider a single Fourier mode of a compression wave sourced by scalar perturbations. The regions where the wave is maximally compressed are potential wells, or troughs. The troughs are effectively cold temperature regions because as the photons leave the potential wells they lose energy. The opposite is the case for maximally rarified regions, which are effectively hot crests. Gradients in the gravitational potential and temperature cause a flow of photons from effectively hot crests to cold troughs.

Tensor perturbations can be viewed as standing gravitational waves formed during inflation. The waves have either “x” or “+” polarization and oscillate at a frequency set by the speed of light. To visualize the effect of gravitational waves, imagine a single plane wave mode passing through this page with a wave vector normal to it. The “x” polarized gravitational wave will cause an “O” to get stretched into an ellipse with a major axis alternating between $\pm 45^\circ$ and the “+” polarization will stretch the major axis between 0° and 90° . Like the scalar perturbations, the tensor perturbations are parametrized by the tilt n_t and amplitude A_t of their spectrum. The amplitude is generally expressed in terms of the ratio of the tensor and scalar amplitudes at some pivot point. For this thesis the convention $r \equiv \frac{\Delta_t^2(k_o)}{\Delta_s^2(k_o)} \propto \frac{A_t^2}{A_s^2}$ evaluated at the wave number $k_o = 0.002 \text{ Mpc}^{-1}$ is used. Unlike scalar perturbations, tensor perturbations have not been detected. They are expected to contribute to the CMB temperature and polarization anisotropy, but not the density fluctuations that develop into structure. Their dominant contribution is to the temperature anisotropy at large angular scales, as shown in figure 1.4. The difference between the observed temperature anisotropy

and that expected from scalar modes alone provides a means to constrain the tensor modes. However, this method is fundamentally limited to a lower bound of $r_{limit} \simeq 0.13$ by cosmic variance [35]. To make a better constraint or to detect tensor perturbations we must look for another signature of tensors, explained below.

The creation of scalar and tensor perturbations is a generic prediction of inflationary models. The perturbation parameters, however, are sensitive to the model for the scalar field that drove inflation. The spectral indices n_s and n_t depend on the energy of inflation and the rate at which it proceeded. In slow roll inflation, n_s is degenerate with r , providing an additional way to constrain r and a consistency check for a primordial tensor detection. The simplest models of inflation predict that the spectrum of perturbations won't be entirely scale invariant, but will have $n_s < 0.98$ [9].¹ Additionally, some models of inflation predict that the spectral index itself changes with the wave number in a so-called running spectral index.

The tensor or gravitational wave background has many observable consequences. If it exists and can be measured, it will be direct evidence of inflation and will further reveal the energy scale that inflation occurred at since $A_t \propto E_{inf}^2$. However, there is no certainty that it is large enough to be detected. If inflation occurred at an energy scale smaller than $3.2 \cdot 10^{15} \text{ GeV}$ it will not be detectable in the CMB polarization since it will be irreducibly obscured by the lensing foreground discussed below [33]. This lower detection limit is $r = 6 \cdot 10^{-5}$, which is much lower than most experiments will be able to reach in the near term.² If inflation occurred at the energy scale of the grand unified theory (GUT), as some theories predict, then $E_{inf} \sim 10^{15-16} \text{ GeV}$ and detectable tensor perturbations would be in the range of $r \sim 6 \cdot 10^{-5} - 4 \cdot 10^{-3}$. There is compelling evidence that tensor perturbations exist at a detectable level. In 2007 the Wilkinson Microwave Anisotropy Probe (WMAP) satellite team announced a significant detection of $n_s < 1$ [57]. Recent analysis of the seven-year data set shows $n_s = 0.968 \pm 0.012$, excluding scale-invariance at the 99.5% confidence limit [36]. For the simplest inflationary models, if $n_s \gtrsim 0.95$ then $r \gtrsim 0.01$, which is within reach of the current generation of CMB polarimeters [9].

1.2 CMB polarization

A CMB polarimeter measures temperature T , and Stokes Q and U on the sky, where $\pm Q$ is oriented parallel and normal to meridian lines and $\pm U$ is oriented $\pm 45^\circ$ to the meridian lines. An alternative

¹For example, in slow-roll inflation $n_s = 1 - 6\epsilon + 2\eta$, where $\epsilon \propto (V'/V)^2$ and $\eta \propto (V''/V)$.

²The scaling relationship for slow-roll inflation is assumed here: $r = 0.0043(\frac{E_{inf}}{10^{16} \text{ GeV}})^4$ [26].

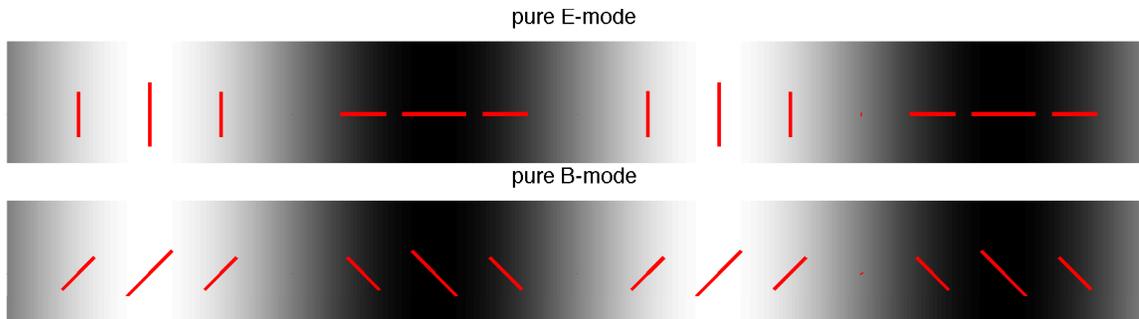


Figure 1.1: Pure E- and B-mode polarization signals for a single Fourier mode on a small (flat) patch of sky. For an E-mode, the polarization is either parallel or normal to the wave vector of the mode. For a B-mode, the polarization is oriented $\pm 45^\circ$ to the wave vector. The background shading indicates the amplitude of the mode. This figure has been adapted from the reference [11].

basis that doesn't have this dependence is the E- and B-mode polarization basis. More importantly, scalar perturbations only contribute to the E-mode polarization signal and B modes are generated only from tensor perturbations. In analogy with electromagnetics, the E-mode polarization is curl-free and invariant under parity change (electric parity), while the B-mode polarization is divergence-free and reverses under parity change. Pure E- and B-mode polarization is shown in figure 1.1 over a small patch of sky for a single Fourier mode with a wave vector parallel to the top of the page. For the E-mode pattern the polarization is oriented either parallel with or normal to the wave vector. For the B-mode pattern the polarization is oriented $\pm 45^\circ$ to the wave vector.

Polarization in the CMB is generated by Thomson scattering of electrons at the surface of last scattering and only occurs when the incident photons possess a quadrupole anisotropy. Circular polarization is produced by scattering of relativistic photons, so the CMB is expected to be only linearly polarized.³ The cross section of Thomson scattering depends on the polarization of the incident photons. A scattered photon is polarized normal to the plane formed by the incident and scattered trajectory. If the incident radiation on the electron is isotropic then no polarization is created since orthogonal polarization states cancel out. Only when there is a quadrupole temperature anisotropy will the scattered radiation possess a net polarization. The local polarization pattern of the scattered radiation follows the projection of the cold lobes of the quadrupole.

If the scattering is rapid then no net polarization is created since any local anisotropy and polarization is randomized. CMB polarization only occurs at the last scattering surface when the mean free path for scattering events has become large. The level of polarization of the CMB is sensitive to the optical depth at last scattering or, effectively, the spatial extent than an electron

³Hence, no measurement of Stokes V is made.

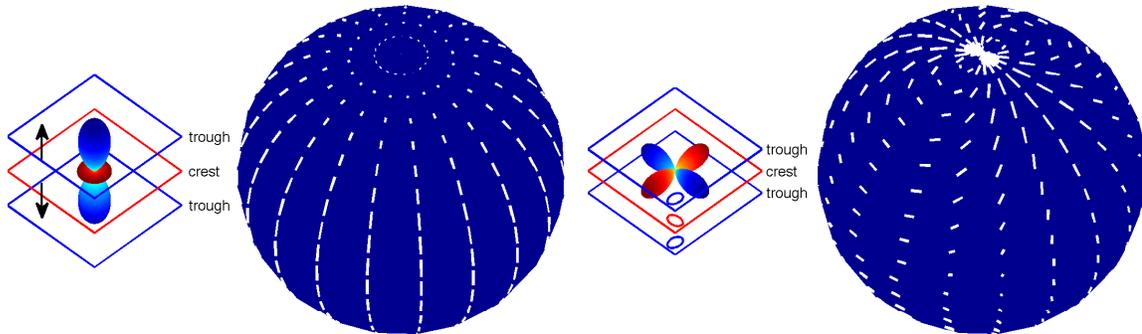


Figure 1.2: Scalar and tensor quadrupole moments and local polarization produced from Thomson scattering. *Far left:* The $m = 0$ quadrupole moment produced by a single mode of a scalar compression wave indicated by the blue and red squares, with wave vector normal to the squares. The arrows represent the flow of photons from the hot crest to cold troughs. *Middle left:* The local polarization pattern created by the $m = 0$ moment through Thomson scattering. The polarization orientation follows the projection of the cold (blue) lobes on the sphere. *Middle right:* The $m = \pm 2$ quadrupole moment produced by a single mode of a gravitational wave. The stretching of space and photon wavelengths by the wave is represented by the blue and red ellipses. *Far right:* The local polarization pattern created by the $m = 2$ moment through Thomson scattering. These figures are adapted from the reference [25].

can sample the quadrupole around it.

Scalar and tensor perturbations generate quadrupole temperature anisotropy at last scattering and give rise to CMB polarization. The anisotropy is distinct for the two and given by the $\ell = 2$ spherical harmonics, with scalar perturbations generating the $m = 0$ moment and tensor perturbations generating the $m = \pm 2$ moments for the two polarizations of the gravitational waves.⁴

The oscillations of the photon-baryon fluid sourced by scalar perturbations generates anisotropy at last scattering. In the scalar compression wave considered earlier photons flow from the effectively hot crests to the cold troughs with velocity parallel to the wave vector of the mode. An electron sitting in an effectively cold trough (over dense region) is surrounded by cold photons, but sees a flow of hot photons above and below it. The opposite case is shown in figure 1.2 for an electron sitting on a hot crest. The wave vector of the Fourier mode is normal to the blue and red squares. The arrows indicate the flow of photons from the hot crest to cold troughs. The photons in the plane of the crest are hot while the photons above and below are colder. The local polarization pattern produced by scattering from this $m = 0$ moment is also shown in the figure. The fact that the photon flow and the wave vector are parallel gives the scattering from scalar perturbations azimuthal symmetry. The scalar pattern has only Q polarization, no U polarization.

Tensor perturbations create the $m = \pm 2$ moments by effectively stretching space and photon

⁴Vector perturbations generate the $m = \pm 1$ moment.

wavelengths at last scattering. One polarization of the gravitational wave considered earlier is shown in figure 1.2. A single Fourier mode is shown with wave vector normal to the blue and red squares. In the crest of the wave space has been stretched like the red ellipse. An electron in the crest will see photons cooled by the stretching of their wavelength along the major axis of the ellipse. Photons incident along the minor axis of the ellipse will arrive hotter. The same is true for the blue troughs but the effect and lobes of the quadrupole are rotated by 90° , as indicated by the blue ellipses. The local polarization pattern for scattering from the $m = \pm 2$ moment is also shown in the figure. Unlike the local polarization pattern of scalar scattering, the tensor scattering pattern has both Q and U polarization.

Both of the local polarization patterns from tensor and scalar quadrupole scattering have electric parity. However, while the scalar pattern has only Q polarization, the tensor pattern contains both Q and U polarization. It is the modulation of these local polarization patterns by a Fourier mode that give rise to the polarization pattern observed in the surface of last scattering. Figure 1.3 illustrates how plane wave modulation creates E- and B-mode polarization patterns for scalars and tensors. A plane wave with a north-south wave vector modulates the scalar and tensor quadrupole moments, which reverse the sign of the lobes between the crests and troughs. As the angle the quadrupole is viewed at changes, so does the projection of the local polarization pattern. The scalar mode creates only E-mode polarization from its Q-only local polarization pattern. Just as in figure 1.1, the polarization is parallel or normal to the wave vector of the mode. The tensor quadrupole, with its local combination of Q and U, creates both E- and B-mode polarization. It's clear that some component of the modulated polarization pattern is $\pm 45^\circ$ to the wave vector, which is the pattern of B-mode polarization shown in figure 1.1.

Scalar perturbations produce only E-mode polarization. Tensors produce E- and B-mode polarization in roughly equal amounts from the “+” and “x” gravitational waves, respectively, [52]. The fact that B-mode polarization is only created by tensor perturbations means that a detection of primordial B-mode polarization corresponds to a detection of the primordial gravitational wave background created during inflation. B-mode polarization is the “smoking-gun” of inflation; a unique signature of inflation imprinted in the polarization of the CMB. This statement is qualified by “primordial” since there are other mechanisms that can produce B-mode polarization, such as weak lensing, reionization and Galactic foregrounds. A measurement of the amplitude of B-mode polarization, or r , provides a measurement of the gravitational wave amplitude, the energy scale of inflation and the shape of the inflaton potential. Even if no signal is detected the constraint on r

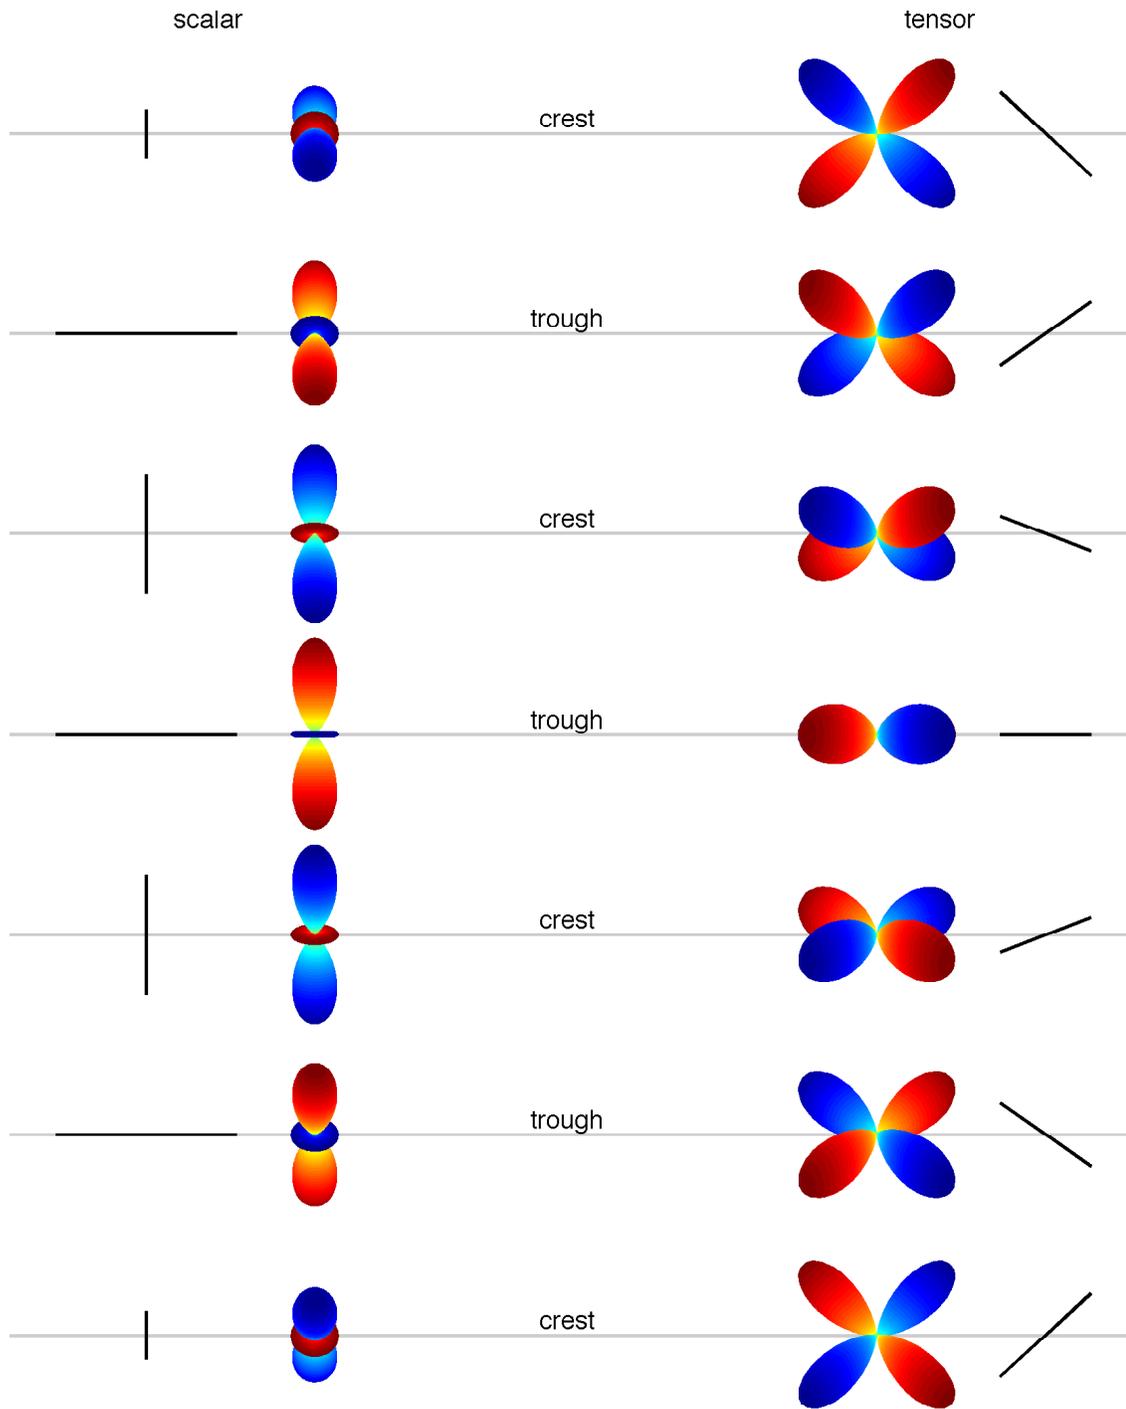


Figure 1.3: Plane wave modulation of the local quadrupole scattering polarization patterns. This figure was inspired by the reference [25] in which the modulation of the $m = 0$ moment was shown.

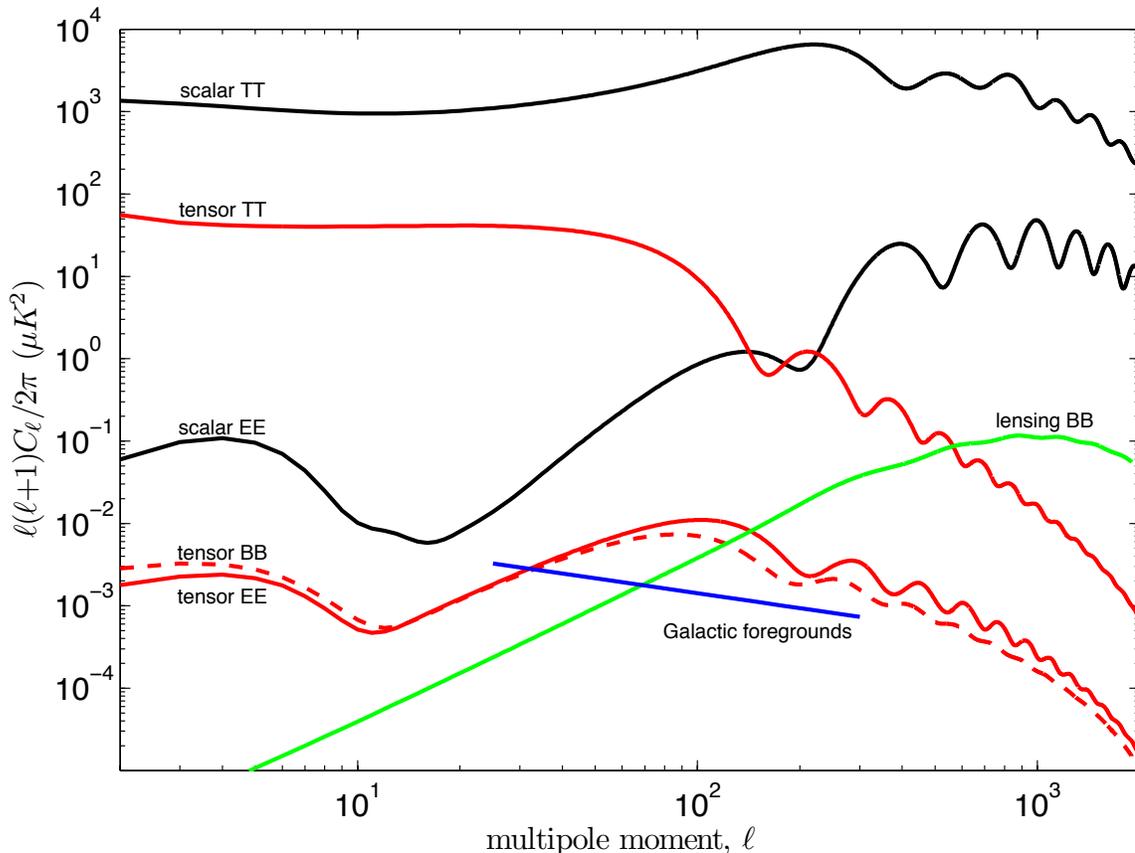


Figure 1.4: Modeled temperature and polarization power spectra for scalar and tensor perturbation. Only the auto-correlations (TT , EE , BB) have been shown. This fiducial model was created using the CAMB modeling code and measured data from the WMAP five-year result [41][19]. The tensor perturbations have been generated for $r = 0.1$. In addition to the primordial spectra, the power spectra for the lensing and Galactic polarized foregrounds at 150 GHz are shown for one of the cleanest observing regions in the sky.

places valuable limits on the allowable models for inflation.

The two-point correlations of the temperature and polarization signals are expressed as power spectra in a spherical harmonic basis, as explained in the analysis chapter. The power spectra take the form $\frac{\ell(\ell+1)}{2\pi}C_\ell^{XX}$, where $XX = \{TT, EE, TE, BB, EB, TB\}$ is one of the auto-correlations or cross-correlations. The prefactor $\frac{\ell(\ell+1)}{2\pi}$ is typically used to preserve the near-scale-invariance of the spectra when they are plotted against the multipole moment ℓ . A fiducial Λ CDM⁵ model for the temperature and polarization power spectra are shown for the auto-correlations in figure 1.4. The model spectra have been created using the CAMB⁶ modeling code and best-fit cosmological parameters from the WMAP five-year result [41][19]. The tensor spectra $C_\ell^{t:TT,EE,BB}$ have been

⁵Lambda or dark energy and cold dark matter

⁶Code for anisotropies in the microwave background

generated with $r = 0.1$, which gives a peak $\frac{\ell(\ell+1)}{2\pi}C_\ell^{BB} \simeq 0.007 \mu\text{K}^2$ at $\ell \simeq 90$. The dominant polarized foregrounds created by Galactic synchrotron and thermal dust, and weak lensing are shown as well (§ 1.3). Parity invariance dictates that there will be no correlations between tensor B-mode polarization and temperature and E-mode polarization, so $C_\ell^{t:TB,EB} = 0$. Correlations are expected between the tensor temperature and E-mode polarization $C_\ell^{t:TE} \neq 0$ [52].

It is worth noting the features in the polarization power spectra that might seem to conflict with the near-scale-invariance of the primordial perturbations. Such effects as photon diffusion, the spatial scale of the quadrupole anisotropy and the oscillation and decay of perturbations impart a transfer function on the primordial spectra and shape the observable spectra. The bumps in the spectra at large angular scales are due to the polarization signals generated by reionization. Without reionization the scalar and tensor polarization spectra would fall at large angular scales. The quadrupole anisotropy at the surface of last scattering is limited to small scales near the electron, dictated by the mean free path at last scattering. At smaller angular scales the scalar temperature and E-mode spectra fall due to damping from photon diffusion. Thus, the scalar E-mode spectrum peaks at the damping scale, $\ell \simeq 1000$. The tensor spectra fall off at small angular scales because the amplitude of the perturbations of those modes have entered the horizon by the time of last scattering. Gravitational waves decay with the expansion of the universe after entering the horizon. Thus, the tensor polarization spectrum peaks at the observed angular size of the horizon at last scattering, $\ell \simeq 90$.

The ripples in the spectra are due to oscillations of the perturbations. The acoustic oscillations of the scalar spectra have a frequency set by the speed of sound in the photon-baryon fluid. The peaks in the scalar E-mode spectrum are out of phase with the peaks in the temperature spectrum since scalar E-mode is sourced by the velocity field, which is out of phase with the compression wave. The peaks in the tensor spectra are caused by the oscillations of the standing gravitational waves. This frequency is set by the speed of light, not the speed of sound, so the peaks of tensor spectra have a factor of $\sqrt{3}$ lower period.

1.3 polarized foregrounds

Polarized foregrounds limit the level to which the primordial B-mode polarization can be measured. These signals are of scientific interest, but are a source of confusion for an experiment pursuing the primordial B-mode signal. The foregrounds include those signals produced by Galactic emission, weak lensing and reionization. The dominant Galactic foreground contamination is from polarized

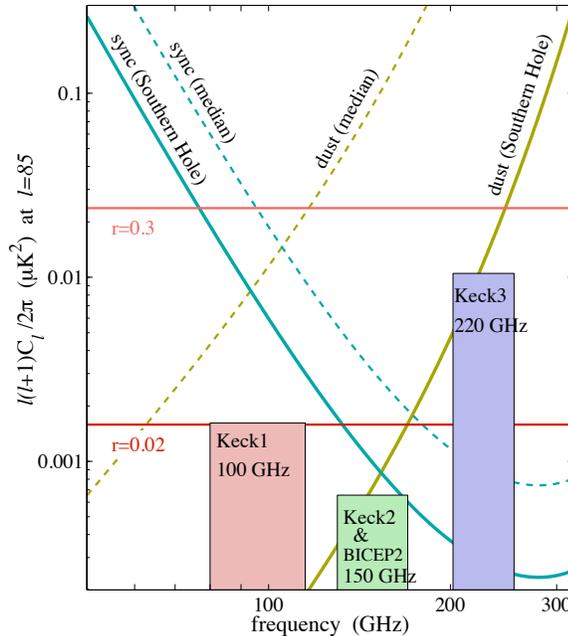


Figure 1.5: The polarized foregrounds that contribute to the B-mode polarization signal. The “Southern Hole” in the dust emission of the Galaxy provides a clean window to observe the CMB polarization (§ 3.4). In this window the polarization contamination is low enough to allow a measurement of primordial B-mode polarization down to a level of $r \sim 0.01$. In this region the optimal frequency band for observing is ~ 150 GHz, which coincides with the peak of the CMB blackbody emission. Note, the median sky has a level of polarization contamination from dust and synchrotron emission that would limit a measurement of r to around 0.3.

emission from thermal dust and synchrotron radiation. Weak gravitational lensing rotates some of the brighter primordial E-mode signal into B-mode polarization. During the epoch of reionization, when hydrogen is ionized by the radiation emitted by the first generation of stars, free electrons produce B-mode polarization through Thomson scattering. However, this signal will occur at large angular scales and peak at $\ell \sim 4$, which is outside the multipole range of ground-based polarimeters. While these late-time sources of B-mode polarization can create considerable foreground confusion, there are ways to discriminate, avoid or even reduce them.

Thermal dust grains emit radiation, which can be polarized by the alignment of the dipole moment of the grains with the Galactic magnetic field. Polarized synchrotron radiation is created by relativistic electrons spinning in the Galactic magnetic fields. Thermal dust emission has been modeled by scaling the standard “FDS model 8” for unpolarized dust emission up to 150 GHz [21]. To estimate the level of polarized signal from thermal dust a polarized fraction of 5% has been assumed, based on the results from the WMAP team’s analysis of three-year polarization data [47]. The synchrotron emission has been modeled from the WMAP three-year 23 GHz (K-band) data,

which is currently the best measurement of synchrotron emission. The power spectrum of dust and synchrotron is expected to scale as $\ell^{-0.6}$, as shown in figure 1.4 [47]. Both of these sources are strongly frequency-dependent, as shown in figure 1.5. A possible primordial or lensing B-mode signal can be distinguished from Galactic foreground confusion by observing in additional frequency bands. The spectral indices for the Galactic emission are based on the findings of the WMAP team. The spectral indices used for modeling the dust emission and synchrotron are $\beta_d = 1.8$ and $\beta_s = -3.0$, respectively, [47]. Galactic foreground confusion can be avoided to some degree by careful choice of the observing region. A minimum in the Galactic foreground emission exists at ~ 150 GHz, as shown in figure 1.5, for the clean patch of sky considered in this thesis (§ 3.4).

Weak gravitational lensing caused by large scale structure between the observer and last scattering surface will shear the relatively bright scalar E-mode signal into B-mode polarization. As shown in figure 1.4, its spectrum tracks that of the scalar E-mode signal peaking at $\ell \simeq 1000$. It doesn't have the same spectral dependence as the Galactic foregrounds, so must be discriminated from primordial B-mode polarization by the relative shapes of their spectra. The lensing signal can be reduced by around a factor of 10 by reconstructing the lensing potential that causes the shear. The residuals from this procedure can still overwhelm the primordial tensor signal if $r \lesssim 6 \cdot 10^{-5}$ [33]. If the tensor amplitude is below this level then it will be undetectable.

1.4 state of the field

E-mode polarization was first detected in 2002 and continues to be precisely measured [37]. B-mode polarization, whether primordial or from foregrounds, has not been detected. The current best constraint on r has been set by the contribution of tensor perturbations to the CMB anisotropy at large angular scales.⁷ An indirect upper limit has been set at $r < 0.36$ by the WMAP team from the seven-year data [40]. When they include additional external priors the limit decreases to $r < 0.24$ [36].⁸ The SPT and ACT teams have slightly improved upon these limits by including their measurements of the temperature anisotropy at high ℓ with the priors. The SPT and ACT teams report constraints of $r < 0.17$ and $r < 0.19$, respectively, [32][18]. These results are approaching the predicted limit of the constraint provided by the temperature anisotropy of $r_{limit} = 0.13$ [35]. In order to probe tensor perturbations below r_{limit} they must be measured directly from primordial B-mode polarization. The current best direct upper limit was set at $r < 0.72$ in 2009 by the BICEP1

⁷All upper limits on r will be quoted at the 95% confidence limit throughout this thesis.

⁸A limit of $r < 0.20$ is also given from different priors, but claimed to be less robust than the one given here.

team from the analysis of the first 2 years of data [13].

An optimal search for primordial B-mode polarization can be made in a very clean patch of sky at the degree-scale ($\ell \sim 100$) at which the signal is expected to peak. In general, the pursuit of B-mode polarization requires low instrument systematics that produce false polarization, high instrument sensitivity and a long, deep integration of a clean contiguous degree-scale patch of sky.

1.5 thesis outline

This thesis is organized in the following way. First the BICEP2 instrument (§ 2) is presented, with particular detail given in this section as a means of documenting some of the knowledge of the author. Then the observing strategy (§ 3) is presented, detailing how BICEP2 is used to pursue B-mode polarization. In the next section, a characterization of the detectors and readout is presented (§ 4). The noise and sensitivity of the instrument is presented in that section with a prediction for the pursuit of B-mode in the absence of systematics. Then, an extensive description of the characterization of the instrument is presented (§ 5). In order to take detector measurements to measurements of polarization on the sky as clean of spurious instrumental polarization as possible, the instrument must be characterized to high precision. The remaining systematic error in characterizing the instrument impacts the purity of the resulting B-mode polarization measurement. In the final analysis chapter (§ 6) an overview is given of the data analysis steps that are used to produce measurements of polarization power spectra. Signal and noise only simulations are used to predict the prospects BICEP2 has for detecting or constraining r with the first year of data.

Chapter 2

Instrument

BICEP2 is a ground-based CMB polarimeter utilizing a new generation of antenna-coupled transition edge sensor (TES) bolometers. It has been designed with the explicit goal of detecting B-mode polarization, with an emphasis on careful control of systematics. It is an on-axis refractive telescope with a 26 cm aperture, operating in a single observing band (148 GHz). BICEP2 builds on the success of its predecessor BICEP1, which observed from 2006–2009 and has the current best direct constraint on the tensor-scalar ratio of $r < 0.72$ [13]. Much of the cryostat and thermal design, optical design, mount, telescope control and housekeeping is based on or exactly the same as that used in BICEP1. Key aspects of the instrument design are summarized in table 2.1.

BICEP2 measures microwave radiation using lithographically fabricated, dual-polarization, antenna-coupled superconducting bolometer arrays. The focal plane has four monolithic detector tiles that have an 8×8 focal plane pixel format, yielding a total of 256 detector pairs. The pixel spacing is $2f\lambda$, with focal plane pixels on a 7.8×7.8 mm pitch. The monolithic, antenna-coupled design has allowed for a five-fold increase in the detector count compared to BICEP1. Monolithic fabrication allows for higher focal plane-packing efficiency, and reduced thermal mass and load on the focal plane, since it eliminates the need for bulky feedhorns.

Each focal plane pixel includes two sets of independent phased-array slot antennas, one per orthogonal polarization. Each antenna array is filtered by an inline microstrip filter and read out by a single transition edge sensor (TES) bolometer. The focal plane has 512 bolometers in total. Current signals from the voltage-biased TESs are read out using time-domain multiplexed (TDM) superconducting quantum interference devices (SQUIDs) [16][29]. Polarization is measured by differencing the gain-corrected detectors within a polarization pair. This differencing scheme removes much of the atmospheric and detector $1/f$ noise and common-mode signal contamination.

The 148 GHz observing band was selected to complement a band of reduced atmospheric emission

Table 2.1: BICEP2 instrument design summary

Band Center	Bandwidth	Beam FWHM	Detector	NET per Detector	Instrument NET
GHz	GHz	degrees	pairs	$\mu\text{K}_{\text{CMB}}\sqrt{\text{s}}$	$\mu\text{K}_{\text{CMB}}\sqrt{\text{s}}$
148	37.5	0.47	256	330	15.1

at the South Pole (§ 5.6.3). Conveniently, the band also corresponds to a null in the Galactic foreground emission (§ 1.3) and is close to the ~ 160 GHz peak of the 2.7 K blackbody emission of the CMB. A high end-to-end bandwidth of $\sim 25\%$ gives the detectors high optical efficiency and high sensitivity (§ 4.1.4).

An on-axis compact refractor has many advantages for careful control and characterization of instrumental systematics. The entire optics chain is housed in the cryostat and is cooled, which gives low, stable optical loading of the detectors. The small aperture allows for the use of a comoving forebaffle to prevent the polarized sidelobes of the beams from scanning the environment. The cryostat can be rotated about the boresight during observations to average down some systematic effects. The small aperture allows for beam-filling calibration measurements in the near field. The far field is located at ~ 70 m, allowing for thorough far-field instrument characterization using ground-fixed sources. The receiver is monochromatic, so the optics have been optimized for a single frequency, reducing distortion, reflections and giving higher optical efficiency. The $\sim 30'$ beams allow BICEP2 to integrate deeply at the angular scale at which the primordial B-mode power is expected to peak.

BICEP2 was deployed in December 2009 to the South Pole, Antarctica, and will observe through the austral winter of 2012. The high, dry, and stable atmosphere and long polar night yield ideal conditions for long-integration, mm-wave observations (§ 3.2). Additionally, the South Pole gives access to a unique patch of sky, the Southern Hole, where polarized foreground confusion from synchrotron radiation and especially thermal dust emission is much lower than the mean sky (§ 3.4). For a single 148 GHz receiver, detection of B-mode polarization should not be obscured in this field for $r > 0.03$.

In December 2010, the Keck array was deployed to the South Pole to complement BICEP2. Three monochromatic receivers operating at 150 GHz were deployed, with two additional receivers deployed in the austral summer of 2011–2012. Eventually, several of the receivers will be modified to observe at 100 and 220 GHz [53]. The multiple frequency bands are necessary for discrimination of the primordial B-mode from the frequency-dependent Galactic foregrounds. The phased-array

antennas and monolithic fabrication allow the detectors to be scaled more easily to these additional observing bands than is possible with absorber bolometers [39].

The range of multipole moments investigated by BICEP2 is $\ell = 30 - 300$, with primordial B-mode peaking at $\ell \simeq 90$. The range is bounded at high multipole (small angular scale) by the resolution of the 0.5° beams and bounded at low multipole (large angular scale) by the size of the observing field. With a scan rate of $2.8^\circ/s$, this multipole range corresponds to a science band of $0.1 - 1 Hz$.

The detector design has been optimized for high instantaneous per-detector sensitivity or low noise-equivalent temperature (NET). The designed per-detector NET is $330 \mu K\sqrt{s}$, which should give roughly background-limited performance at the South Pole. The device parameters are carefully controlled for high uniformity and yield and the multiplexer has been optimized for high yield as well (§ 4). The high yield and per-detector sensitivity should give BICEP2 an instrument sensitivity of $NET_{inst} = \frac{NET_{det}}{\sqrt{N_{det}}} = 15.1 \mu K\sqrt{s}$. This corresponds to a factor of ~ 9 increase in mapping speed over BICEP1. The designed per-detector NETs of the two instruments are roughly the same, so much of this sensitivity increase is due to the fivefold increase in detectors from the higher packing efficiency provided by monolithic arrays. Such a high sensitivity is necessary to integrate down for a signal-to-noise sufficient to detect the faint B-mode polarization signal. With this projected array sensitivity, BICEP2 should achieve the raw signal-to-noise in a single season needed to detect $r > 0.01$, or constrain r if it is below that level.

2.1 TES bolometers

BICEP2 uses antenna-coupled TES bolometers that are read out by time-domain multiplexed SQUIDs. The bolometers have a dual-transition design, meaning they have both a titanium and aluminum TES on each bolometer island. The titanium TES is used for standard low-loading observations, while the aluminum TES allows for measurement at higher loading levels, such as in the lab or when observing at low elevation. This a large step forward in technological development from the spider-web-coupled NTD thermistors used in BICEP1. In this section a brief introduction to the TES bolometer is given. The actual implementation of the bolometers is given in § 2.4.2.4, and the low-temperature SQUID readout is discussed in § 2.5.2.

Bolometers are used to measure millimeter-wave radiation. Incident radiation is deposited on a thermally isolated structure, which heats up in response to the optical power. A thermometer is then used to measure the temperature increase of that island, providing a measurement of the incident power. The optical signal can be directly absorbed by the bolometer or can be coupled to

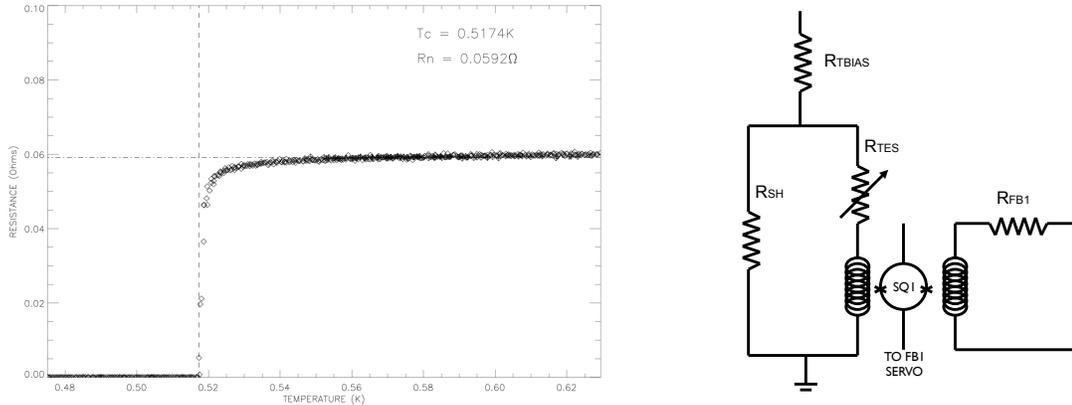


Figure 2.1: The superconducting transition of a TES and its readout circuit. *Left:* The steep change in resistance with temperature of a superconductor makes it ideal for use as the thermistor in a bolometer. This resistance versus temperature profile was acquired for a BICEP2 titanium TES by ramping the bath temperature of the bolometer while measuring the resistance. The transition from the superconducting state to normal state with resistance R_n occurs at the critical temperature T_c . *Right:* The TES bolometer is voltage-biased by a shunt resistor R_{sh} that is much smaller than the operational resistance of the TES, R_{tes} . A SQUID ammeter is used to measure the change in output current of the TES, which is inductively coupled to the magnetically sensitive SQUID. The SQUID is linearized by a feedback servo that nulls the input flux from the TES.

a beam-forming antenna, which then directs the power to the island.

A TES bolometer capitalizes on the steep change in resistance for small changes in temperature of a superconductor in its critical transition. A transition for a BICEP2 TES is shown in figure 2.1. This characteristic makes TESs highly sensitive thermistors, with a sensitivity set by the steepness of the transition. This is expressed by the dimensionless parameter $\alpha = \frac{T}{R} \frac{\partial \log(R)}{\partial \log(T)}$. The temperature at which the TES passes from a normal state with resistance R_n to superconducting state is the critical temperature T_c .

The BICEP2 bolometers are made from free-standing islands of silicon nitride (SiN) lithographically patterned on silicon wafers (§ 2.4.2.4). The island has some heat capacity C and thermal conduction to the lower temperature bath of the silicon. The power flow to the bath (P_{bath}) is given by:

$$P_{bath} = K(T^{\beta+1} - T_{bath}^{\beta+1}) = \frac{G}{(\beta+1)T^\beta} (T^{\beta+1} - T_{bath}^{\beta+1}) = \frac{GT}{\beta+1} \left(1 - \left(\frac{T_{bath}}{T}\right)^{\beta+1}\right) \quad (2.1)$$

where T is the island temperature and T_{bath} is the bath temperature. The differential thermal conductance $G \equiv \frac{dP_{bath}}{dT} = (\beta+1)KT^\beta$ is derived from equation 2.1 and prefactor K . The term β is the thermal conductance exponent and is set by properties of the conducting material but is generally 2–3 for phonon-mediated materials. The G of the bolometer is defined by the geometry of the legs that support the released island.

Optical power (P_{opt}) is absorbed by a phased-array antenna and deposited on the island causing it to heat up. In addition, the TES located on the island is voltage-biased into its transition so that it can be used as a thermistor. Electrical, or Joule, power (P_J) is dissipated on the island and is set by the applied bias (V_{tes}) and resistance of the TES (R_{tes}). A power balance exists in this arrangement between the optical power, Joule power and power flowing to the bath P_{bath} given by the thermal differential equation:

$$C \frac{dT_{tes}}{dt} = -P_{bath} + P_J + P_{opt} \quad (2.2)$$

In steady state the Joule and optical power balance the heat flow to the bath. If the optical load increases the temperature of the island increases and the resistance decreases. The Joule power V_{tes}^2/R_{tes}^2 , with constant voltage bias, goes down and compensates the increase in optical power. The corresponding output signal is the change in output current of the TES from its change in resistance, $I_{tes} = V_{tes}/R_{tes}$. This current is then read out by a SQUID ammeter as described in § 2.5.2. The TES is inductively coupled to the SQUID with total inductance L_{tot} .

This self-regulation is the negative electrothermal feedback (ETF) that results from a voltage-biased TES [28]. It has many advantages that will be discussed below. The voltage bias is implemented by placing the TES in parallel with a shunt resistance R_{sh} chosen so that $R_{sh} \ll R_{tes}$. The standard bias circuit of the TES is shown in figure 2.1 along with the SQUID readout.

The TES response is governed by equation 2.2 along with the electrical differential equation [30]:

$$L_{tot} \frac{dI}{dt} = V_{tes} - IR_{sh} - IR_{tes} \quad (2.3)$$

These two equations are coupled differential equations for the thermal and electrical circuits of the TES. The TES resistance is dependent on both its temperature and the applied current. The resistance versus temperature profile is only a constant current slice of a surface $R_{tes}(T, I)$ that describes the TES resistance. The entire performance and noise properties of a TES can be determined from these differential equations. The most important for the present discussion are parameters that effect the stability and responsivity of the devices. The results for the noise are given in § 4.1.6.1 in context of fitting a model to the measured detector noise.

The eigenvalue solutions to equations 2.2 and 2.3 describe the effective rise time (τ_+) and fall time (τ_-) of a TES to a delta-function energy impulse [30]. These time constants involve combinations of the standard electrical and thermal time constants, $\tau_{el} = \frac{L_{tot}}{R_{dyn}}$ and $\tau_{nat} = \frac{C}{G}$, that have been

enhanced by a loop gain from ETF. The low-frequency loop gain is defined as $\mathcal{L}_I \equiv \frac{P_J \alpha_I}{G T_{tes}}$, where the subscript denotes constant current. In the limit of small L_{tot} and $\tau_+ \ll \tau_-$, when the poles are widely separated and don't interact, the rise and fall times reduce to τ_{el} and τ_{eff} , respectively. The latter time constant, τ_{eff} , is the effective thermal time constant or the natural time constant enhanced by \mathcal{L}_I and the electrical circuit parameters.

The time constants τ_{\pm} impact the stability of the TES. When the two poles converge they interact and give rise to thermal oscillations in the TES. This in turn has an impact on the selection of the inductance in the readout circuit and the design parameters C and G of the TES. It also impacts how they are biased and operated. The TESs can be damped or oscillating, either stable or unstable with thermal oscillations that grow with time and eventually drive the TES out of the transition. The TES will be critically damped if the two time constants are equal and will be overdamped if $\tau_+ < \tau_-$. A voltage-biased TES is always stable under these conditions as long as $R_{tes} > R_{sh}$ [30]. Given $R_{sh} \ll R_{tes}$ this requirement is met for a wide range of TES operating resistance, all but very low in the transition. The damping requirement can be met by a careful choice of inductance given the desired TES parameters. For fixed circuit and TES parameters the time constants then depend mostly on R_{tes} for a relatively uniform transition, so stability then depends on the bias point in the transition. As the TES is biased lower into the transition the larger τ_- decreases while the smaller τ_+ increases. Eventually low in the transition the time constants converge and the TES becomes underdamped. At that point the TES is stable when $\tau_{nat} > (\mathcal{L}_I - 1)\tau_{el}$ or when $\mathcal{L}_I > 1$ and $L < \frac{\tau_{nat}}{\mathcal{L}_I - 1} \cdot (R_{sh} + R_{tes}(1 + \beta_I))$ [30].

The power-to-current responsivity — the ratio of the current signal to change in optical power — is also derived from equations 2.2 and 2.3. In the case of voltage bias and strong thermal feedback from a high loop gain, the DC power-to-current responsivity is:

$$s_{I,dc} = -\frac{1}{I_{tes}(R_{tes} - R_{sh})} \quad (2.4)$$

which is approximately the reciprocal of the voltage bias V_{tes} for $R_{tes} \gg R_{sh}$. In this limit, the DC responsivity of the detector doesn't depend on the TES parameters only on the bias circuit parameters. Negative ETF suppresses the effect of fluctuations in device parameters that might lead to gain fluctuations between channels. In addition to this “self-calibrating”, TESs also “self-bias” into the superconducting transition through negative ETF when the bath temperature is much lower than T_c . Sensitivity to bath temperature fluctuations is also suppressed for negative ETF when $T_{bath} \ll T_c$. This is particularly important since differential temperature fluctuations in a detector

Table 2.2: BICEP2 bolometer design parameters for the titanium and aluminum TESs

TES	T_c K	R_n Ω	P_{sat} pW	G_c pW/K	β	C pJ/K	α
Ti	0.45	0.06	9	80	2.5	0.4	100
Al	1.3	0.5	600	1760	2.8	0.4	30

pair produce false polarization measurements.

Negative ETF also increases the bandwidth of the TESs. This speed enhancement is particular useful for scanning experiments such as BICEP2 in which the science band is projected to higher frequencies. The natural time constant of the TES may not provide sufficient bandwidth and may otherwise filter some of the science signal.

The designed TES parameters are shown in table 2.2 for the two TESs used in BICEP2 bolometers, assuming a bath temperature of 250 mK. The T_c is determined by the transition temperature of the superconducting film. The R_n — and the R_{tes} in operation — is determined by the geometry and properties of the superconducting film. The choice of T_c determines the operating temperature and thus the thermodynamic noise of the detectors. A high-sensitivity device should use as low a T_c as allowed by the cryogenic cooling system, so that $T_c \gg T_{bath}$.

The choice of G is set by a balance between the requirements of noise and optical loading. Apart from the noise from the arrival rate of photons, thermal fluctuation noise is the dominant source of detector noise power. Thermal fluctuation noise power is proportional to G (§ 4.1.6.1), so better detector sensitivity calls for lower conductance. However, the amount of optical power that the bolometer can accept without saturating is set by $P_{sat} = P_{bath} \propto G$ from equations 2.1 and 2.2. An additional safety factor is generally included in the power to account for observing during periods of poor weather and higher loading. The P_{sat} is thus set to the product of the safety factor and typical loading conditions and the G is determined from equation 2.1 for an assumed β . The exponential dependence of G on β can be used to refer the conductance to a given temperature. In table 2.2 the thermal conductance is quoted at the transition temperature G_c . The values in the table are derived from $P_{sat} = 9$ pW for an assumed total optical load of ~ 3.7 pW and a safety factor of 2.5. The G_c is designed for the titanium device and then scaled up to conductance at T_c for the aluminum device. The conductance exponent changes from $\beta = 2.5$ to $\beta = 2.8$ for the aluminum due to the dependence on β with temperature over that range. Note the factor of ~ 77 increase in the saturation power of the aluminum TES over the titanium device, which makes it ideal for high loading conditions.

The heat capacity C is set by the materials deposited on the island. It plays an important role in the stability of the detector through τ_{nat} . In order to improve the margin on stability — increase the separation between τ_{\pm} — additional heat capacity can be added to the bolometer to slow it down and increase τ_{-} or τ_{eff} relative to τ_{+} or τ_{el} .

2.2 cryostat and thermal design

The BICEP2 receiver is cryogenically cooled to provide a low operating temperature for the detectors and other superconducting components, and to cool the optics to a low and stable temperature. The receiver is cooled to a nominal base temperature of 4 K by the evaporation of liquid helium.¹ The focal plane is cooled to a base temperature of ~ 300 mK by a three-stage helium sorption refrigerator operating from the 4 K base. This temperature is well below the superconducting critical temperatures of the TESs and SQUIDs. The focal plane is mechanically isolated and thermally controlled to prevent temperature fluctuations that would produce false polarization signal. Superconductors are used to shield the TESs and SQUIDs from magnetic pickup and must also be cryogenically cooled below their critical temperatures. The cryostat and receiver architecture is shown in figure 2.2.

During typical CMB observing conditions the receiver consumes 22 L of helium per day. The liquid helium level is monitored with a cryogenic liquid level monitor and the flow rate of helium gas from the cryostat. A superconducting-filament probe orientated parallel to the cryostat axis gives an accurate measure of the liquid level. The level is corrected in software for the change in level with the elevation tilt of the mount. The boil-off of the liquid helium is monitored by a flow meter to give an independent measure of the helium level, which is in good agreement with the corrected liquid level. Due to the expense and the painstaking of transport and management of liquid helium, BICEP2 is scheduled to be the last liquid-helium-cooled cryostat operated at the South Pole. The transition to using mechanical cryo-coolers to reach 4 K base temperatures has already been made by the Keck Array and South Pole Telescope.

2.2.1 cryostat

The optics chain and focal plane are housed in an evacuated, helium-cooled cryostat. The cryostat was designed in cooperation with and built by Redstone Aerospace.² It consists of an outer vacuum shell and toroidal liquid helium storage vessel separated by two vapor-cooled shields (VCSs), which

¹The mean pressure at the South Pole is 681 mb, so the base temperature is actually ~ 3.7 K.

²www.redstoneaerospace.com

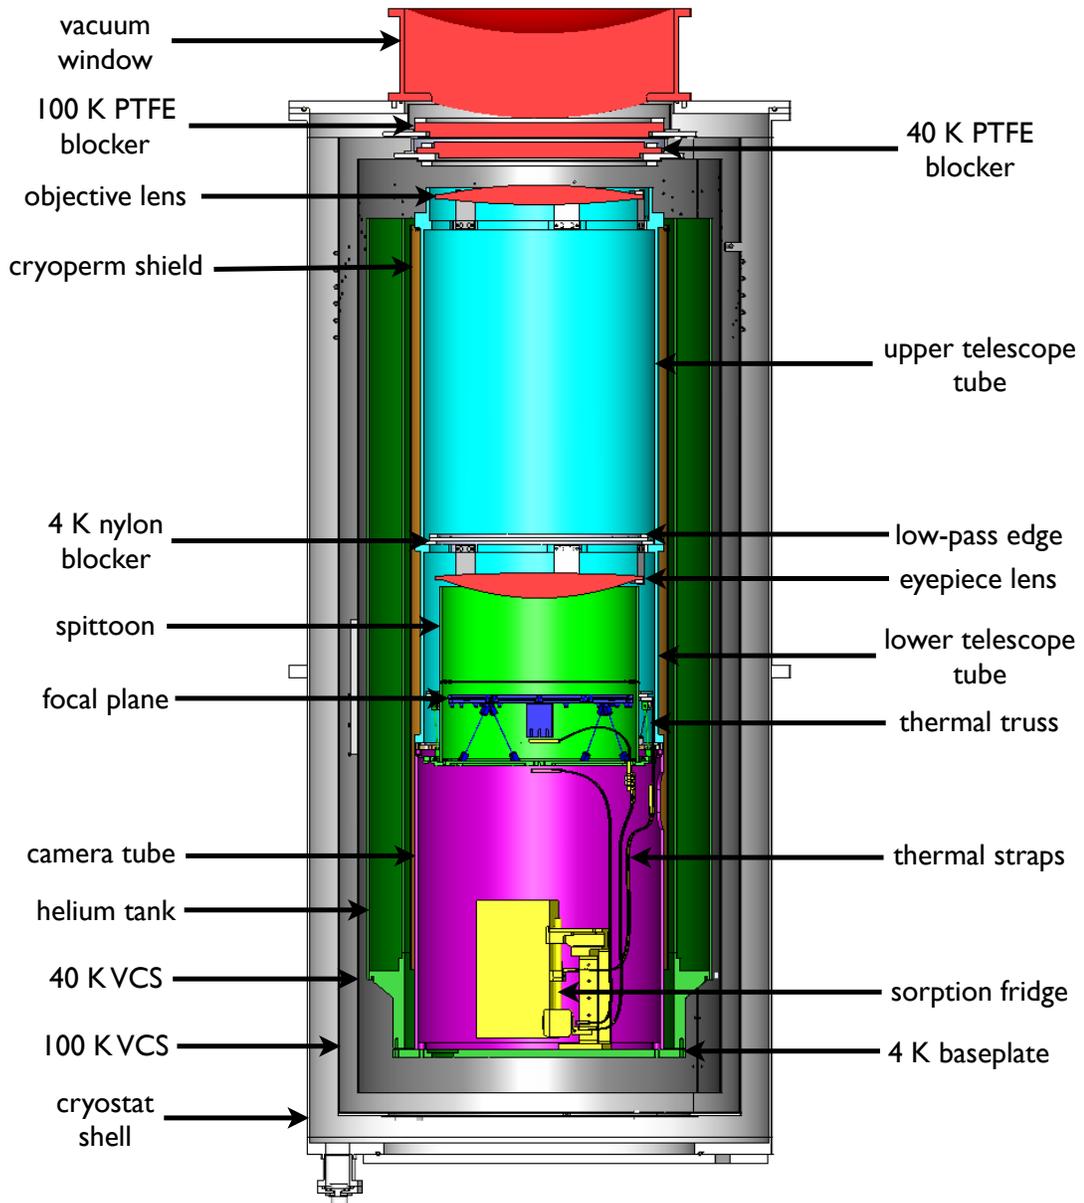


Figure 2.2: A cross-sectional view of the BICEP2 receiver. The cryostat includes the vessel shell (white), outer 100 K VCS (light gray), the inner 40 K VCS (dark gray), and the 4 K helium reservoir (dark green). The cryostat insert sits inside the helium reservoir and includes the camera tube (magenta) and lower and upper telescope tubes (cyan). The insert is also surrounded by a 4 K Cryoperm magnetic shield (orange). The sorption fridge and the heat straps (both shown in yellow) connecting to the sub-Kelvin stages (green and blue) are shown. The 1 K spittoon and two-stage thermal truss are shown (light green), with the 300 mK focal plane and thermal components (dark blue) inside. The optical elements (red) include the cryostat window, 100 K and 40 K teflon VCS filters, the objective lens, 4 K low-pass edge and nylon filters and the eyepiece lens.

reduce the radiative heat load on the storage vessel. The VCSs are cooled by the helium boil-off vapor leaving the vessel, which is routed through coils of small-diameter tubes wrapped around the shields. The VCSs were designed to achieve 100 K and 40 K temperatures under normal cryostat loading and boil-off. To further reduce the radiative load, the spaces between the shell, shields and vessel are filled with low-emissivity, high-reflectivity multilayer insulation (MLI). The conductive heat load between stages is reduced through the use of rigid, low-conductivity G-10 standoffs. The VCSs replace the liquid nitrogen cooling stage used in the BICEP1 cryostat. Liquid nitrogen is only used to precool the liquid helium reservoir during initial cool down of the cryostat. The designed helium consumption rate is ~ 20 L per day during normal observation conditions.

The outer vacuum shell is cylindrical, measuring 27.6 cm in diameter and 62.6 cm in length. The ends are closed off using aluminum plates with o-ring seals. The cryostat is fixed in the telescope mount by a flange located roughly at its midpoint. The upward-looking end of the cryostat is closed off with a plate containing the transparent vacuum window that allows radiation to enter the cryostat. The plumbing and electrical feedthroughs are located at the downward-looking end next to the rear closeout plate and are easily accessible when the cryostat is in the mount. The feedthroughs do not exceed the diameter of the outer cryostat wall to prevent accidentally damaging them when the cryostat is inserted in the mount.

A cylindrical insert containing the focal plane, 4 K optical elements and refrigerator is located inside the helium vessel toroid. The base of the insert is a 431 mm diameter, 12.7 mm thick oxygen-free high-conductivity (OFHC) copper disk. It bolts directly to the base of the helium vessel and provides an isothermal 4 K heat sink for the insert. The top of the insert is constrained by standoffs attached to the top of the helium vessel to provide rigid alignment of the insert along the boresight axis. The standoffs are made of low-conductivity G-10 in order to reduce the temperature gradient between the base and top of the insert caused by the warmer vessel top.

2.2.2 sub-Kelvin refrigerator

The refrigerator is a closed-cycle, three-stage ($^4\text{He}/^3\text{He}/^4\text{He}$) sorption refrigerator provided by collaborators at CEA-Grenoble.³ Each stage is composed of a closed system including an adsorbent pump, a still to collect condensed liquid helium and a low-conductivity pump tube connecting the two. The first stage contains ^4He , and the second and third stages contain ^3He , which has a lower condensation temperature. The first and second stages together are referred to as the intercooler

³The group is led by Lionel Duband.

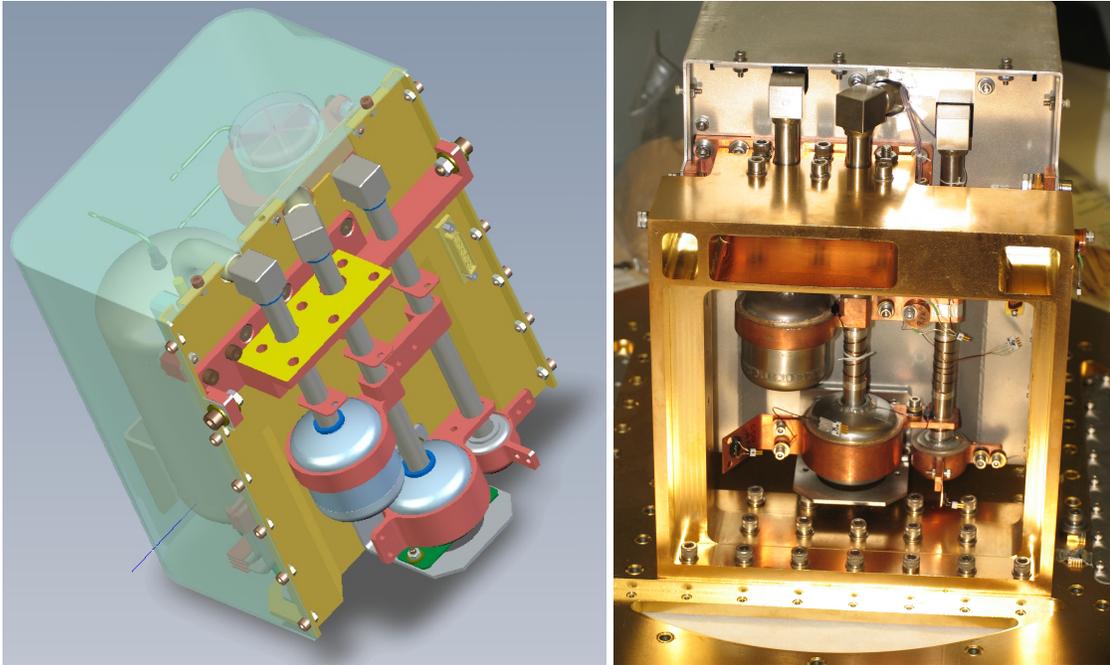


Figure 2.3: The three-stage He sorption fridge. *Left*: A rendering of the sorption fridge with the aluminum shield transparent to show the pumps it shields, and the stills and pump tubes on the front side. *Right*: An image of the fridge installed in the fridge bracket and mounted to the 4 K base plate.

(IC). The third stage is referred to as the ultracooler (UC). The pumps are cooled using gas-gap heat switches connected to the 4 K base plate. They are housed in an aluminum box, shown in the transparent view of figure 2.3, to prevent them from radiatively heating the insert when turned on. The fridge is supported by a gold-plated OFHC fridge bracket that bolts to 4 K base plate. The operation of a similar three-stage sorption refrigerator is explained elsewhere [5]. The fridge was designed with the following specifications: a minimum hold time of 48 hours, a maximum recycle time of 4 hours, an IC ${}^4\text{He}$ temperature of $\lesssim 2.8\text{ K}$ under a $45\ \mu\text{W}$ load, an IC ${}^3\text{He}$ temperature of $\lesssim 340\text{ mK}$ under a $25\ \mu\text{W}$ load, and a UC ${}^3\text{He}$ temperature of $\lesssim 250\text{ mK}$ under a $1.5\ \mu\text{W}$ load. Under typical loading conditions the UC ${}^3\text{He}$ reaches $\sim 250\text{ mK}$, the IC ${}^3\text{He}$ reaches $\sim 300\text{ mK}$ and the IC ${}^4\text{He}$ reaches $\sim 1\text{ K}$. The fridge is easily cycled within 4 hours and has a hold time nearly double the specification at ~ 80 hours during standard observing conditions. This has allowed for an improvement in the observing efficiency of BICEP2 (§ 3.10).

2.2.3 sub-Kelvin thermal architecture

The sub-Kelvin structure is supported by a set of thermal trusses that reduce the parasitic loads while maintaining rigid optical alignment. The thermal straps that connect the sub-Kelvin thermal

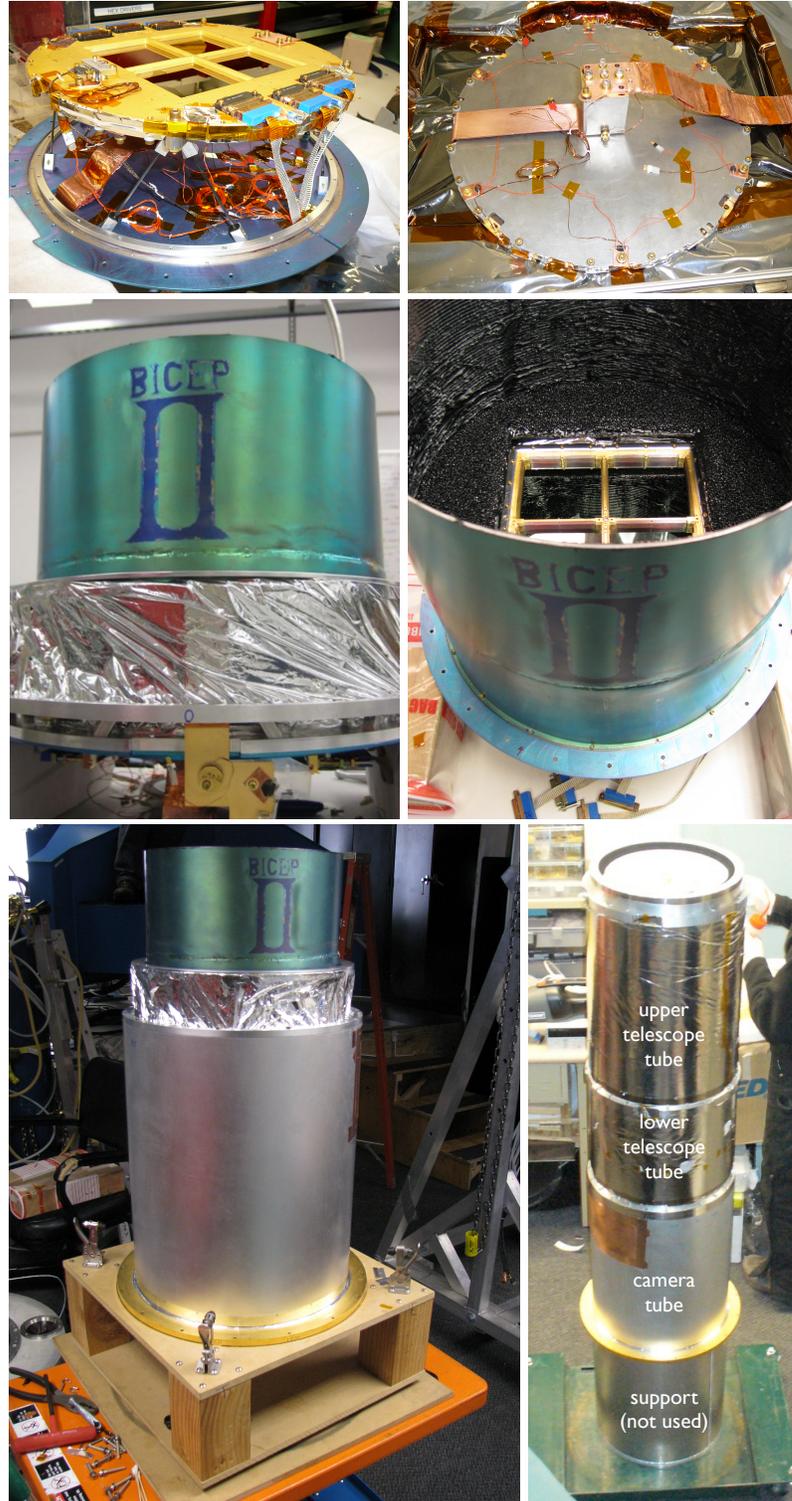


Figure 2.4: The components of the sub-Kelvin structure and 4 K insert. Moving through the images left-to-right, top-to-bottom: The focal plane connected to the spittoon bottom by carbon fiber posts; The backshort with the filter block and UC thermal strap; The spittoon integrated with the two-stage thermal truss (obscured by MLI), note two levels of the truss bottom; The interior of the blackened spittoon and focal plane apertures; The sub-Kelvin structure mated with the camera tube; The full 4 K insert with camera tube and lower and upper telescope tubes.

components to the fridge have been designed to minimize microphonic dissipation at the focal plane. Furthermore, the temperature of the focal plane is controlled by two temperature control servos to minimize scan-synchronous thermal drifts that mimic optical signals in the TESs.

The sub-Kelvin structure includes the focal plane, a superconducting shield called the spittoon and standoffs for thermal isolation. The sub-Kelvin structure mounts to the top of the camera tube, which is a thick aluminum tube that surrounds the fridge and bolts to the 4 K base plate. The focal plane is cooled to ~ 300 mK, the spittoon is cooled to ~ 400 mK and the entire sub-Kelvin structure must be attached to the 4 K camera tube. Carbon fiber posts are used to rigidly attach the different components to one another while minimizing the load on the lower temperature stages.

The spittoon surrounds the focal plane and provides magnetic shielding and a low-temperature surface on which to terminate any spillover of the beams. The spittoon has two parts: an open cylinder of niobium and a niobium base that attaches to the end of the cylinder opposite the illuminated end. The focal plane is mounted to the base using four sets of standoffs each containing two thin carbon fiber posts. The spittoon cylinder is lowered over the focal plane and mounted to the base, creating a nearly continuous superconducting enclosure. A niobium platform has been added just above the focal plane to improve the magnetic shielding. The platform has a cutout that matches the aperture of the focal plane. The interior of the spittoon is blackened with microwave absorber to terminate the beam spillover.

The sub-Kelvin structure is attached to the camera tube using a two-stage thermal truss. The truss has three independent flat aluminum annuli that are connected by thin carbon fiber posts. The upper annulus connects to the other two annuli, which do not connect to one another. This creates a long thermal path between the two lower annuli in a relatively compact space. One of the lower annuli connects to the spittoon while the other connects to the top of the camera tube, as shown in figure 2.4.

The two-stage truss is cooled to ~ 1 K by the IC ^4He still. The spittoon is cooled to ~ 400 mK by a thermal link to the IC ^3He still. The focal plane is cooled to ~ 250 mK (without thermal control) by a thermal link to the UC ^3He still. The thermal connections are made using flexible heat straps composed of multiple layers of OFHC copper sheets. The ends of the straps are terminated in gold-plated OFHC copper blocks that are electron beam welded to the sheets. To prevent microphonic vibrations in the cryostat from depositing thermal energy in the straps, each heat strap is restrained by a G-10 and carbon fiber support assembly. Thermal loading is further reduced by using niobium-titanium wiring so that the thermal conduction is limited to the phonon conduction

of the superconductor.

The focal plane is temperature controlled to 280 mK during observations, buffering it around 30 mK above the base temperature of the fridge. This provides a stable bath temperature for the TESs and SQUIDs well below their superconducting transition temperatures. Fluctuations in the TES bath temperature are controlled in two ways. A passive thermal filter buffers the focal plane from fluctuations originating in the fridge or 4 K thermal bath. The filter is a $5.5 \times 2.5 \times 2.5$ cm block of 316 stainless steel that is bolted to the focal plane backshort, as shown in figure 2.4. The UC heat strap is connected to the opposite end to cool the focal plane through the 5.5 cm dimension of the block. The stainless steel combines high heat capacity and low thermal diffusivity to serve as a filter with high conduction and a long time constant. It filters the focal plane from thermal fluctuations faster than ~ 1300 s [45].

In addition, the FPU temperature is actively maintained by two temperature control servos. A high-bandwidth servo is maintained at the connection between the strap and filter block to control broadband temperature fluctuations. A low-bandwidth servo is maintained on the focal plane base to control slow drifts in temperature. No active thermal control is performed on the focal plane within the 0.1–1 Hz science band. The servo is performed by the housekeeping hardware (§ 2.9.3) using temperature control modules (TCMs) that include redundant pairs of high-sensitivity neutron-transmutation doped (NTD) thermistors and resistive heaters.

2.3 optics

The multipole range $\ell = 30 - 300$ allows for a small 26 cm aperture to be used, with all the benefits of a compact refractor described above. The cryostat houses the entire refractive optics chain. The refractive optical elements are cooled to provide low, stable loading of the detectors. The optics have been designed for telecentricity and low aberration in the focal plane in order to control polarization artifacts. The optical elements have been carefully chosen so as not introduce spurious polarization or birefringence into the optics chain. A combination of infrared blocking filters are used to reduce the radiative load on the cryogenic stages. The receiver is monochromatic and the optics have been optimized for the single passband centered at 148 GHz. This improves the efficiency of the antireflection (AR) coatings and reduces aberration. The bolometers are coupled to phased-array slot antennas that form raw $14'$ beams. The small aperture and $2f\lambda$ detector spacing give a large field of view with high throughput. As shown in figures 2.2 and 2.5 the optics chain consists of a vacuum window, three infrared blocking filters, a low-pass edge filter, refractive lenses, and the

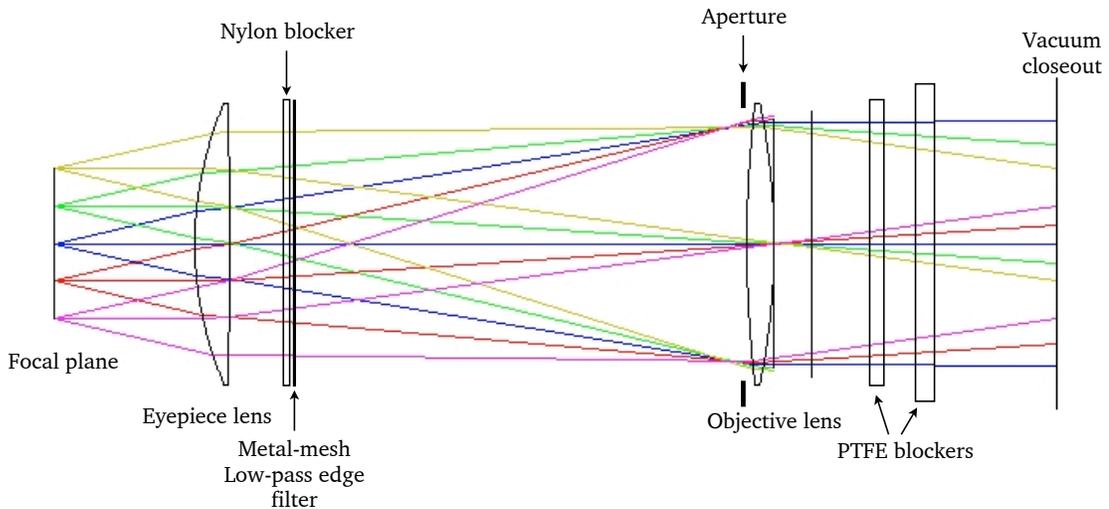


Figure 2.5: A schematic of the optics chain. This image was made by R. Aikin and is taken from the BICEP2 optics paper [2].

phased-array antennas at the focal plane.

The 4 K optical elements are aligned and cooled by the 4 K telescope insert, shown in figure 2.2. The insert has three independent sections: the camera tube, the upper telescope tube and the lower telescope tube. The camera tube includes the sub-Kelvin structure and bolts directly to the 4 K base plate. The lower telescope tube bolts on top of the camera tube and includes the eyepiece lens and the nylon and low-pass filters. The upper telescope tube bolts to the top of the lower telescope tube and includes the eyepiece lens and aperture stop. The tubes are made from aluminum, which has adequate thermal conductivity to keep the entire insert cooled to a stable 4 K temperature. The interior of the telescope tubes is baffled with a 10 mm thick layer of Eccosorb HR-10 epoxied into place using Stycast 2850.⁴

2.3.1 vacuum window

Radiation enters the cryostat through a low-loss, 32.5 cm diameter vacuum window at the upward-looking end of the cryostat. The window is formed from a closed-cell, nitrogen-expanded polypropylene foam, Propozote PPA30.⁵ Four layers of PPA30 are heat-laminated by the manufacturer for a total thickness of ~ 13 cm. The outer rim of the foam cylinder is bonded to a 41 cm diameter aluminum housing using Stycast 1266. The aluminum housing is mounted to the outside of the top

⁴Both products are from Emerson & Cuming Microwave Products: www.eccosorb.com.

⁵www.zotefoams.com

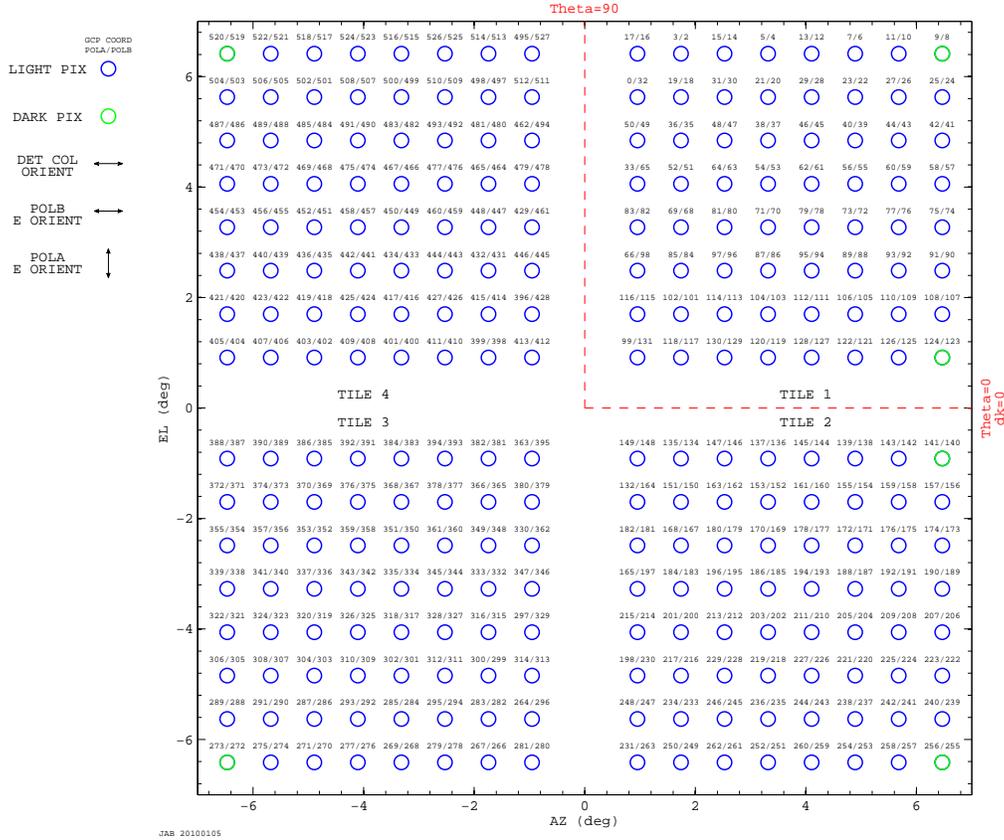


Figure 2.6: Layout of the BICEP2 beams with nominal locations, beam widths and polarization orientations. This is the projection of the channel centers relative to the telescope pointing onto a flat sky at a boresight angle of 0° . The circles represent the nominal 0.5° beam widths of the detectors. The tile boundaries, light and dark channels and polarization orientations are also shown. The tile wire bonds are along the left and right edges. Columns of detectors in an array are parallel to the top edge.

closeout plate with an o-ring seal. The window is expected to radiate a minimum of ~ 4 W into the cryostat [61]. This level of loading is far too high for the sub-Kelvin stages, so a series of cooled filters is used to intercept the load. Transmission through the window has been measured to be $>98\%$ using a chopped thermal source. The moisture in the warmer environment forms a layer of frost on the cold exterior surface of the window. To prevent this, a thin membrane is used to create a warm, dry space above the window. The membrane is stretched tight in a frame and excess slack in the film is prevented from vibrating by purging the space with nitrogen gas at 0.35 psi above ambient pressure. A set of holes blow warm air from the control room over the top of the membrane to sublimate snow accumulated on its surface. An infrared camera mounted near the window is used to monitor residual snow accumulation.

2.3.2 infrared blocking filters

To reduce out-of-band radiative loading of the focal plane, a system of absorptive dielectric infrared blockers are used. They are designed for high in-band transmission, high out-of-band absorption and good thermal conductivity. The filters must transmit as much of the in-band radiation as possible to ensure adequate optical efficiency. This adds a corollary requirement that they have a low refractive index or are easy to add an antireflection coating to in order to reduce loss from reflections. They must also absorb as much out-of-band radiation as possible to reduce optical loading of the detectors and the sub-Kelvin stages. The decreased loading reduces the photon noise and allows lower bolometer conductivity for lower thermal fluctuation noise. It also lowers the heat load that degrades cryogenic hold time and observing efficiency. The heat absorbed by the filters must be sunk to the cold cryostat to prevent the filters from heating and reradiating on the colder stages. So, the filter material must have adequate thermal conductivity.

Teflon or polytetrafluoroethylene (PTFE) is an absorptive blocker with an appealing combination of the above requirements and excellent transparency at millimeter wavelengths. The BICEP2 filter stack builds on the successful design implemented in BICEP1, with one modification. An additional absorptive blocker has been added at 4 K to further reduce loading of the focal plane. A 4 K PTFE filter would do a poor job of blocking infrared radiation, so a nylon filter has been used instead [4]. Two PTFE filters are mounted to the ends of the 100K and 40K VCSs near the vacuum window. The former filter is 1.3 cm thick and the latter is 3.4 cm thick. A 5.2 cm thick nylon filter is mounted to the top of the 4 K lower telescope tube just above the eyepiece lens.

2.3.3 low-pass filter

The infrared blockers do not filter high-frequency radiation leaks. To mitigate an observed response of the detectors to radiation above the pass band (§ 5.6.3), a low-pass filter is included in the optics chain. The filter is a capacitive-grid, low-pass edge filter with a cutoff wavenumber of 8.3 cm^{-1} provided by collaborators at Cardiff University [1].⁶ High-frequency radiation above the 249 GHz cutoff frequency reflects from the filter and is not transmitted to the focal plane. The filter is made by hot-pressing layers of metal-mesh and dielectric substrate. The metal-meshes have different characteristic impedances to create a filter with a sharp cutoff and low in-band ripple. The 300 mm diameter filter is mounted to the 4 K lower telescope tube just above the nylon blocker.

⁶The group is led by Professor Peter Ade.

2.3.4 lenses

Two cooled, biconvex lenses are used to re-image the focal plane at infinity. The design is based on the optics used in BICEP1, with several optimizations. The lenses are optimized for the single 148 GHz band, allowed to be asymmetric with respect to the curvature of the two surfaces, and the eyepiece lens is optimized for telecentricity rather than aberration at the image-plane. The design and optimization of the lens separation and curvature is described in detail in the BICEP2 optics paper [2].

The lenses are made from high-density polyethylene (HDPE) for manufacturing ease and due to their low reflective loss. They are constrained by the insert, which provides rigid optical alignment and cools the lenses to 4 K. The diameter of the lenses is limited to ~ 310 mm by the interior diameter of the insert tube. The optical path between the objective and eyepiece lenses is 550 mm or 563 mm physical separation. The optical path between the focal plane and eyepiece lens is 150 mm or 153 mm physical separation.

The 550 mm focal length and 264 mm diameter aperture give BICEP2 an f-number of $f/2.1$ and plate scale of $0.1^\circ/\text{mm}$. The physical layout of the antennas has been determined from the detector array lithography mask and a solid model of the focal plane. The antenna locations include the effect of thermal contraction of the focal plane. The plate scale has been combined with the physical model to determine the ideal pointing of the detector beams, as shown in figure 2.6. The instantaneous field of view along an edge of the square focal plane is $\sim 13^\circ$ or $\sim 18^\circ$ corner-to-corner.

The lenses are mounted to the insert using six OFHC copper flexures. The flexures are flexible enough to account for the differential thermal contraction of the lenses and insert tube, but sufficiently rigid to reliably align the lenses in the optics chain. They provide an adequate thermal path to the insert tube. The eyepiece and objective lenses are mounted at the top of the lower and upper telescope tubes, respectively.

2.3.5 antireflection coatings

Antireflection (AR) coatings are used on the filters and lenses to reduce reflective loss. This increases the throughput of the optics system and also reduces crosstalk from beams reflected back to the focal plane. Each surface is coated with a single membrane of expanded PTFE, which has a low index of refraction of ($n = 1.2$), low absorption at 150 GHz and only 0.8% loss per surface. A 406 μm thick membrane of Mupor⁷ PM23D is used on the lenses and PTFE filters, and a 381 μm membrane of

⁷Interstate Specialty Products: www.interstatesp.com

Mupor PM23J is used on the nylon filter. The AR coat is heat-bonded to the optical components using a 51 μm film of low-density polyethylene (LDPE) with a melting temperature of 120 C, which is well below the melting point of the lenses and filters.

2.3.6 cold aperture stop

A 264 mm diameter cold stop is used to define the aperture size of the telescope and terminate spillover of the beams. The stop is located on the interior side of the objective lens, as shown in figure 2.5. It is made from a 3 cm thick annulus of microwave-absorbing Eccosorb HR10. The interior of the annulus is tapered at $\sim 45^\circ$ to a final 264 mm diameter to decrease diffraction at the aperture. The stop is bonded to the objective lens with Stycast epoxy and cooled by the lens to a stable 4 K.

2.4 focal plane unit

The heart of the BICEP2 polarimeter is the focal plane unit, which uses novel antenna-coupled TES bolometer arrays to measure 2 mm radiation. The large detector count (528 channels) and parasitic loading requirements make a SQUID multiplexer ideal for reading out the TESs. The focal plane unit includes a 27 cm diameter isothermal copper detector plate, four detector tiles each with a z-cut crystal quartz antireflection (AR) tile, 16 sets of NIST-fabricated SQUID multiplexer and Nyquist interface chips, a printed circuit board for routing of electrical connections, beryllium-copper spring clips to fix the arrays and AR tiles to the detector plate, and a superconducting niobium quarter-wave backshort. The science-grade focal plane is shown in figure 2.7, open in a clean room at the South Pole for a final inspection just prior to integration in the telescope insert. A view of the fully assembled focal plane is shown in figure 2.8, which clearly shows the focal plane apertures on the illumination side and the backshort on the opposite side.

The focal plane is cooled to 250 mK and temperature controlled to 280 mK for normal observations. This provides a stable bath temperature for the TESs and SQUID multiplexer chips, well below their superconducting transition temperatures. Fluctuations of the TES bath temperature resemble false changes in input power to the antennas, so the temperature of the detector plate is carefully controlled.

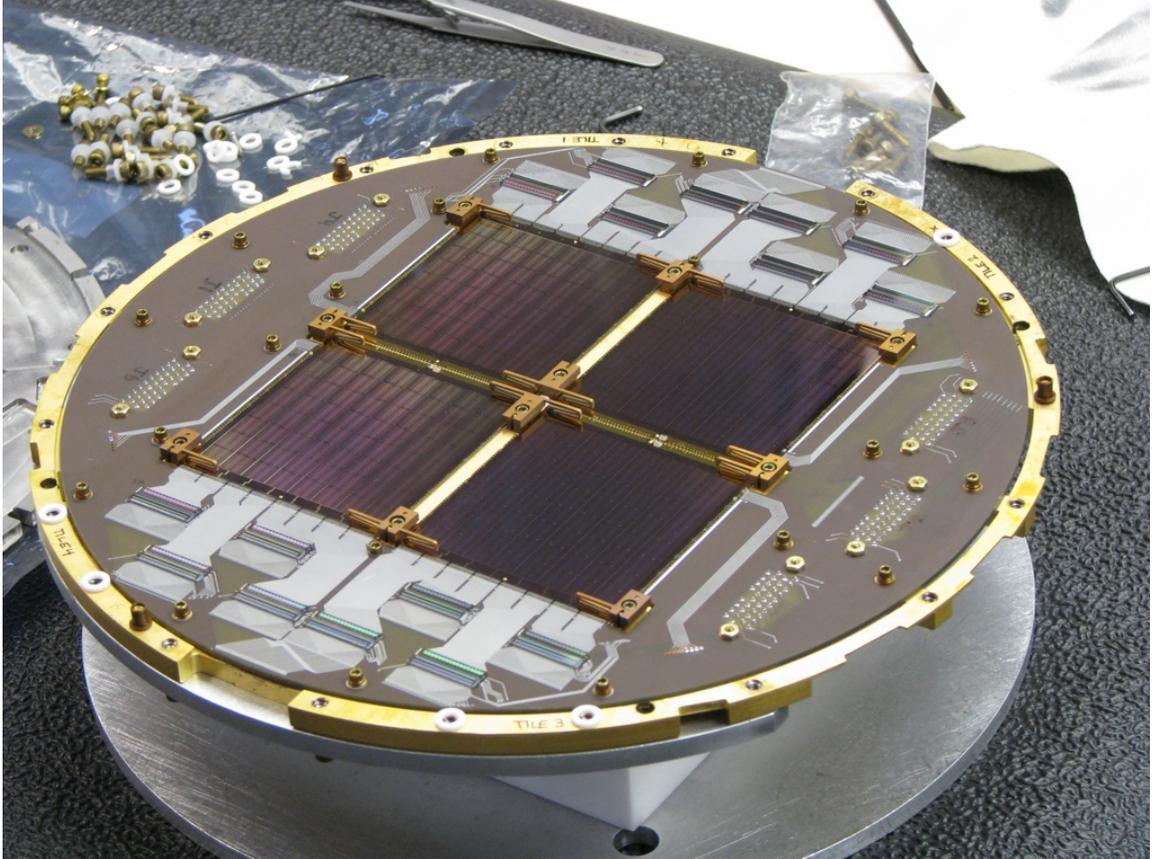


Figure 2.7: An image of the focal plane unit, lovingly hand built by the author. The four detector arrays are shown, clamped to the gold-plated detector plate by 10 tile clips. The circuit board is shown with superconducting aluminum traces leading from the arrays to 16 pairs of multiplexer and Nyquist chips. A few of the lapped spacers that separate the backshort from the arrays are shown. They are the white disks resting on the perimeter of the detector plate.

2.4.1 detector plate

The central support structure of the focal plane is an isothermal, gold-plated OFHC copper detector plate. The plate is a 27 cm diameter, 7.5 cm thick disk. The manufactured plate was required to be no more than 20 μm out of plane. This planarity specification ensures beams from the tiles are sufficiently parallel to one another and to the optical axis of the telescope. The plate is thick enough so that with proper handling this specification can be met by vendors and maintained during integration. The OFHC copper provides adequate conductivity to ensure the plate is an isothermal heat sink for the detector tiles and circuit board. The detector plate also provides mechanical registration and threaded attachment points for most of the focal plane components.

Four rectangular aperture windows are cut out of the detector plate for illumination of the detector arrays. These windows were made as large as possible to reduce distortion of the beams due

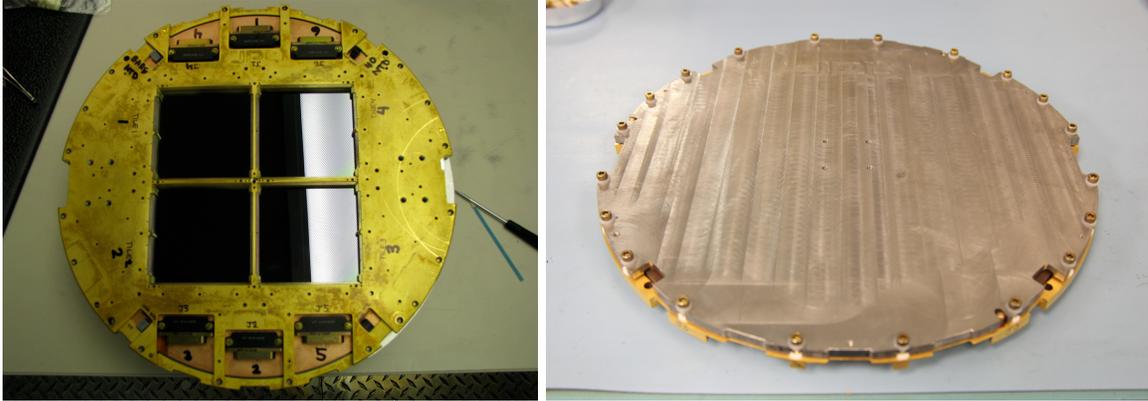


Figure 2.8: The focal plane unit fully assembled and ready for installation. Left: An image of the illumination side of the focal plane, showing the gold plated copper detector plate, the illumination side of the detector arrays with antireflection coating, and connectors. Right: An image of the superconducting backshort, which is precisely separated from the detector plate using lapped spacers (white) around the perimeter of the focal plane.

to interaction with the metal aperture edges. Due to design restrictions, the separation between the antenna patches of the edge pixels and aperture edge couldn't be increased above ~ 2.5 mm. Electromagnetic simulations showed this proximity would differentially affect the orthogonal polarizations of the edge pixels. To mitigate this effect corrugations were added to the apertures, running parallel to the plane of plate as shown in figure 2.9. The corrugation geometry was informed by electromagnetic simulations and selected to be 0.25 mm thick on a 0.50 mm pitch with a 0.60 mm depth.

2.4.2 detector arrays

Radiation is measured using planar, antenna-coupled TES arrays [39]. The detector arrays are lithographically fabricated on 4" silicon wafers in the Micro Devices Laboratory at the Jet Propulsion Laboratory [8]. Each detector tile includes an 8×8 square array of dual-polarization focal plane pixels, a superconducting wiring bus to bring the detector output to wire bond pads along one tile edge, a border of gold along the remaining three edges to promote thermal conduction, and a set of through and slotted holes for alignment of the tile in the focal plane. The detector tiles have a rectangular format with 76 mm and 72 mm sides. The rectangular arrays are cut from the circular wafers using a deep trench etching step. A schematic of a detector tile and its major components are shown in figure 2.10. The detectors are separated by $2f\lambda$ with a physical spacing of ~ 7.8 mm. The bolometers have a dual-transition TES design, using both a titanium and an aluminum TES. The detectors can be biased on the titanium transition for low-loading operation, or the aluminum



Figure 2.9: A close-up image of the illumination side of the focal plane unit, superconducting niobium center spacer screw and corrugated detector windows. Newton's rings are barely visible in several of the dark detector windows. They are created by small variations in the spacing between the antireflection and detector tiles.

transition under high-loading conditions.

The silicon wafers are provided by the manufacturer with a $1.0\ \mu\text{m}$ layer of silicon nitride (SiN) on each side and have a total thickness of $515\text{--}530\ \mu\text{m}$. Four additional layers are deposited by the JPL fabrication team: a $0.15\ \mu\text{m}$ silicon dioxide (SiO₂) protect layer is deposited on top of the front-side SiN, a $0.15\ \mu\text{m}$ layer of niobium is deposited on top of the protect layer to form the ground plane, a $0.3\ \mu\text{m}$ (SiO₂) interlayer dielectric (ILD) is deposited on top of the ground plane, and finally a $0.4\ \mu\text{m}$ layer of niobium is deposited to form microstrips. The antennas, wiring layer and bolometer elements are patterned on the front side of the wafers. The wafers are illuminated from the backside through the silicon nitride and bare silicon. The ground plane is anodized to prevent electrical shorts from developing between the two niobium layers due to holes created in the ILD by contaminants during fabrication. This arrangement has the added benefit that the summing network for the antenna is patterned on the opposite side of the ground plane and is protected by incoming radiation.

Each focal plane pixel includes a pair of independent planar antenna arrays for beam formation and coupling to incident microwave radiation, microstrip networks to sum power from each sub-antenna coherently, inline microstrip filters for spectral band definition, and TES bolometer islands to measure the absorbed power. The two independent, colocated antenna arrays measure orthogonal

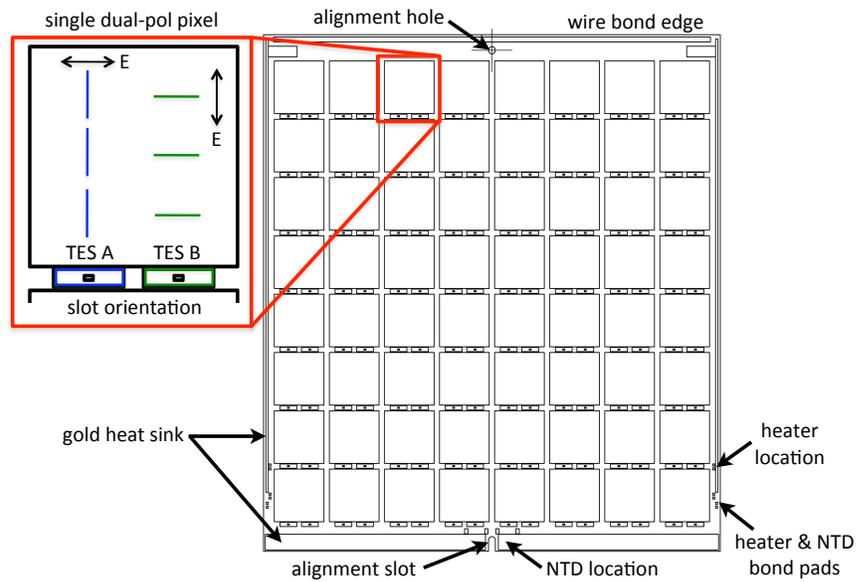


Figure 2.10: A schematic of a single detector tile with polarization orientation of the two bolometers. The tile includes an 8×8 array of focal plane pixels, superconducting wiring bus to bring the bolometer signals to wire bond pads, a gold border for thermal conduction and a set of through and slotted alignment holes. The schematic for the polarization sense shows only three slots, but in reality each pixel has a 12×12 array of orthogonal slots.



Figure 2.11: A micrograph (100x magnification) of a portion of a focal plane pixel, including the slot antenna, microstrip line and detector island.

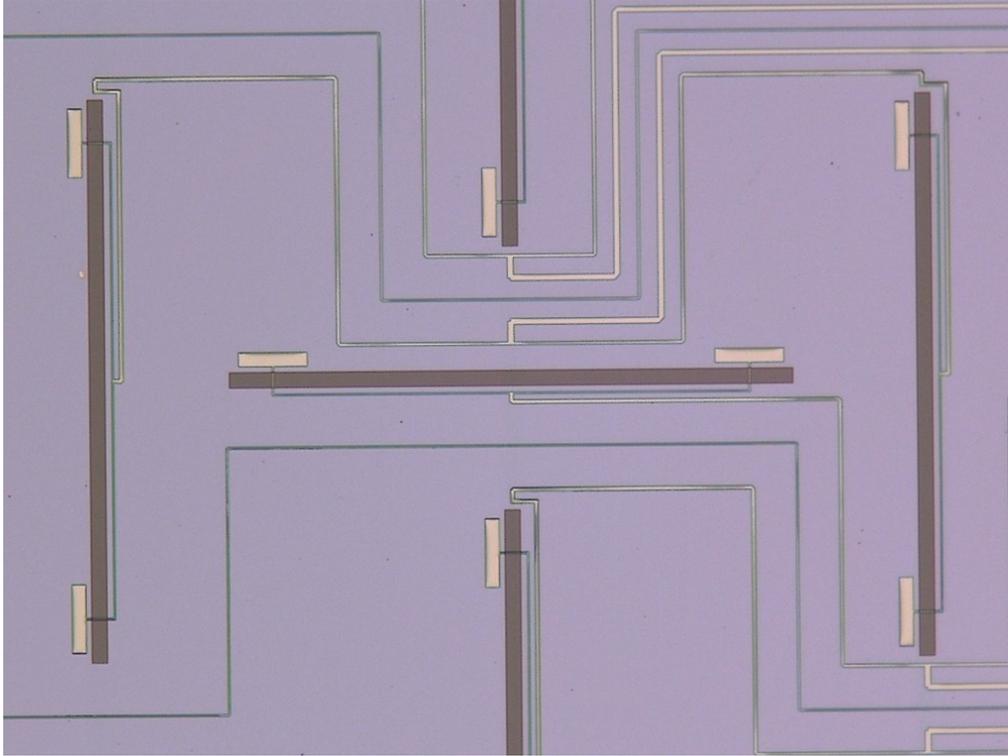


Figure 2.12: A micrograph (800x magnification) of a portion of the antenna, showing four slots (subantennas) of the two antenna arrays and the summing network feeding the slots.

polarization. The components of a focal plane pixel are shown in figure 2.11. The following naming convention is used to describe the polarization orientation of each bolometer: A-detectors measure horizontally polarized light and B-detectors measure vertically polarized light, with respect to the pixel orientation in figure 2.10. The polarization orientation with respect to the focal plane axes is shown in figure 2.6.

2.4.2.1 antenna

Planar phased-array antennas couple to incoming radiation and form collimated beams. There are two colocated antenna arrays, one measuring vertical polarization and the other measuring horizontal polarization with respect to the detector tile axes shown in figure 2.10. Each antenna is composed of a 12×12 array of subantennas, covering a 7.4×7.4 mm square patch. The antenna centers are spaced 7.82 mm horizontally and 7.80 mm vertically in figure 2.10. The subantennas are slots ($12 \times 412 \mu\text{m}$) that are etched through the ground plane to the protect layer, as shown in figure 2.12. The electric field of incoming radiation excites a voltage difference across the short axis of the slots. This is read out inductively by a pair of coupling capacitors that terminate superconducting

microstrip lines. There are two offset-fed stubs per slot to maintain symmetry and reduce cross-polar response. The width of the beam is set approximately by λ/d , where λ is the wavelength in vacuum and d is the width of the antenna. This yields a raw beam width of $14.4'$ FWHM. The slots are long in order to achieve $\sim 35\%$ fractional bandwidth prior to the band-defining filters.

2.4.2.2 summing network

A summing network of superconducting niobium microstrip lines adds the power from the 144 subantennas of the array coherently. As the summing trees from each of the subantennas add together the microstrip lines become progressively wider in order to keep the line impedance uniform and reduce reflections back into the antenna, which requires tapering at the summing branches. The reflections in the summing network have been measured to be less than 1% over a large bandwidth [39]. The subantennas are fed uniformly; each slot contributes equally to the coherent power relayed to the bolometer. This results in a uniform, or a two-dimensional top hat aperture illumination that has side lobes in a fourfold symmetry in the far-field antenna pattern. The side lobes are at a level of -15 dB and are terminated on the cold aperture stop.

The microstrip lines run on top of the $0.3 \mu\text{m}$ ILD layer parallel to the ground plane. The use of a slot antenna array and backside illumination shields the summing network from incoming radiation with the ground plane. The dielectric loss tangent of the ILD has been measured to be $\sim 6 \cdot 10^{-4}$ at 145 GHz in sub-Kelvin engineering arrays [39]. The length of the microstrips leading from the subantennas to the TESs is typically 2-3 cm. This results in only a negligible degradation of the optical efficiency due to loss in the microstrip.

2.4.2.3 filter

Signals from the summing network pass through a compact, inline microstrip filter shown in figure 2.13. The filter defines the spectral pass band of the detectors beyond what is set by the geometry of the antenna array. The filter is a lumped-element, 3rd-order bandpass Chebychev LC filter, using coplanar waveguide (CPW) inductors and microstrip stub capacitors [39]. The pass band was designed for a center frequency of 148 GHz and 25% fractional bandwidth to complement an atmospheric transmission window at the South Pole. The band edges have very sharp cutoffs and the in-band transmission is high. The filter is naturally nonresonant and not believed to have fundamental harmonic leaks. The high bandwidth gives high optical efficiency and detector sensitivity.

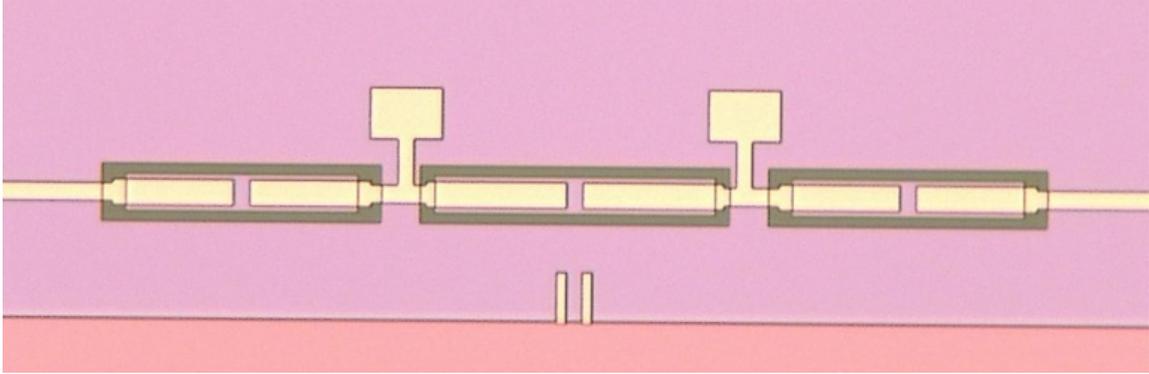


Figure 2.13: A micrograph of the inline microstrip filter. A CPW inductor is formed by islands of niobium released from the ground plane. Capacitive stubs separated from niobium by the ILD form series capacitance over the island and shunt capacitance over the ground plane

The design is a combination of three series CPW inductors, six series microstrip capacitors, and two shunt microstrip capacitors. The CPW inductors are formed by three islands ($14.0 \times 90.6 \mu\text{m}$) of niobium ground plane that are separated from the bulk ground plane by a $4.5 \mu\text{m}$ gap. The microstrip entering the filter from the summing networks is $7 \mu\text{m}$. It is deposited on top of the $0.3 \mu\text{m}$ ILD layer covering the niobium islands and ground plane. Capacitive stubs are created where the microstrip terminates in pads above either the island or the ground plane. The stubs ($10 \times 47.6 \mu\text{m}$) deposited above the islands create series capacitive elements, while the stubs ($20.0 \times 28.3 \mu\text{m}$) deposited above the ground plane form shunt capacitors.

The geometry of the CPW islands and gap sets the value of the inductive elements in the filter. The geometry of the stubs and the size of the gap and dielectric constant of the ILD sets the series and shunt capacitances. The band center is controlled by the thickness and dielectric constant of the ILD. The thickness of the ILD must be carefully controlled during the fabrication process to maintain repeatable band centers. This is crucial to ensure the pass band doesn't overlap with oxygen and water lines located below and above it, respectively, which would negatively impact the loading and noise of the detectors. Spectra for a typical pair of detectors are shown in § 5.6.3.

2.4.2.4 bolometer

The filtered microstrip continues on to a thermally isolated bolometer island, shown in figure 2.14. It includes a lossy gold meander, titanium and aluminum TESs, a layer of gold to increase the bolometer time constant, electrical leads for the TESs and three rectangular holes for the promotion of etching below the island. The free-standing island is composed of a $150 \times 310 \mu\text{m}$, $1.0\text{-}\mu\text{m}$ -thick SiN membrane. Six released SiN legs support the island and the electrical leads.

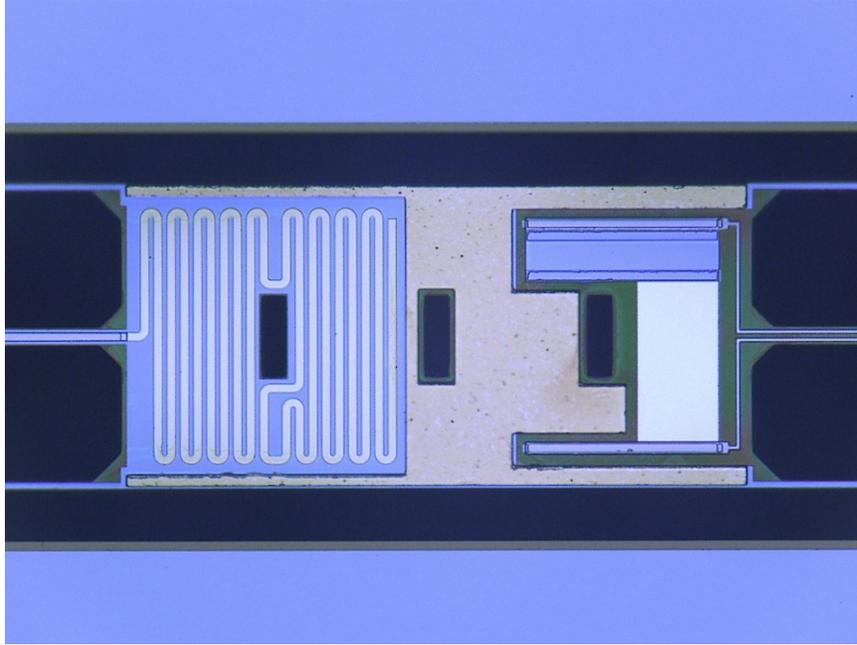


Figure 2.14: A micrograph (700x magnification) of the free-standing bolometer island, showing the legs, input and output lines, the gold meander, gold deposition layer, and the aluminum and titanium TESs. Power from the antenna-coupled microstrip enters the island on the left-center leg and is thermalized on the island by the gold meander. The change in output current of the voltage-biased TES is read out from the bias lines that enter the island on the right-center leg

Unlike the summing network the detector islands are not shielded from incident radiation, they are exposed through the bulk silicon of the array. Earlier island geometries showed a significant direct coupling to high-frequency radiation, with the island and leads forming an antenna. To address this issue the geometry of the island was changed to the one shown in figure 2.14. The ground plane was moved closer to the long edges of the island and the incoming microstrips and much of the island are now shielded with niobium.

The support legs are $535\ \mu\text{m}$ in length and $1.15\ \mu\text{m}$ thick. They are composed of a $1.0\ \mu\text{m}$ layer of SiN covered by the $0.15\ \mu\text{m}$ niobium ground plane layer. The four legs connecting to the corners of the island are $4\ \mu\text{m}$ wide, while the two legs connecting to the midline are $9\ \mu\text{m}$ wide to accommodate the input and output microstrips. The island is completely released from the bulk silicon wafer leaving a scalloped ditch between the island and legs, and the bulk silicon. The etching extends approximately $350\ \mu\text{m}$ into the bulk silicon, but does not extend all the way through the wafer. The legs are chamfered where they attach to the island and the bulk wafer to reduce the risk of leg breakage. A corner without a chamfer creates a stress point in the SiN that can shear the leg when the island is exposed to excess force. The chamfered legs make the islands strong enough to sustain a fairly high rate of air flow from an air duster, or the forces encountered during handling

and transport of the arrays.

The microstrip runs from the filter on to the island along one of the $9\ \mu\text{m}$ wide center legs. It is step-tapered from $7\ \mu\text{m}$ at the filter to $4\ \mu\text{m}$ where it connects to the bolometer resistor. The resistor is a $4\ \mu\text{m}$ wide, $0.15\ \mu\text{m}$ thick lossy-gold meandering microstrip. Incoming power is dissipated in the resistor so that the temperature of the island increases with increased optical power. Given the measured loss tangent of the gold, the resistor must be at least $2.2\ \text{mm}$ long to properly dissipate the incoming power, including reflection from the unterminated end. Simulations show that the design, which uses a $2230\ \mu\text{m}$ meander, should dissipate $\sim 99\%$ of the power on the island [39]. The dissipation time constant is much shorter than the bolometer time constant, ensuring efficient thermal coupling of the antenna and bolometer.

The bolometer island includes thin-film, elemental titanium and aluminum TESs connected in series. The titanium TES is $15.8 \times 95.0\ \mu\text{m}$ and $0.25\ \mu\text{m}$ thick, orientated with its long axis parallel to the long axis of the island. The aluminum TES is $78.5 \times 39.5\ \mu\text{m}$ and $0.05\ \mu\text{m}$ thick and orientated with its long axis perpendicular to the long axis of the island. It is offset from the electrical connection bars to accommodate an etching hole. The transition temperature of thin titanium films is approximately the same as the $0.4\ \text{K}$ critical temperature of bulk titanium and is insensitive to variations in the thickness of the film. Unexplained excess detector noise has been reported by numerous groups using TESs, and has been demonstrated to depend on the geometry of the TES [60]. A 1/6-square geometry was selected for the titanium TES and a 2-square geometry for the aluminum because they showed less excess noise during engineering tests of the detectors [8]. The niobium ground plane that is deposited on most of the island is not placed under the TES films. This prevents unwanted alteration of the transition temperatures of the TES films by the proximity of the niobium and prevents flux pinning in the niobium that might lead to increased noise in the TESs.

The two TESs are electrically connected by horizontal bars of aluminum and niobium vias formed from the ground plane layer. Two $2\ \mu\text{m}$ niobium microstrip lines connect to the niobium vias and complete the electrical connection to the TESs. The voltage bias is applied to the TES and the output current is measured from these electrical lines. The niobium lines leave the island along one of the $9\text{-}\mu\text{m}$ -wide center legs. These lines connect to $5\ \mu\text{m}$ microstrip lines, which then join a wiring bus that runs in the gaps between the antenna patches of adjacent pixel columns. The wiring bus connects to niobium bond pads located along one edge of the array. There are nine wiring buses in total; the two edge wiring buses have 16 microstrips and service the outer eight TESs in the outer

pixel columns and the seven inner wiring buses have 32 microstrips and service the nearest 16 TESs of the two adjacent pixel columns. The microstrips of the wiring bus are deposited on top of the ILD layer that covers the niobium ground plane. The ground plane is anodized to prevent electrical shorts between the two niobium layers. Preventing shorts in the wiring layer is critical, since 32 TESs are connected in series in the Nyquist interface chips through the TES bias lines (§ 2.19). A short in one of the lines will short all of the TESs, requiring that TES to be electrically disconnected. This issue was seen in earlier engineering arrays and addressed by a change in fabrication procedure [8].

The use of two different TESs allows for the use of the same bolometer under different loading conditions. The titanium TES is designed with a 450 mK superconducting transition temperature and 9 pW of saturation power. It is used during standard observations where the telescope is directed at high elevations where the background loading is low. The aluminum TES is designed with a 1.3 K superconducting transition temperature and 600 pW of saturation power. The higher loading margin allows it to operate under the higher background loading conditions of the laboratory or, for example, during low-elevation beam mapping of ground-fixed sources.

A 2.5 μm thick layer of gold is deposited on the island to increase its heat capacity. It covers most of the island, except for the locations of the TESs and resistor. A 4–5 μm gap is left between the gold and those areas to avoid parasitic electrical effects. Distributed heat capacities on bolometer islands can lead to increased excess detector noise, so the gold layer is deposited directly on the niobium ground plane to ensure good thermal contact with the island [20][30]. The additional heat capacity was added to allow for an increase in the bandwidth-limiting inductor of the Nyquist interface chip, while maintaining margin on stability of the TES (§ 2.19).

2.4.2.5 alignment holes

The detector tiles are aligned in the focal plane using a pair of alignment through holes that mate with alignment pins on the detector plate. The through holes are created during the deep trench etching step that cuts the array from the wafer. One hole is circular with a 1.04 mm diameter, while the other is a slotted notch of the same diameter that extends 1.58 mm into the tile. The holes are located at the midline of the wire bond edge and the opposite edge, as shown in figure 2.10. During cool down the detector plate contracts 0.22 mm along an edge relative to the low-CTE silicon array. The circular hole fixes the center of the wire bond edge with respect to the detector plate and circuit board and prevents gross relative motion of the wire bonds connecting the circuit board

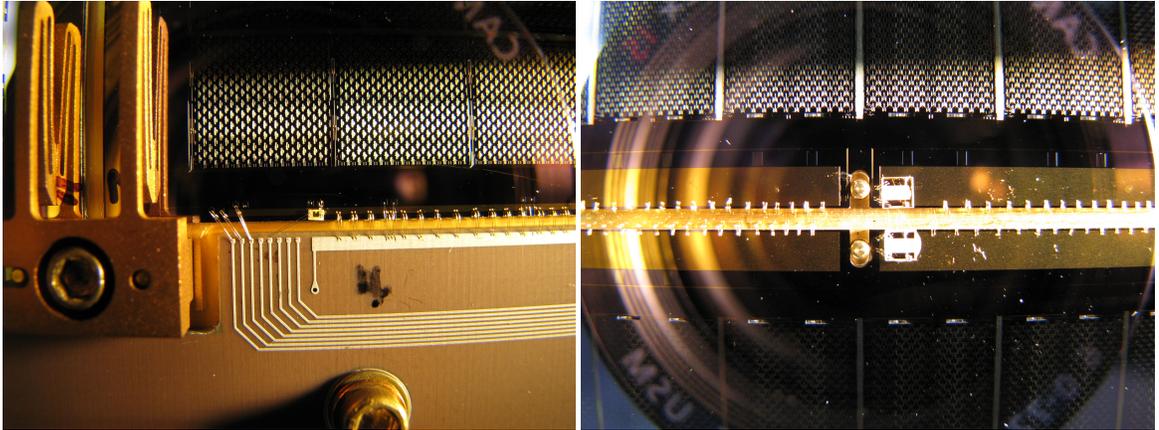


Figure 2.15: The detector tile thermal components. *Left:* A close-up image of a tile edge showing a tile clip, resistive chip heater (white square), the gold border with heat sink bonds connecting it to a grounding layer on the circuit board, and the electrical lines and wire bonds that provide bias to the heaters and NTDs. *Right:* A close-up image of the edges of two detector tiles, in the middle of the tiles opposite the wire bond edges. The NTD thermistors for the tiles are the white squares in the center. The gold border and heat sink bonds connecting it to the detector plate are shown. The alignment slots and pins are also shown.

to the detector tiles. The 0.11 mm contraction on either side of the circular hole is compensated by the loop heights of the wire bonds. The slot in the opposite edge accommodates the 0.22 mm of relative contraction of the plate. The alignment holes mate with 1 mm diameter pins press fit into the detector plate. A close-up view of the slotted alignment holes opposite the wire bonds is shown in figure 2.15.

2.4.2.6 thermal components

Thermometers and heaters have been added to the detector tiles to monitor the tile temperatures, heat the TESs into the superconducting transition for operation, and measure the thermal conductivity between the tile and detector plate. A pair of NTD thermistors and resistive chip heaters have been added to each detector tile. The NTDs are monitored for scan-synchronous thermal fluctuations that have not been suppressed by the thermal control (§ 2.2.3). They are attached to the front side of the detector tiles on the gold heat sink border. Wiring leads and bond pads route the connections to the edge of the tile and are superconducting to minimize thermal conduction back to the circuit board along the lines. The thermal components of the tiles are shown in figure 2.15.

TESs can generally be biased on transition through self-heating by applying the maximum TES bias. For some bolometers, especially the dark TESs that receive no additional optical load, this technique does not work and they remain superconducting. To drive these detectors from the

superconducting to normal state, the maximum bias is applied and the tile heaters are pulsed for one second to deposit a small amount of heat on the tiles. This is sufficient to bias all operational detectors into the transition.

The antireflection and detector tile substrates absorb some incident radiation and the tile temperature can fluctuate with changes in optical loading. Therefore it is important to properly heat sink the detector tiles to the large thermal reservoir of the detector plate. To facilitate conduction a thick border of gold is deposited on the front side of the detector tiles as shown in figure 2.10. The array edges adjacent to the wire bond edge each have a 0.35×63.7 mm continuous gold strip. The edge opposite has two 2.2×32.5 mm strips that are separated by the slotted alignment hole. In addition, there are two 1.5×4.3 mm strips along the wire bond edges that connect to the borders along the edges adjacent to the wire bond edge. These two smaller strips are placed under the feet of the tile clips, as is the border along the edge opposite the wire bond edge, to connect the gold border to the well-sunk tile clips. The bulk of the heat-sinking, though, is provided by approximately 1200 1.5 mil gold wire bonds. These bonds connect the gold border and the detector plate along two edges and connect to a strip of aluminum-clad copper on the circuit board along the third edge.

2.4.3 antireflection tiles

Antireflection (AR) tiles are used to reduce reflections from the illumination side of the detector tiles. They are located between the detector tiles and plate, and can be seen through the focal plane apertures in figure 2.9. The cross section of the AR tiles matches the detector arrays and they have the same alignment holes for registration. The thickness of the AR tiles is $241 \mu\text{m}$. In order to increase thermal conduction between the tile stack and the detector plate, the AR tiles are made from z-cut crystal quartz oriented so that the high-conductivity axis is normal to the tile cross section.

2.4.4 tile clips

The four sets of detector and AR tiles are held to the detector plate using nine beryllium-copper spring clips. A clip consists of a cuboid base and one or more legs extending from the base. The legs have a levered profile to create a spring force at the foot of the leg. The bases are machined and then the leg profile is formed using wire electrical discharge machining (EDM). The base is undercut to prevent interference with the tile edge during thermal contraction of the detector plate. A pair of 1 mm diameter slotted and through holes in the base maintain alignment of the spring clips. They

also prevent the feet from rotating into wiring layers on the detector tiles when the clips are fastened to the detector plate with 2-56 screws.

The tile clips come in four configurations; there are four clips with one leg at the four corners of the tile array block, two clips with two parallel legs at the midline of the array block, two clips with two antiparallel legs for the other midline, and one clip with two sets of two parallel legs at the block center. The tile clip configurations and arrangement are shown in figure 2.4 and a close-up view of one tile clip is shown in figure 2.15.

The tile clips apply ~ 0.5 N of force to each corner of the tile stack. Directly below the feet of the spring clips the detector plate has landings to support the force. The gold heat sink on the tile is wider at the landing of the foot to promote conduction from the tile into the tile clips.

2.4.5 printed circuit board

Routing of electrical signals in the focal plane is performed by a multilayer, G10 printed circuit board (PCB) manufactured by Murietta Circuits.⁸ The PCB is 265 mm in diameter and composed of multiple layers of G10 and copper wiring layers. There is no protective solder mask on the top layer of the board and the G10 and traces are visible in figure 2.4. It is fixed to the detector plate by 4-40 screws and lock washers to maintain adequate thermal contact through thermal contraction. In addition, the attachment holes are radially slotted to allow the copper detector plate to contract relative to the lower CTE G10 without bowing the circuit board. Electrical signals are routed from the circuit board using six 37-pin MDM connectors and four 4-pin microdot connectors.

A continuous superconducting circuit is required between the TESs and input to the SQUID multiplexer chips and between the output of the SQUID chips to the series arrays, to prevent parasitic resistance in those circuits. The top layer of copper traces are coated with aluminum, which has a critical temperature (1.2 K) well above the focal plane operating temperature. The remaining portions of the circuit, including the wire bonds and wiring layers on the multiplexer chips and detector arrays, are all superconducting. Although the bias and flux feedback traces are not required to be superconducting, they are also coated with aluminum.

Bias lines for the first-stage SQUIDs are connected to each of the 16 multiplexer chips in series. To service all of the chips, these bias or address lines must be routed around the board on several layers connected by vias. To minimize pickup in the bias circuit loops the return traces are routed directly below their counterpart on lower layers of the circuit board, doubling back from the final

⁸Murietta Circuits: www.murietta.com

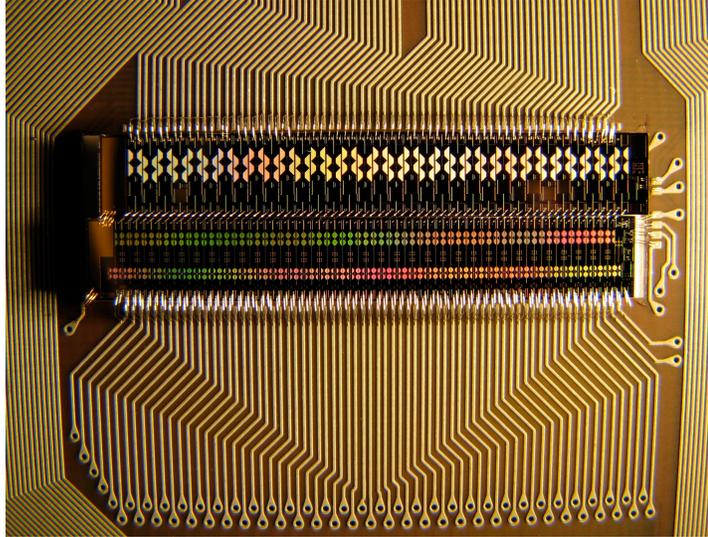


Figure 2.16: A pair of SQUID multiplexer and Nyquist interface interface chips with wire bonds and electrical routing. The NYQ chip is on the top and connects to aluminum traces leading to 32 detectors on one of the arrays. The NYQ chip connects to the MUX chip on the bottom, which contains 33 first-stage SQUIDs that are coupled to a single second-stage SQUID. The lines leading to the MUX chip on the bottom are the first-stage SQUID bias lines, which are wired in series for all the chips. The additional connections on the right of NYQ chips are the single pair of TES bias lines for the 32 detectors in the column. The additional bonds on the right of the MUX chip are the second-stage bias and flux bias lines, and the servoed first-stage flux feedback lines. Gold wire bonds connect gold heat sinks on the left of the chips to grounding strips on the circuit board.

chip all the way to the originating connector.

2.4.6 Nyquist and multiplexer chips

The small current signals from the bolometers are amplified and transduced to an output signal using a three-stage, time-domain SQUID multiplexer. The central components of the multiplexer are the version MUX07a SQUID multiplexer (MUX) chips and the SQUID series arrays (SSAs). Signals from the bolometers are bandwidth-limited by Nyquist interface (NYQ) chips to prevent aliasing of high-frequency detector noise. The NYQ chips also provide shunt resistors for biasing the TESs. The MUX and NYQ chips and SSAs are provided by collaborators at NIST-Boulder.⁹

The operation of the cold readout electronics is detailed in § 2.5, but the NYQ and MUX chips are introduced here as components of the focal plane. There are 16 pairs of chips used in the focal plane, which are located around the perimeter as shown in figure 2.4. A pair of chips are placed with their long axes parallel to each other and to the wire bond edges of the arrays, with the NYQ chips closer to the arrays. Ideally, the magnetically sensitive chips would all be equally spaced from the

⁹The group is led by Kent Irwin.

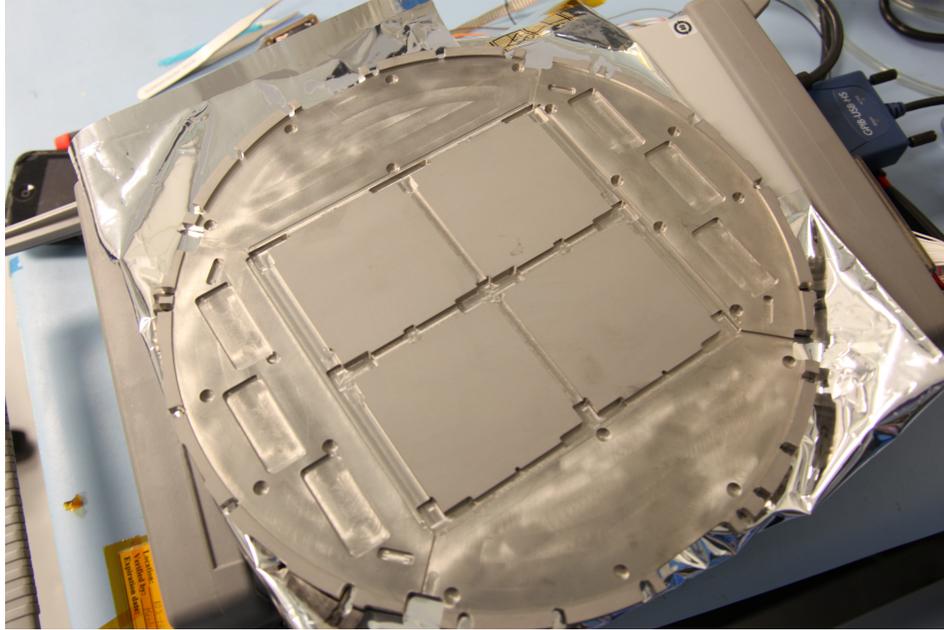


Figure 2.17: The superconducting niobium backshort. The quarter-wave backshort surface that faces the detector tiles is shown. Trenches separate the tile backshort surfaces to allow clearance for the wire bonds and tile clips. The surfaces above the multiplexer chips (top left and bottom right) are smooth to prevent modulation of the local magnetic field. Note this backshort is an older version that doesn't include a threaded hold in the center that mates with a spacer and screw to control the backshort separation.

unshielded edges of the array apertures and backshort. However, space is needed to fan out traces in the circuit board and the chips are forced into two unequally spaced rows.

A pair of NYQ and MUX chips are attached side-by-side to an alumina carrier using Stycast 2850 epoxy with a low-viscosity catalyst (24LV). The chips are 3×20 mm in dimension. This modular mounting scheme allows sets of chips to be easily replaced by removing the carrier without damaging the circuit board. The alumina carriers are attached to the PCB using a bead of Stycast 2850 (catalyst 9) along the short edges of the carriers.

2.4.7 backshort

A superconducting niobium backshort is used to magnetically shield the TESs and SQUIDs and provide a quarter-wave backshort behind the antennas. The backshort has the same 27 cm diameter as the detector plate and is approximately 5.5 cm thick. It is machined from a single piece of unalloyed, reactor grade niobium. The lapped side of the backshort that faces the arrays is shown in figure 2.17 and the opposite side is shown in figure 2.8.

The backshort is mated to the detector plate by 20 fasteners and spacers placed along the

perimeter. The quarter-wave spacing of the arrays and backshort is enforced by lapped macor¹⁰ spacers. The mating surfaces along the edges of the backshort and plate are coplanar with the backshort surface behind the arrays and the surface of the plate the tile stack is clamped to. These surfaces are lapped and so provide a very stiff and planar reference surface for the focal plane. A small-diameter macor spacer is used in the center of the focal plane to prevent bowing of the copper plate away from the rigid backshort. The detector plate and backshort are clamped against the center spacer using a 2-56 screw that threads into the backshort. To prevent thermally driven supercurrents from generating spurious magnetic fields, the center screw is made from niobium and electrically isolated from the detector plate by a nylon washer.

The spacer thickness was determined by optical metrology measurements. The relative height between the antenna slots and detector plate surface was measured for the four arrays. Due to manufacturing variations in the thickness of the silicon wafers and the AR tiles, the stack up heights vary between the arrays. With $\lambda/4 = 507 \mu\text{m}$ at 148 GHz and a mean relative height of $790 \mu\text{m}$ the spacers were lapped to $1297 \mu\text{m}$.

Additional features were machined into the face of the backshort that mates with the detector plate. To prevent the wire bonds from shorting to the backshort, grooves were cut around the detector array locations. Care was taken to make sure the backshort overlaps the antennas as much as possible. The overlap that was achieved, given design constraints, was 1.8–2.9 mm around the perimeter of the array. The overlap is smaller for the corner pixels where additional clearance is needed for the tile clips. Clearance is also needed for the MUX and NYQ chips that sit proud of the circuit board due to the alumina carriers, so pockets are cut into the niobium at those locations. The SQUIDs are sensitive to gradients in magnetic field applied normal to them. The surface of the backshort near the chips is free of corners that could introduce a normal component to the magnetic field in those regions.

2.4.8 assembly

There are approximately 6900 aluminum and gold wire bond interconnects in the focal plane. These were made using a Westbond¹¹ 4546E semi-automatic wire bonder. The 1 mil aluminum bonds provide electrical connections between the circuit board, detector arrays and the MUX and NYQ chips. The 1.5 mil gold wire bonds provide thermal heat sinking of the arrays and chips to the detector plate and circuit board.

¹⁰Corning, Inc.: www.corning.com/docs/specialtymaterials/pisheets/Macor.pdf

¹¹West Bond, Inc: www.westbond.com

There are 5572 aluminum wire bonds: 1024 bonds between the arrays and circuit board, 1024 bonds between the circuit board and the NYQ chips, 1024 bonds between the NYQ and MUX chips, 2048 bonds between the address lines of the MUX chips and circuit board, 256 bonds from the other feedback and bias lines of the MUX and NYQ chips to the circuit board, 64 bonds for the array NTDs and heaters, and 132 Faraday cage bonds. Since the address lines of the MUX chips are connected in series, a failure of even one bond can result in the loss of 16 detectors. Redundant bonds were used on those connections to prevent single-point failure. Aluminum wire bonds connect the bond pads of the array heaters and NTDs to the circuit board to minimize conduction through the leads and to maximize the heat delivered to the arrays. These bonds were also doubled to prevent the loss of those components due to wire bond breakage.

Thermal contraction of the detector plate and circuit board relative to the arrays during cool down creates relative motion of the wire bond pads of the arrays and circuit board. This was compensated by the loop heights of the wire bonds. Bonds were made between the niobium ground plane of the arrays and the detector plate or circuit board to create a Faraday cage between the tiles and focal plane. This prevents radio-frequency radiation from entering the focal plane through the gaps between the ground plane of the arrays and the detector plate.

There are approximately 1300 gold wire bonds for heat sinking the arrays, MUX and NYQ chips. Superconducting aluminum bonds have reduced thermal conduction due to the absence of electron conduction, so gold bonds are used for heat sinking. The array bonds connect the gold border of the tile to the detector plate on two edges and to the circuit board on the third edge. The wire bond edge doesn't include heat sinking bonds, but is adequately sunk through the densely-packed detector bonds. Two array heat sink bonds were placed every ~ 1 mm along the three edges, for approximately 100 bonds per array edge. Another 100 gold bonds connect gold heat sink bond pads on the MUX and NYQ chips to the circuit board.

The gold wire bonds were made by heating the entire focal plane unit to approximately 80 °C and heating the bonding wedge. When Josephson junctions are exposed to temperatures in excess of 125 °C their critical currents and normal state resistance begins to degrade with time [43]. To prevent annealing of the Josephson junctions the focal plane was minimally heated for as short a duration as possible. Hence, the gold wire bonding was performed at only 80 °C, though a higher temperature would have made bonding easier. The gold wire bonds from the arrays to the detector plate were made before the circuit board, which contains the MUX chips, was installed. The only bonds that must be made with the MUXs at an elevated temperature are the heat-sinking bonds

for the MUX and NYQ chips, and the gold array to circuit board wire bonds. The array NTD and heaters were installed using silver epoxy and cured at 180 °F for four hours. This was done to the bare arrays before installation in the focal plane to prevent heating the MUX chips.

2.5 cold readout electronics

The cold readout electronics consists of the components introduced earlier: the NYQ chips, the MUX chips and the SSAs. The NYQ chips provide the detector bias and bandwidth-limit the TES signals. The MUX chips contain the first two stages of the SQUID multiplexer. The SSAs provide the final SQUID amplifier stage. The NYQ and MUX chips are located on the focal plane at ~ 300 mK, while the SSAs are attached to the 4 K base plate. The chips are completely lithographically fabricated by NIST-Boulder. The multiplexing architecture is 33×16 : 32 TESs are read out in a multiplexer row in addition to a “dark” first-stage SQUID and there are 16 multiplexer columns. So, 528 channels are read out, but only 512 of those are coupled to a TES.

2.5.1 time-domain SQUID multiplexing

The small current output of a TES would be difficult to read out if not for SQUID current amplifiers. The designed output impedance of the BICEP2 TESs during operation is $30 \text{ m}\Omega$. This low resistance would be mismatched to the noise of available FET amplifiers. SQUIDs, however, are easily impedance-matched to TESs. In addition, they offer a natural means to multiplex a large number of detectors with fewer wires. If each of BICEP2’s 528 readout channels had a separate pair of wires running to 4 K the loading of the focal plane would overwhelm the UC fridge stage. The multiplexer reduces the wiring count to the focal plane from 528 to 97 pairs of wires.

2.5.1.1 SQUID amplifier

Before discussing the operation of a SQUID multiplexer, first consider how a SQUID amplifier can be used to read out a single TES, as illustrated in figure 2.1. The current output of a TES is inductively coupled to a single SQUID by an input coil. A feedback coil is also inductively coupled to the SQUID. The amplifier is operated in flux-lock loop (FLL) to linearize the periodic output and increase the dynamic range of the SQUID’s response to input flux. A lock point is selected on a linear, high-gain portion of the SQUID’s $V - \Phi$ response curve. As the flux from the input coil changes in response to the TES current, a compensating flux is applied by the feedback coil to

cancel it. The amount of flux feedback needed to cancel the TES signal serves as the output of the detector channel. Below a minimum bias current, I_c^{min} , the SQUID is completely superconducting and generates no output signal for changes in input flux. So, the SQUID must be biased at least above I_c^{min} for operation. Above a second bias current, I_c^{max} , the SQUID always generates an output voltage for changes in the input flux. The dynamic range is largest just above I_c^{max} and grows smaller as the bias is increased above it. The SQUID output signal is coupled to an amplifier that boosts and buffers it to room temperature electronics. An analog-to-digital (ADC) converter samples the difference between the output of the SQUID amplifier and the output expected from the lock point. The difference, or error, signal is fed to a digital-to-analog (DAC) converter with a series feedback resistance (R_{fb1}) that produces a compensating current and flux in the feedback coil.

The first-stage SQUIDs dominate the noise of the amplifier, so have been designed to provide low noise and high gain. Their low output impedance must be matched to an intermediate, higher-output impedance stage. The 100 SQUIDs of the SSA each have $\sim 1 \Omega$ dynamic resistance for a $\sim 100 \Omega$ total output impedance that is well matched to room temperature electronics. The first-stage could be coupled directly to the SSA, but the NIST SQUID amplifier uses an intermediate (SQ2) stage for several reasons. The SSA cannot be placed on the focal plane because its high dissipation is too great of a load for the UC stage, so it is placed at 4 K. The first-stage SQUIDs could be connected directly to the SSA, but this would require routing the summing coil to 4 K. The superconducting loop would need to be carefully shielded, since it would generate screening currents for any changes in flux through the large loop area, which would couple flux to the SSA input. Also, the parasitic inductance from the cabling would reduce the coupling efficiency of the summing coil [51][16].

2.5.1.2 SQUID multiplexer

When a SQUID is biased below I_c^{min} it will not generate output in response to changes in input flux. This allows a number of front-end SQUIDs, each coupled to a TES, to share a common readout amplifier. The TESs are continuously biased but are only sampled when the SQUID they are coupled to is biased normal. The off-state SQUIDs are superconducting and contribute no noise, heating or power to the common amplifier chain. In a time-domain multiplexer the SQUIDs are biased on, the output is sampled and then the bias is turned off, in sequence. When all rows in the multiplexer have been visited — a full frame — the multiplexer returns to the first row in the sequence. Due to the finite open-loop bandwidth of the multiplexer each row must be visited long enough for the

amplifier to settle before it is sampled.

In the NIST three-stage SQUID multiplexer 33 first-stage SQUIDs (SQ1s) are coupled to a common amplifier consisting of two additional stages: a second-stage SQUID (SQ2) and a SQUID series array (SSA). There are 16 SQ2s and SSAs, for a 33×16 multiplexing architecture. The 33 SQ1s are coupled to a single SQ2 on a MUX chip. The output of the SQ2 is routed to the SSA at 4 K. The 16 SQ1s in a row are wired in series with a room temperature bias resistor (R_{sq1}). This common bias turns one row or SQ1 on in each column, and is turned on and off sequentially for the 33 rows.

Each SQ1 is shunted by a 1.5Ω address resistor and an output inductor (L_{int}). The address resistance is comparable to the $\sim 1-3 \Omega$ dynamic resistance of the SQ1s, which means the SQ1s operate between a voltage and current bias. Changes in the SQ1 input flux cause a modulation in the current through the resistor and output coil. Each output coil is coupled to one of 33 inductors (L_t) connected in series in a superconducting transformer loop called the summing coil. In addition to the common transformer loop, all of the SQ1s are coupled to inductors (L_{fb1}) in series on a common feedback line. The feedback line is biased with a room temperature series resistor (R_{fb1}) to provide dynamic flux feedback. Since only one SQ1 in a column operates at a time, the transformer and feedback can be shared.

The transformer loop is coupled to the SQ2 by an inductor (L_{in2}). The SQ2 is flux biased to a useful lock point by an additional inductor (L_{fb2}) that is biased by a room temperature resistor (R_{fb2}). The SQ2 is shunted by a bias resistor ($R_{sq2bias}$) and its circuit is biased by a room temperature resistor (R_{sq2}). The $90 \text{ m}\Omega$ SQ2 bias resistance is much smaller than the $\sim 3 \Omega$ dynamic resistance of the SQ2, so it is operated in a voltage-biased mode. Input flux from the transformer inductor modulates the output current of the SQ2, which is carried to the SSAs located on a circuit board mounted to the 4 K base plate. The SQ2 bias resistors are also located on this circuit board.

The SQ2 output lines are wired in series with 100 input inductors (L_{ina}) each coupled to one of the 100 SQUIDs in the SSA. Changes in the output current of the SQ2 modulate the input flux to the SSA. The voltage response of each SQUID in the SSA adds in phase and boosts the gain and dynamic range of the final amplifier stage. The SSA is current-biased by a room temperature resistor (R_{sqa}). Flux modulation from the 100 input inductors generates a voltage response that is fed to the room temperature electronics through the SSA bias lines. Flux bias to control the SSA lock point is provided by 100 series inductors (L_{fba}) connected in series to a room temperature bias resistor (R_{fba}).

2.5.1.3 bandwidth considerations

The SQUID amplifier has limited bandwidth due to combination of the poles from the LRC circuits at the different stages of the amplifier. The lowest pole dictates the open-loop bandwidth of the amplifier and how long it takes the amplifier to settle τ_{settle} . This in turn affects how quickly the multiplexer can cycle through rows without introducing significant crosstalk from the preceding row. Each row in a column has a different DC output level at the final amplifier stage, even in the absence of input flux from the TESs. The amplifier must be given adequate time to slew between these levels. In addition, switching transients at the beginning of each row visit generally increase the amount the output must slew (§ 4.1.5.2).

In order to ensure the amplifier has adequately settled before sampling, the multiplexer must spend a time ($\tau_{dwell} = x\tau_{settle}$) on the row, where x is a safety factor. Each row is sampled at the multiplexer frequency $f_{mux} = \frac{1}{\tau_{dwell}}$. The rate at which every row is sampled, or frame rate, is then dictated by the number of channels multiplexed (N_{mux}) by $f_{frame} = \frac{1}{N_{mux}\tau_{dwell}}$. The amplifier bandwidth must be a factor N_{mux} larger than the readout bandwidth required for each detector. The noise bandwidth — assuming a one-sided PSD — of the amplifier is $f_{amp} = \frac{f_{mux}}{2\tau_{dwell}}$, and the Nyquist frequency of the detector is $f_{Nyq} = \frac{f_{frame}}{2}$.

The SQUID amplifier noise power S_{amp}^2 is roughly white, which means that the total aliased noise power in the detector readout will just be the ratio of the noise bandwidth to the Nyquist frequency of the detector:

$$S_{tot}^2 = f_{amp}/f_{Nyq} = \frac{f_{mux}}{2\tau_{dwell}} \cdot \frac{2}{f_{frame}} = N_{mux} \quad (2.5)$$

This N_{mux} increase in the amplifier noise is the so-called multiplexer disadvantage inherent in a time-domain multiplexed readout. In order to compensate for this increase in amplifier noise the mutual inductance of the input coil must be increased so that the current noise referred to the input coil of the TES is not degraded, or $S_{I,amptot}^2 = S_{\Phi,amptot}^2/M_{in}$. The required level of coupling is dictated by the relative amplifier and detector noise levels, but to get the same current noise level of an unmultiplexed amplifier the mutual inductance must be a factor of $\sqrt{N_{mux}}$ higher.

The bandwidth of each amplifier stage has been designed to be as large as possible while equal to the other stages. If the bandwidth of one stage greatly exceeds another the noise bandwidth will go up but without a corresponding increase in the multiplexing bandwidth, which is set by the slowest time constant. Thus the aliased noise penalty will increase. The input-referred noise level decreases

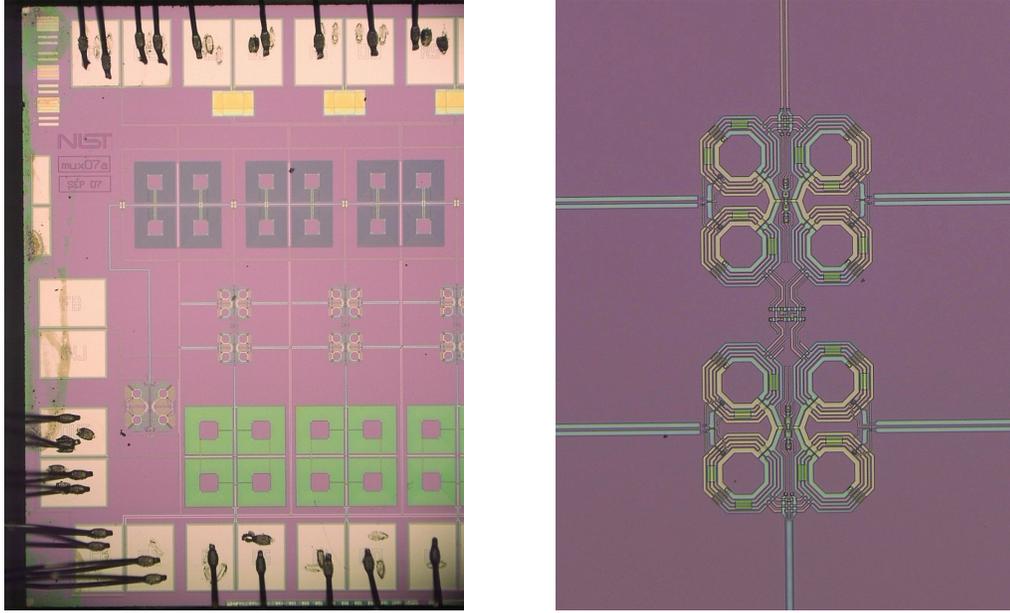


Figure 2.18: Micrographs of a multiplexer chip and close-up of the first-stage and dummy SQUID. *Left:* The MUX chip with the first two stages of the SQUID multiplexer is shown. In this view the NYQ chip is below the MUX chip and the lower wire bonds connect to the output of the NYQ chip circuit. The SQ1 bias wire bonds are shown at the top of the image, but are generally double-bonded to avoid single point failure in the series SQ1 bias lines. The 1.5Ω address resistors, gradiometric summing transformer coil, real and dummy SQ1, SQ2 and gradiometric input coil transformer are shown. Note a discrepancy: in this image the input coils of the SQ1 nearest the SQ2 on the far left is connected. *Right:* The first-stage SQUID is located on the top and the unbiased dummy SQUID is on the bottom of this micrograph. The four, octagonal windings of the SQUID loop are shown, partially covered by the input and feedback coils. Bias for the SQUIDs and output for the SQ2 enters from the top. The input coil signal enters from the connected horizontal lines on the lower right and the flux feedback lines are overlapping and enter from the bottom. The Josephson junction used in the real SQUID is the first connection — at the center of the cloverleaf — that the address lines make, which creates an asymmetry in the inductance of the two arms of the SQUID loop.

with stages further from the first stage. So the additional penalty is decreased if, for example, the SSA bandwidth exceeds that of the SQ2.

2.5.2 multiplexer chips

The version MUX07a multiplexer chips used in the BICEP2 focal plane are described in more detail in this section. The DC SQUIDS are a planar washer design with high coupling efficiency to the input and feedback coils. The SQUID loop is wound in a cloverleaf pattern to create a second-order gradiometer. The inductance is placed entirely on one leg of the loop to create an asymmetric $V - \Phi$ curve with two unique lock points.

The SQUID loop consists of four octagons wound in a cloverleaf pattern with opposing current

orientation, to create a second-order series gradiometer as shown in figure 2.5.2. This rejects external magnetic fields lacking spatial variation, but leaves the SQUID sensitive to the second spatial derivative of the field. To achieve good static field rejection the geometry of the gradiometer is designed to be as balanced as possible. The planar washer design helps to achieve this balance by focusing the flux to the $\sim 40 \mu\text{m}$ diameter interior of the loop through the Meissner effect.

The layout of the SQUID loop is asymmetric with respect to the location of the Josephson junctions. This creates an asymmetry in the inductance of the two branches of the SQUID loop. All of the self-inductance is placed on one branch and all of the bias current flows in one direction around the SQUID. The corresponding imbalance in the self-induced flux in the loop causes a distortion in the $V - \Phi$ response of the SQUIDS. There is an increase in the slope on one side of the curve and a decrease on the other.¹² This allows for two different lock points on the SQUID, one with higher gain and dynamic resistance and the other with lower gain and dynamic resistance, but longer flux base line.

The planar washer SQUID design allows for the input and feedback coils to be deposited directly above the SQUID loop for high coupling efficiency [31]. This arrangement leads to resonances in the SQUID $V - \Phi$ response, which are mitigated by the use of intracoil damping resistors [27]. Each turn of the feedback or input coils is shunted by a $\sim 1 \Omega$ resistor to alter the resonant frequency of the circuit. Resonances create low-gain lock points on the SQUIDS, which is undesirable for high amplifier gain and low and stationary noise. Despite the damping resistors, the SQ1s used in BICEP2 showed pronounced resonances, which required careful selection of the SQ1 biases and lock points (§ 4.1.8.2).

A mutual inductance naturally exists between the feedback and input coils due to their colocation. As the current in the common feedback line is switched from row to row this mutual inductance induces screening currents in the superconducting input transformer coils of all SQ1s in the same column. This in turn induces screening currents in the input coils of all the other TES bias circuits. The current in the superconducting transformer is persistent, but the TES circuits include resistance and so their flux response decays with the L/R time constant of the circuit. This flux decay modulates the transformer currents, thereby inducing a decaying input flux in each of the SQ1s. This is a potential crosstalk source between all channels in the same column (§ 4.2). To prevent this, a second set of input and feedback coils are wound in the opposite direction. Thus, screening currents induced in one coil by the other are nulled by the same currents generated in the second set of coils. To

¹²The same effect occurs when the Josephson shunt resistors are asymmetric or using additional positive feedback (APF) [14].

match the mutual inductance of the SQUID loop and coils as closely as possible, the second set of coils are wound around a dummy SQ1. This dummy SQ1 is unbiased and only serves to match the geometry of the functional SQ1 as well as possible. The SQ2 does not have a dummy SQUID, so a higher coupling coefficient exists between its feedback and input coils. The dummy SQ1 effectively reduces the coupling coefficient k_{fb-in} between the feedback and input coils. NIST measured a reduction from $k_{fb-in} \simeq 0.6$ to $k_{fb-in} \simeq 1.6 \cdot 10^{-2}$ with the introduction of the dummy SQ1 [16].

Newer versions of the multiplexer chips, including the MUX07a, use a double-transformer input coil [17]. In addition, the MUX07a and older versions used a half-integer number of windings for the input coil and feedback coils around the SQUID loop. The input and feedback coils are wound around the cloverleaf SQUID loop with 1.5 turns for SQ1 and 4.5 turns for SQ2. Since the input coils of both stages are coupled to transformer loops, an increase in the effective area ($A_{eff} = \Phi_s / \Delta B$) of the SQUIDS has been measured compared to older designs that don't include transformers [58]. The superconducting transformer loops generate screening currents for flux threading the loop that couples to the SQUIDS. The half-integer windings around the lobes of the cloverleaf leave a portion of the loop unshielded outside of the lobes of the SQUID loop. The measured effective areas of SQ1 and SQ2 for the MUX07a are $883 \mu\text{m}^2$ and $483 \mu\text{m}^2$, respectively. The earlier MUX06a version, had a factor of ~ 180 less A_{eff} for SQ1 and the same A_{eff} for SQ2, since it used an input transformer on SQ2 but not SQ1. Beginning with version MUX09a, NIST now uses whole-integer windings of the input and feedback coils around the gradiometer lobes. Now the SQ1 has two windings as opposed to 1.5 and the SQ2 has four windings as opposed to 4.5. The A_{eff} has been reduced by a factor of ~ 1471 and ~ 22 for SQ1 and SQ2, respectively.

One of the SQ1s on every column is not connected to a TES and its input coil is left open. This “dark” SQ1 has been designed to allow the removal of a common-mode signal from the other 32 channels in the column. This is potentially useful for reducing a common-mode low-frequency noise contribution of the SQUID multiplexer [16]. Due to inductive pickup between the input coils of the SQUIDS, the dark SQ1 is placed next to the SQ2. This prevents TES signals in the SQ1 input coil from coupling to all other channels in the column through the SQ2 (§ 4.1.5.1).

2.5.3 series arrays

The final stage of the multiplexer is an array of 100 series SQUIDS, as explained in § 2.5.1.2. Each 100-element SSA is lithographically fabricated on a single chip. Eight SSA chips are packaged together in a module that supports the chips, routes their electrical connections and provides magnetic

shielding. The chips are connected by wire bonds to a die that brings all electrical connections to 48-pin sockets. The modules are shielded by an inner cryogenic magnetic shield, an outer superconducting niobium shield and a multilayer high-permeability shield (§ 2.6). The SQUIDs in the arrays dissipate too much power to be located on the focal plane. Instead, they are integrated into a circuit board that is sunk to the 4 K base plate near the fridge. Superconducting cables connect the SQ2 output from the focal plane to the SSA input coils. BICEP2 uses two series array modules each amplifying the output of eight multiplexer chips.

The SQ2 to SSA circuit was believed to be the lowest pole in the amplifier and thus dominating the settling time. This in turn limits the readout rate of the detectors with a corresponding increase in the aliasing of out-of-band detector and readout noise. The BICEP2 TESs have an excess noise plateau that extends to 10 kHz and requires a high readout rate $f_{nyq} \gtrsim 10$ kHz to avoid significant aliased noise. Thus, the SQ2 to SSA circuit has been optimized for maximum bandwidth. The SQ2s are biased on the higher dynamic resistance lock point for a lower L/R time constant. The SQ2–SSA cables are shortened as much as possible to reduce the parasitic inductance. The largest contribution to the inductance in that circuit is the SSA input coil inductance (L_{ina}). At the time of the BICEP2 deployment, SSAs were available with either 1-turn or 3-turn input coils. The 3-turn coils give higher coupling and gain for the SQ2 output to the SSA, which decreases the input-referred amplifier noise. The increased inductance leads to longer settling times and lower readout rate. The increased amplifier noise is offset by the much larger decrease in aliased noise. In the end, the settling time measurements between the 1-turn and 3-turn input coils were inconsistent and 3-turn arrays were used.

2.5.4 Nyquist chips

The Nyquist interface chips low-pass filter the detector signal and provide bias resistors for the TESs. The BICEP2 focal plane uses 16 version NYQ09a chips. Each chip can service 33 TES bias circuits, but one of the circuits is left open to provide a diagnostic dark SQ1.

The detector signals are low-pass filtered prior to being sampled at the frame rate by the multiplexer. A single-pole, low-pass filter is created from the L/R circuit from the detector operating resistance in series with a total inductance $L_{tot} = 1.6 \mu\text{H}$. The Nyquist chips include $L_{Nyq} = 1.35 \mu\text{H}$ inductors with an additional inductance $L_{in} = 0.25 \mu\text{H}$ from the SQ1 input coils. The TESs were designed to operate at $R_{tes} = 30 \text{ m}\Omega$, so the bandwidth of the TES circuit, excluding ETF, is $f_{3dB} \simeq 3 \text{ kHz}$. Low in the transition, where τ_{\pm} begin to interact, the electrical bandwidth diverges

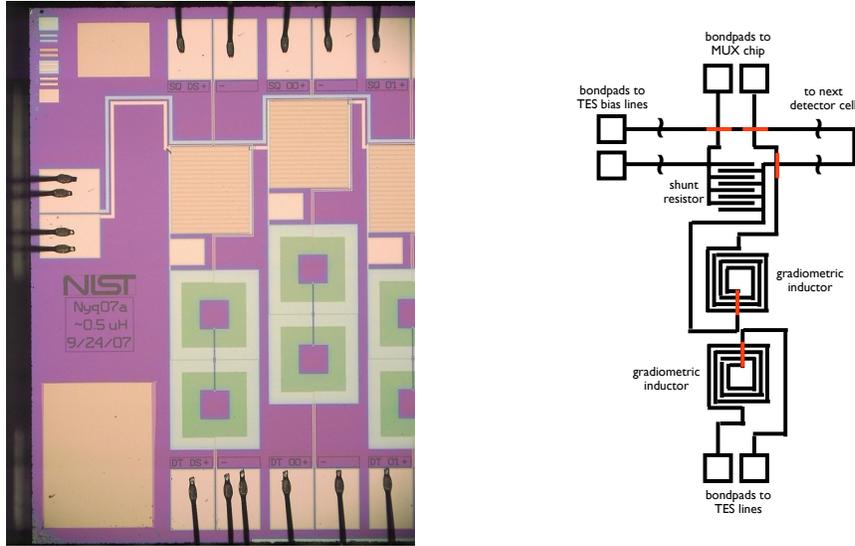


Figure 2.19: An image of a Nyquist interface chip and the interface circuit for one TES. *Left:* A micrograph of a NYQ chip showing the circuit elements and wire bonds for two-and-a-half detector cells. This is an older version of the Nyquist chip with a lower inductance value. *Right:* A schematic describing the circuit elements seen in the micrograph for one of the detector cells. Each cell includes a $3.0 \text{ m}\Omega$ resistor, a gradiometric $1.35 \text{ }\mu\text{H}$ inductor separated between the two input lines, superconducting microstrips and bond pads. The microstrips are shown in red where one crosses without making connection to the other. Note that the first circuit cell on the left is generally left open input for the dark SQ1.

from τ_{el} . High in the transition, where the poles are well separated, this is a reasonable approximation.

The inductive elements are rectangular spiral inductors wound from niobium microstrip. The total inductance is split between two spiral inductors, which are wound in opposite directions to create a first-order gradiometer that rejects ambient fields but is sensitive to their gradient. There are additional microstrips that cross over the windings to connect to the center termination of the inductor. The inductors are deposited above a niobium ground plane with square holes in the center. This geometry ensures balance of the gradiometer by the focusing of flux through the washer center due to the Meissner effect. The inductance is determined by the number of windings. Due to space constraints on the chips, the size of the ground plane limits the inductance to $1.35 \text{ }\mu\text{H}$. Higher inductance chips can be fabricated with the same form factor, but the ground plane must be eliminated.

The choice of inductance on the NYQ chips is a balance between aliasing of noise and detector stability. For an ideal detector a single-pole filter with $f_{3dB} = 3 \text{ kHz}$ should provide readout reasonably free of aliased noise for $f_{sample} \gtrsim 6 \text{ kHz}$. The high-frequency excess noise plateau of the BICEP2 detectors is not sufficiently attenuated by the filter. A higher L_{Nyq} , if it could be fabricated,

would lower the bandwidth and decrease the detector noise aliasing. However, this increases the electrical time constant of the detectors toward the effective thermal time constant. When those two time constants approach one another the detector risks instability. So, the choice of L_{Nyq} was limited both by the available inductance but also by the desire for a large margin of detector stability.

The TESs are voltage-biased by shunt resistors chosen to be $3\text{ m}\Omega$ or $\sim 1/10$ the designed TES operational resistance. The shunt resistors are located on the NYQ chips, where they are wired in series. This means 32 TESs are biased by one pair of bias lines, which has implications for the bias optimization and crosstalk (§ 4). The shunt resistors were confirmed to be $3.0\text{ m}\Omega$ in separate four-wire measurements conducted at $\sim 300\text{ mK}$.

2.5.5 electrical routing

The cold electronics are distributed between the $\sim 300\text{ mK}$ focal plane and a circuit board that is sunk to the 4 K base plate. Superconducting niobium-titanium, twisted-pair cables connect the focal plane and 4 K circuit boards, terminating on either end in 37-way MDM connectors. There are six cables that connect to the focal plane. Two cables each provide 32 wires needed for the bias and flux bias for the SQ2s on eight multiplexer chips. Two additional cables provide 32 wires for the SQ1 flux bias and TES bias on eight chips. The remaining two cables provide the 66 wires needed for the SQ1 bias, with 34 channels on one cable and 32 on the other.

Two circuit boards are mounted to the 4 K base plate near the fridge. They break out and heat sink electrical connections leading to the cryostat feedthrough, break out and heat sink the series array modules, and provide mounting points for the SQ2 bias resistors. The circuit boards are manufactured by Murietta circuits. Two boards are used, each containing a series array module, eight SQ2 bias resistors, two 100-way MDM connectors, and two 37-way MDM connectors for the bias and flux bias cables. One of the boards is used to break out the wiring for the SQ1 bias lines and includes two additional 37-way connectors.

Shielded, isothermal 100-way cables connect the 4 K boards to three 100-way pi-filtered connector modules that feed the cables through 4 K base plate of the insert (§ 2.6). Three additional shielded, isothermal 100-way cables connect the modules to the hermetic cryostat feedthrough. The feedthrough is made from circuit boards epoxied into an aluminum plate that bolts to the cryostat wall, and is equipped with 100-way MDM connectors on both sides.

2.6 shielding

The TESs and SQUIDs are sensitive to magnetic fields and must be adequately shielded. SQUIDs are the most magnetically sensitive device in the instrument, by design. The planar SQUIDs are sensitive to fields applied normal to them. The gradiometric design of the SQUID loop and the SQ1 and SQ2 input transformers reject all but second-order spatial gradients in the field ($\delta^2 B_z / \delta x \delta y$). Magnetic pickup is dominated by the SQ1 input transformer in the MUX07a chips used in BICEP2 (§ 2.5.2). Magnetic fields affect the T_c , G , and the transition width and α of TESs. However, TESs are much less sensitive to magnetic fields than SQUIDs. Both static and time-varying fields are problematic. Static differences in the flux across the series array SQUIDs can decrease the coherence and gain of the array. Spatial gradients and time-varying field fluctuations across the detector arrays can lead to variations in the TES parameters. Flux can be trapped in the superconductors during initial cool down and thermal cycling. The Lorentz force from time-varying fields or applied current can cause motion of the trapped flux, leading to increased noise. During observations scan-synchronous contamination is introduced by the changing attitude of the telescope and modulation of the local field by the ferromagnetic mount. Fluctuations that are common-mode for detectors are removed by pair-differencing or removal of a scan-synchronous template. However, differential fluctuations within a polarization pair will contaminate the polarization measurement.

The shielding requirements are set by the sensitivity derived for the TESs and SQUIDs and maximum desired contamination signal. If the contamination is limited to $\lesssim 100$ nK the magnetic field must be attenuated to < 10 pT at the TESs and < 0.1 pT at the SQUIDs of the MUX chips. To attenuate earth's $50 \mu\text{T}$ field to these levels the shielding factor must be $\sim 7 \cdot 10^6$ at the TESs and $\sim 5 \cdot 10^8$ at the SQ1s.

The magnetic shielding has been designed to attenuate stray fields at the focal plane and the SSA modules without impeding the optics. In addition, precautions have been taken to prevent the use of ferromagnetic materials near the focal plane which might distort the local field. A combination of superconducting and high-permeability cryogenic shields are used to achieve high shielding factors. Superconductors shield by expelling flux through the Meisner effect, but trap flux in strong fields when it's more energetically favorable than expelling it. The superconducting shields are all made of niobium, which has a critical temperature of 9.2 K. High-permeability materials effectively absorb nearby magnetic fields by pulling them into their bulk. Cryoperm¹³ 10 is a nickel-iron alloy optimized for high-permeability at low temperatures. It is annealed for use at 4 K, has

¹³Amunel Manufacturing Corp: www.amunel.com

a maximum DC permeability $\mu \sim 250,000$ and high saturation flux density. Metglas¹⁴ 2714A is a cobalt-based amorphous metal, with a high DC permeability of $\mu \sim 1 \cdot 10^6$ and saturation flux density of ~ 0.6 T.

The first stage of shielding is an outer 4 K Cryoperm 10 shield that surrounds the insert. This stage reduces the overall field inside the insert and reduces the potential for trapped flux. Its high saturation flux is adequate to provide the first stage of shielding. The shield is an open-ended cylinder mounted to the interior wall of the helium vessel with the insert inside it, as shown in figure 2.2. This open-ended geometry gives the best performance for fields applied along the cryostat axis. The field component parallel to the axis is the least attenuated, while the component normal to the axis is most attenuated.

The next level of shielding is the superconducting niobium spittoon that surrounds the sub-Kelvin structure (§ 2.2.3). The focal plane is further shielded by a combination of high-permeability and superconducting shields. These shields are especially important for reducing the field parallel to the cryostat axis, which is the field component the focal plane SQUIDs are sensitive to. The backshort is made of superconducting niobium and provides shielding for the MUX chips and the detector arrays (§ 2.4.7). A Metglas shield is used on the opposite side of the MUX chips from the backshort to completely enclose them in magnetic shielding. Fifteen layers of $15 \mu\text{m}$ thick Metglas 2714A were bonded together and cut into crescent shapes that fit in a recess between the PCB and detector plate. The alternating layers of Metglas were crossed to prevent overlapping seams in the layers that might allow fringing fields. To prevent flux from being trapped in the superconducting shields the spittoon and focal plane are cooled from the center. This center-point cooling is intended to push trapped flux radially outward as the shields are slowly cooled.

Variations in the input flux of the series array SQUIDs must be carefully controlled. The output voltages of the SQUIDs add in series, so differences in the input flux will cause the signal to add out-of-phase and the gain of the array will be diminished. The series arrays are heavily shielded from external magnetic fields to prevent this. The SSA modules are surrounded by a Cryoperm sleeve that slides into a superconducting niobium box. An additional shield made of several layers of Metglas 2714A is wrapped around the niobium shield. As a further precaution the location and orientation of the arrays in the 4 K Cryoperm shield was optimized. The modules were installed so that the flux-sensitive axis of the SSAs is normal to the boresight axis. In addition, the modules were elevated off the base plate so they are located further from the ends of the magnetic shields of the cryostat.

¹⁴Metglas, Inc: www.metglas.com

Simulations were run in COMSOL Multiphysics to assess the magnetic shield design.¹⁵ The simulations demonstrated that the 4 K Cryoperm shield would not saturate in a 50 μT ambient field, so would perform adequately as the first-stage magnetic shield. They also showed that for a moderate scenario — the field applied 45° to the cryostat axis — the ambient field would be attenuated to ~ 10 pT normal to the focal plane at the SQUIDs, with only a very small spatial gradient. The ambient field would be attenuated to a total magnitude at the detector array of ~ 50 nT, with a very small spatial gradient. The desired attenuation of the simulated shields did not meet the ideal shielding requirements to limit fluctuations to $\lesssim 100$ nK. However, any pickup that is common-mode in a detector pair will be rejected by pair differencing and scan-synchronous template subtraction. The major concern is for differential fluctuations or sensitivity within a detector pair that would contaminate the pair-difference polarization measurement.

A number of measures are taken to shield the focal plane and readout wiring from radio frequency interference (RFI). Radiation entering the cryostat window is filtered by the low-pass edge in the optics chain before reaching the focal plane. The focal plane is protected by a Faraday cage of aluminized Mylar that extends from the focal plane aperture to the base of the interior of the spittoon. Loading of the focal plane by the spittoon through the single layer of aluminized Mylar is negligible. The Faraday cage at the focal plane aperture is completed by wire bonds connecting the niobium ground plane of the arrays and the detector plate and PCB. The readout and thermometry wiring is protected by the use of twisted pair cables and another aluminized Mylar shield extending from the spittoon wall to the top of the optics tube, shown in figure 2.4. The readout and thermometry wiring has two levels of filtering to prevent RFI from entering the cryostat. Filtered connectors in the MCE and housekeeping backpack filter the wiring at the cryostat feedthrough. Pi-filtered MDM connector modules located at the 4 K base plate of the insert provide additional RFI protection.¹⁶

2.7 mount

BICEP2 uses the BICEP1 mount, which was left in place after BICEP1 was decommissioned in 2009. More detailed information about the mount can be found in the BICEP1 literature [61]. The mount has three rotational axes — azimuth, elevation, and theta. The mount is capable of scanning a $\pm 210^\circ$ range at a rate of up to $5^\circ/s$ in azimuth. The elevation axis has a range of 50° – 90° , and can tilt at a rate of $1^\circ/s$. These high slew rates are achieved while maintaining precise telescope

¹⁵Formerly known as FEMLAB. Comsol, Inc.: www.comsol.com

¹⁶Cristek Interconnects, Inc: www.cristek.com/pdf/6.0/FILTRRevD.PDF

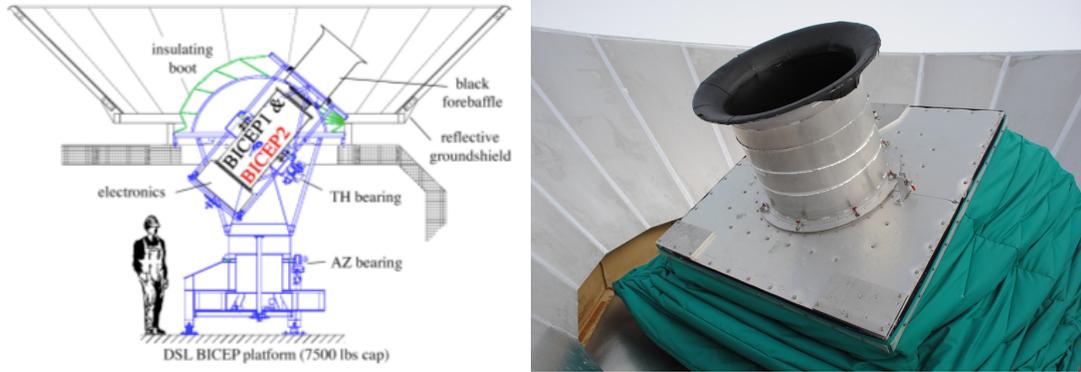


Figure 2.20: A schematic of the BICEP2 telescope and exterior view of the telescope looking out. *Left:* The schematic of the telescope shows the cryostat, mount and mount platform, exit of the telescope through the roof of DSL and the ground shield and forebaffle. *Right:* An image of the exterior of the telescope from the roof of DSL, showing the boot platform, fabric bellows, ground shield and forebaffle.

pointing and introducing only minimal microphonic vibration. The theta axis, also known as deck, rotates the mount about the boresight axis and has a 360° range. A schematic cross section of the telescope and mount, and an external view of the telescope is shown in figure 2.20.

The telescope is located on the second floor of the Dark Sector Laboratory (§ 3.2). The ~ 2.2 m diameter base rests on a platform in the telescope control room. The upper portion of the mount and cryostat passes through a 2.5 m diameter hole in the roof of the building. A flexible environmental seal allows the entire mount and cryostat to be housed in a warm environment.¹⁷ The seal is a fabric bellows attached to a support structure, the boot platform, that moves with the azimuth and elevation motion of the mount. A brush seal between the bottom of the boot and the roof top provides an environmental seal to the control room.

The low-vibration mount motion is accomplished using ultra quiet crossed roller bearings, gearless cycloidal motor reducers, and toothed belt drives. The motion of the mount is driven by three DC brush motors each powered by an amplifier. The mount position is measured by three encoders. Limit switches prevent extreme rotation of the axes that might damage the telescope or cabling. Two orthogonal tilt meters are mounted to the azimuth stage to monitor the tilt of the azimuth axis. Each rotational axis has a cable wrap to route cabling to the receiver. The azimuth cable wrap is located in the base of the mount and limits the azimuth range to $\pm 210^\circ$. A custom theta cable wrap replaces the slip ring used in BICEP1 that allowed for continuous theta motion. The cable wrap restricts the theta axis range to 360° and requires it to be rotated at zenith at low speed. Observations occur at fixed deck angles, so low-vibration theta rotation is not required. Control of

¹⁷The author is well acquainted with cryogenic transfers outside when the temperature is -100°F .

the mount motion is described in § 2.9.2.

2.7.1 ground shield and forebaffle

A combination of stationary reflective ground shield and comoving forebaffle protect the receiver from stray radiation, as shown in figure 2.20. Any radiation from the ground must diffract at least two times before entering the cryostat. This prevents the polarized side lobes of the beams from coupling to the external environment, including the ground, the roof of the building and the neighboring South Pole Telescope.

The ground shield is a 45° truncated cone of aluminum panels with a diameter of ~ 4 m at the top of the structure. The shield extends 2.2 m above the roof of DSL and prevents any line of sight coupling to the receiver window. Seams between aluminum panels in the ground shield are taped with reflective aluminum tape to present a virtually uninterrupted reflection of the sky to the telescope. The rim of the shield is rounded to minimize diffraction.

The comoving forebaffle is composed of an aluminum cylinder that opens to a flared lip to minimize diffraction at the forebaffle aperture. The interior is baffled with a layer of microwave absorber, Eccosorb HR10, to prevent scattering of stray light. The Eccosorb is encapsulated in a 1.6 mm thick, tear-resistant polyethylene foam (Volara) to protect it from the environment. The forebaffle is heated slightly above ambient temperature to prevent snow from accumulating on the surface, which would increase the optical loading of the detectors. The forebaffle is attached to the top of the cryostat and is flush with the plane of the boot platform. As the cryostat performs a deck (theta) rotation the forebaffle rotates as well, so there is no relative motion between it and the detector beams.

2.8 optical star camera

An optical star camera is used to precisely determine the boresight pointing (§ 3.6.2). This is the same camera used by BICEP1 and additional information can be found elsewhere [62][59]. There are 10 free parameters in the pointing model so the camera was designed to find at least 20 stars, even in the summer. The camera consists of an Astrovid¹⁸ StellaCam EX CCD and 100 mm diameter refractive lens enclosed in a rigid housing. The camera housing is mounted to the top of the cryostat near the window and is coaligned with the boresight axis. The sensor and antireflection coatings

¹⁸AVA Astro Corporation: www.astrovid.com

have been optimized for infrared light, in order to view stars in daylight. The camera has a 2' resolution, 0.5° field of view and is sensitive enough to detect magnitude 2.8 stars in daylight.

2.9 warm electronics

The warm electronics include the MCE, housekeeping electronics and motion control electronics. These interface with three control computers that command the components and manage the data acquisition.

2.9.1 control overview

Control and data acquisition are managed by a combination of the generic control program (GCP), the MCE acquisition system (MAS) and various drivers. GCP is a modular software package, which can be distributed across multiple control computers.¹⁹ It combines data from the different packages and provides a control interface. GCP includes four main software packages that communicate via socket connections: control, mediator, antenna and the MAS daemon. A graphical interface called bicepViewer is used to control GCP. The control software is the top-level layer that manages the bicepViewer connections, interprets and executes user commands and scheduling scripts, and saves data. The mediator software is an intermediate layer between the control and low-level antenna layer. It combines data from multiple sources into a single data stream that is passed back to the control for archiving and distributes commands sent by the control layer. The low-level antenna layer communicates with the device drivers that control the telescope (PMAC), housekeeping (BLASTbus) and GPS timing. The MAS daemon interfaces GCP with the MAS software that controls the MCE and collects detector data.

There are three Linux-based control computers: bicep31, bicep32 and bicep33. The bicep31 computer runs the antenna software and includes four PCI cards that provide motion control, GPS timing and housekeeping data. The bicep32 computer runs the MAS daemon and MAS software and includes a PCI card that interfaces with the MCE. The bicep33 computer runs the control and mediator software packages and has the frame grabber card that interfaces with the star camera.

The data from the various sources — housekeeping, motion control and MCE — must be combined synchronously and stamped with an acquisition time. The master timing is provided by the MCE synchronization box (§ 2.9.5). Every ~ 100 Hz the sync box produces a timing signal and unique frame sequence number. The number and timing signal are transmitted to the MCE and

¹⁹GCP was originally developed by Erik Leitch.

BLASTbus card on bicep31. The timing signal triggers the MCE to report detector data, stamped with the sequence number, back to the PCI card on bicep32. It also triggers the BLASTbus card to pull housekeeping data stamped with the sequence number from the BLASTbus crate. The timing signal on bicep31 also pulls the time from the GPS card, and pointing data from the PMAC Turbo card. The data streams are collated and archived by GCP in a frame-based file format referred to as an “arc file”.

2.9.2 motion control

Motion control of the telescope is controlled by the programmable multiaxis controller (PMAC) system that integrates commercial and custom components. The system includes a custom PMAC breakout box, Turbo PMAC and Interpolator PCI cards²⁰, interlock box and a manual motion control paddle. The Turbo card contains the computer that controls the servos for the three mount axes. It controls the motor amplifiers based on the requested axes motion and feedback from the encoder positions. The Interpolator card accepts the output of the analog encoders and provides encoder position data to the PMAC card. The PCI cards are located in the bicep31 computer and are connected to the PMAC breakout box. The breakout box accepts additional inputs from the tilt meters, encoders, flow meter, and auxiliary inputs that can be used to read in reference signals during calibration tests.

Undesired motion of the mount is curtailed by a custom servo interlock box. The interlock monitors signals from limit switches and emergency stops distributed on the mount. In the case of a fault the interlock will lock out the motor amplifiers and signal the PMAC breakout box. When the telescope is run in manual mode instead of servo mode, the motion can be controlled by motion control paddle.

2.9.3 housekeeping

The housekeeping subsystem is used for control and monitoring of the thermometers and heaters in the telescope. It includes heaters and thermometry, a custom housekeeping chassis called the backpack, the BLASTbus and its PCI card. The housekeeping backpack contains custom cards to bias and read out thermometers, control heaters and provide temperature control servos. It mounts to the bottom of the cryostat and interfaces with MDM connectors located at the cryostat wall. Signals from the backpack are routed to the rack-mounted BLASTbus crate. The BLASTbus is

²⁰Both PCI cards are products of Delta Tau Data Systems, Inc.: www.deltatau.com.

a data acquisition and control system built for the BLAST²¹ experiment by collaborators at the University of Toronto [48].²² Signals from the crate are routed through a three-way hydra cable to the BLASTBus PCI card located in bicep31.

Resistive and diode thermometers and resistive heaters are mounted at various locations on the cryostat, optics, fridge, and sub-Kelvin structure. During engineering tests the heaters and thermometers are used to measure the thermal conductance between temperature stages. Temperature control modules (TCMs) are located on the focal plane detector plate, thermal strap connecting to the focal plane filter block and the UC fridge still. The TCMs include two NTD thermistors, two heaters, and bias and readout circuitry. They are used for either high-sensitivity temperature readout or a temperature control servo. The use of the focal plane TCMs and tile thermistors is explained in § 2.4.2.6. The JFET readout for the NTDs is located in a readout module mounted to the 4 K base plate. The module also provides a breakout circuit board for the thermometry wiring.

To reduce low-frequency noise in the thermometer readout, the resistive thermometers and NTDs are AC biased. AC bias sine waves are generated by the BLASTbus and converted to analog AC bias by the housekeeping backpack. The thermometer output is demodulated by the BLASTbus. Analog signals are passed from the housekeeping backpack to the BLASTbus crate. The signals are digitized by ADCs and pulled from the BLASTbus by the BLASTbus PCI card at ~ 100 Hz.

2.9.4 GPS timing

An accurate time stamp for the data is provided by the GPS timing subsystem. It consists of a Symmetricom²³ GPS PCI card and an external antenna that receives timing information from the GPS system. The PCI card is plugged in to the bicep31 computer and controlled by a custom driver on that computer. The card outputs an 8 kHz square wave and a 1 pulse-per-second signal that are input to the PMAC breakout box for timing in the motion control system.

2.9.5 multi-channel electronics

The multi-channel electronics (MCE) provided by collaborators at UBC²⁴ provides the room temperature electronics that manage the SQUID multiplexer [3]. The MCE provides the bias and flux bias for all of the SQUID stages, handles the FLL of the SQ1 feedback and packages and returns

²¹Balloon-borne large aperture submillimeter telescope

²²The group is led by Barth Netterfield.

²³Symmetricom: www.symmetricom.com

²⁴University of British Columbia. The group is led by Mark Halpern.

Table 2.3: MCE electrical output

	SSA bias	SSA fb	SQ2 bias	SQ2 fb	SQ1 bias	SQ1 fb	TES bias
card	readout	bias	bias	address	address	readout	bias
DAC bits	16	16	16	14	14	14	16
R (k Ω)	3.1	5.2	2.0	10	2.1	6.3	0.4
V	0.5	2.5	2.5	1.0	1.0	1.0	2.5
I (mA)	0.2	0.5	1.0	0.1	0.4	0.1	4.0

the detector data to GCP. In addition, a synchronization box provides unique serial numbers for merging of the different data streams.

2.9.5.1 hardware

The MCE consists of a custom 6U subrack with eight custom cards. It is controlled by the bicep32 control computer through a PCI²⁵ interface card. The computer and MCE communicate using two fiber optic cables. The MCE is powered by a 24 V linear power supply. The chassis mounts directly to the bottom flange of the cryostat at the cold electronics feedthrough (§ 2.5.5). Three filtered, 100-way male MDM connectors on the back of the MCE engage with the female MDMs of the feedthrough. The crate is populated with one power card, one clock card, two bias cards, two readout cards and two address cards. The cards plug into two backplane circuit boards — the instrument backplane and bus backplane boards — built into the chassis that distribute signals and power among the cards. Synchronization of the merging of detector data with data from the other warm electronics is managed by the MCE Synchronizer, or sync box, which communicates with the MCE over a third fiber optic cable.

The power card connects to the power supply and distributes power to the other cards through the backplane. The clock card controls timing in the MCE, distributes commands to the other cards, and packages and transmits data to the control computer. It connects to the fiber optics from the PCI card and sync box. One address card controls the bias for all SQ1s, while the other biasing-address card provides the fast SQ2 flux bias for all 16 columns. One bias card controls the TES bias and the other bias card controls the SQ2 bias and SSA flux bias. The two readout cards bias and sample the SSAs, control the SQ1 FLL servos and pass data to the clock card. There is one readout card per eight multiplexer columns.

The instrument backplane circuit board routes signals from the cards to the three filtered con-

²⁵peripheral component interconnect by Astronomical Research Cameras: www.astro-cam.com

nectors. The SQ1 bias lines are routed to a single filtered connector, P1. The SQ1 flux feedback and bias, and TES bias lines for columns 0–7 are routed to connector P2, and to connector P3 for columns 8–15. The bias resistors for the output lines are located on the backplane, connected in series with the signal lines of each output channel. The return lines for each channel are grounded to the chassis. The series bias resistors, maximum full scale voltage output and corresponding bias current are summarized for each channel in table 2.3. The connector filters are circuit boards with 100-way MDM connectors housed in an aluminum Faraday enclosure. The filter connectors are on a drive shaft to engage the cryostat connectors after installation of the MCE. The circuit boards connect to the instrument backplane by flexible kapton-encapsulated connectors.

The cards include a combination of ADCs and DACs and an FPGA to control them. The bias cards have 33 2.5 V, 16-bit DACs to apply detector and SQUID bias and flux bias. The address cards have 41 14-bit DACs capable of switching the SQ1 bias or SQ2 flux bias at the row visitation rate. The readout cards include eight pre-amplifiers and 14-bit video ADCs to sample the output of the SSAs, a 14-bit DAC to output the servoed SQ1 flux feedback, and two 16-bit DACs to control the SSA bias and cable offset. The FPGA computes the necessary flux feedback using a firmware PID servo, and relays filtered data to the clock card for transmission back to the control computer. The pre-amplifier has four stages to provide gain to and low-pass filter the output of the SSAs before they are digitized by the ADCs. The filter has six poles that combine for a bandwidth of ~ 6 MHz in the revision B9 cards. The first stage is a differential amplifier, with the negative input coming from the “cable offset” 16-bit ADC to null the voltage offset from the SSAs and keep them from railing the ADCs.

The sync box is a separate rack-mount chassis that generates a data-valid pulse and sequential 32-bit frame sequence number to synchronize the acquisition of the MCE data and housekeeping data. The frame sequence number and timing pulse are transmitted to the MCE clock card over a fiber optic. The detector data are packaged into “frames” along with the frame sequence number and are transmitted to the PCI fiber card at every data-valid pulse. The pulse indicates when the multiplexer has cycled through all of the rows and returned to the beginning of the frame every $N_{mux} \times row_len$ clock cycles (20 ns each). The sync box accepts commands from the bicep31 computer over an RS-232 serial cable and reports the frame sequence number to bicep31 through the BBC PCI card. The *row_len* and *num_rows* parameters are assigned by bicep31. The sync box operates on a 25 MHz clock, so the specified *row_len* parameter must be half the value specified for the MCE’s 50 MHz clock. When the sync box is enabled the timing in the MCE cards and transmission of data to the

PCI card is controlled by the data-valid pulse, rather than the synchronization pulse generated by the clock card.

2.9.5.2 software

The MCE software includes MAS, firmware programmed into the card FPGAs, PCI and sync box firmware and shell scripts to send higher level commands to the MCE. The MCE is controlled and communicated with via low-level control software called the MCE acquisition system (MAS). The MAS software is run on bicep32. Basic commands and responses are initiated at the command line in linux using the *mce_cmd* program, which can be run in an interactive mode or through shell scripts. The PCI card driver can also be communicated with through shell commands.

2.9.5.3 FLL servo

For a small input coil signal there is a small change in the output voltage of the SSA. The signal is amplified, low-pass filtered, and digitized by the readout card. The output signal is passed to the FPGA. The error signal, which is the difference between the open-loop amplifier lock point (*adc_offset*) and the amplified output, is sampled at the 50 MHz clock rate. The error signal is oversampled, with (*sample_num*) samples coadded together. The output is passed to the PID servo algorithm that calculates the flux feedback to be applied to the channel on the next visit after the completion of the frame. The algorithm is:

$$fb_{n+1} = P \times coadd_n + I \times \sum coadd_i + D \times (coadd_n - coadd_{n-1}) \quad (2.6)$$

where fb_{n+1} is the feedback applied at the next visit ($n+1$) of the row, as derived from the current (n) and prior ($n-1$) visits. The coadded error signal is $coadd_i = \sum err_j = \sum adc_output_j - adc_offset$, with the sum j running over the number of samples, set by the *sample_num* parameter. The prefactors P , I , and D are the proportional, integral and derivative terms for the PID loop and are specified by the user. The PID loop is typically operated in a pure-integral state. The 32-bit output is truncated to 14 bits and stored until the next row visit, at which point it is passed to the 14-bit DAC to apply the nulling flux feedback to the SQ1 of the row. Two additional processing steps are performed on the same output, depending on the acquisition settings. If flux jump correction is turned on an additional correction to the signal is made before it's passed to the DAC. If filtered feedback is requested fb_{n+1} will be low-pass filtered. Each of the 528 output channels has an independent PID servo loop running, with 33×8 loops running on each readout card.

2.9.5.4 DAC output correction

The 14-bit signed DAC output has an output range of -8192 to +8191. When the output of the feedback DAC approaches the limits of this range the DAC output must be offset toward the center of the range to prevent raling. When flux jump correction is enabled ($en_fb_jump = 1$) the corrected DAC output uses the measured flux quanta of the SQ1s to relock the servo at the same point on the SQ1 $V - \Phi$ response. This will maintain continuity of the amplifier output. The quanta are measured from the $V - \Phi$ response of the SQ1s and specified as the period of the SQ1 at the lock point in feedback DAC units. The adjusted flux feedback DAC output for the next frame visit is:

$$fb_{adj,n+1} = fb_{n+1} - j_n \times Q_{sq1} \quad (2.7)$$

where j_n is a flux jump count updated for the current visit (n) and Q_{sq1} is the measured flux quantum for the SQ1 of interest. The flux jump count j is incremented by 1 when $fb_{adj} > 7800$, which is 95% of the maximum DAC output and is decremented by 1 when $fb_{adj} < -7800$, with a corresponding adjustment of the DAC output. The PID servo output is a signed 32-bit number, so the 12 least significant bits (LSBs) are truncated and equation 2.7 is implemented on the remaining most significant bits (MSBs). The flux jump counter is thereby limited to 2^7 . The 14 least significant bits of the adjusted output are then passed to the DAC. When flux feedback is disabled ($en_fb_jump = 0$), the output is adjusted when the limit of the DAC is reached. The procedure then is to take the 32-bit PID output, drop the 12 LSBs and pass the 14 LSBs to the DAC. This DAC adjustment brings the DAC output back in to range with no regard to regaining the SQ1 lock point in a continuous manner.

2.9.5.5 filtered feedback

The feedback data is low-pass filtered in certain data modes so that it can be downsampled without aliasing before it is transmitted to the control computer. The low-pass filter is a four-pole Butterworth filter implemented in firmware as two cascaded biquads. The downsampling factor is specified by the parameter *data_rate* and gives a final sampling rate of the detectors:

$$f_{sample} = \frac{50 \text{ MHz}}{N_{mux} \cdot row_len \cdot data_rate} \quad (2.8)$$

The filter coefficients are hard-coded in firmware to give a constant ratio of the 3 dB cutoff frequency to the final sampling frequency, $f_{cutoff}/f_{sample} = 0.0025$. Thus the filter scales with the user-defined

sampling parameters to maintain a cutoff below the Nyquist frequency of the sampling rate. The DC gain of the filter has been simulated to be 2044, accounting for quantization effects.

2.9.5.6 raw and rectangle readout modes

The MCE can be configured for special data taking modes in order to achieve readout rates higher than the frame rate. These readout modes are especially useful for assessing the high-frequency components of the detector and readout noise (§ 4.1.6). A raw readout mode allows a single row for each column to be sampled at the 50 MHz clock rate. The PID servo is turned off since these calculations cannot be made at the clock rate. The raw data are just the direct output of the ADCs that have not been differenced by the *adc_offset*. The data are limited to the size of the buffer and only 8192 samples can be acquired per column. If a single column is read out then 65536 samples can be acquired. The MCE can be run in either a multiplexed or non-multiplexed configuration when acquiring raw data. The former allows for the study of the settling of the amplifier as the rows are switched. The latter allows the open-loop output of the amplifier or detectors to be studied without interference from switching transients. The readout can be configured for rectangle mode, which allows a subset of rows and columns to be multiplexed and sampled at an increased frame rate. This mode is typically run at a 400 kHz sampling rate on a single row. The settling of the multiplexer must be considered in this mode. The readout is restricted by the hardware to 4 MB/s.

2.9.5.7 data modes

The user-specified *data_mode* determines the data output from the readout cards sent back to the control computer. The data output is a 32-bit word for each readout channel at each row visit. These data are packaged together for all channels into frames at the frame rate. When data are acquired in raw mode the output is the amplified signal from the ADCs prior to being differenced with the *adc_offset*. In error mode the data are the differenced and coadded error signals. In feedback data mode the output is the result of the PID algorithm in equation 2.6. The output of the low-pass filter is the filtered feedback, which is not corrected for the DC filter gain. The flux jump correction setting alters the feedback data that is reported relative to that applied by the DAC. In some data reporting modes the flux jump counter j is reported, which allows the user to calculate the adjusted DAC output using equation 2.7. Different combinations of these data can be output to the control computer depending on the *data_mode* selected.

2.9.5.8 timing sequence

Timing units are given in terms of the number of 20 ns clock cycles. The time the multiplexer dwells on each row is given by a user-configurable parameter, the *row_len*, which corresponds to τ_{dwell} from § 2.5.1.3. After a frame period of $N_{mux} \times row_len$ clock cycles a synchronization pulse is sent by either the sync box or clock card and the multiplexer begins a new frame. Before turning on the SQ1 bias for a row the multiplexer waits for *row_dly*, which is fixed in the firmware at 4 clock cycles. The bias is then turned on for that row for the remainder of the row visit. This delay prevents the simultaneous switching on and off of two SQ1s on a column, which can lead to larger switching transients. The multiplexer then waits *sample_dly* clock cycles before the ADCs on the readout cards are sampled. This parameter is configured by the user and can be as little 0 clock cycles and as long as *row_len*. After *sample_dly* clock cycles the ADCs are sampled for *sample_num* cycles, which can be set between 0 and *row_len*. The user must take care when selecting *sample_dly* and *sample_num* so that they sum to no more than the *row_len*. Additional samples over the *row_len* will not be averaged. The flux feedback from the previous row visit continues into a new row for *fb_dly* cycles before the flux feedback is switched to the value for the new row. The *fb_dly* can be as little as 7 cycles when flux jump correction is turned off and 18 cycles when it is turned on. The 11 cycle difference is due to the additional calculation requirement for the DAC output with flux jump correction enabled. After the N_{mux} SQ1s have all been multiplexed, which takes a frame period, a synchronization pulse is sent by the clock card and a new frame is started.

2.9.5.9 SQUID composition

In order to operate the multiplexer, bias and flux bias values must be selected for each stage of the multiplexer. The biases are selected to optimize the gain, linearity and dynamic range of the SQUID amplifiers. The flux biases are selected to set the SQ2s and SSAs at optimal lock points. The maximum gain of a SQUID ($\delta V/\delta\Phi$) occurs somewhat below I_c^{max} , while there is still some superconducting current flow without voltage for a range of input flux. However, the modulation depth and the linearity is reduced below I_c^{max} so the bias selection is a compromise between gain, linearity and dynamic range. Since the MUX07a SQUIDs are asymmetric, two lock points exist: a lock point with high gain and dynamic resistance and one with low gain and dynamic resistance. High gain improves the output gain of the amplifier and thus reduces its input-referred noise. In addition, the higher dynamic resistance lowers the L/R time constant of the SSA-SQ2 circuit. The decreased settling times allow for a higher detector sampling rate and reduced aliased noise. However, this

steep slope has a drawback.

Selection of the lock points must also take into account the composition of the $V - \Phi$ response of the three SQUID stages. It is important to avoid wrapping in the total composed $V - \Phi$ curves because it can present multiple lock points with the same amplifier output value to the FLL servo loop. This can result in unwanted jumping between lock points, or locking on a point with lower gain or dynamic range than the intended composed lock point. Avoiding nonstationary behavior of this sort is especially important to the noise model in the final analysis.

The voltage response of SQ1 projects across the flux input of SQ2 and the SQ2 voltage response projects across the SSA flux input. If the flux projection of SQ1 onto SQ2 is large compared to the flux required to go from peak to trough on the lock slope of SQ2 then the composed SQ1–SQ2 response will have kinks at the turn around points of the SQUIDs. For large SQ1 modulation depths it's best to select a lock point half way in flux space between the peak and trough to minimize the wrapping. Selecting the lock point on the shallower slope of the asymmetric SQ2 will further reduce wrapping since it has a longer flux baseline. This must be balanced with the need for high gain and decreased settling time. If the SQ1 flux projection onto SQ2 is small, then the steep slope can be selected for higher dynamic resistance, gain and linearity, without increased wrapping. The amplitude of the composed SQ1–SQ2 response is limited by the modulation depth of the SQ2. If it is large compared to the peak-trough flux baseline of the SSA, then the SSA lock point should be selected at the midpoint of the baseline to limit wrapping. Otherwise, a lock point with higher gain, linearity and dynamic resistance should be selected. The BICEP2 SSAs are symmetric, so the positive and negative lock points are equivalent.

2.9.5.10 fast SQ2 flux bias

The mean projection of a SQUID stage onto the subsequent stage depends on the current and coupling of its output inductor with that stage. The flux bias for subsequent stages must be selected after the bias of the forward stage is turned on. Only one SQ2 couples to each SSA, so only one flux lock point needs to be selected for each SSA. However, there are 33 SQ1s coupled to each of the 16 SQ2s, each having a unique mean flux projection onto the SQ2. This can lead to suboptimal projection of the SQ1s onto the SQ2s if only one flux bias is selected for each SQ2. The problem is prevalent if the projection of the modulation depth of SQ1 is large compared to the SQ2 flux quanta. Therefore an optional fast SQ2 flux bias switching mode has been implemented to use a different SQ2 flux bias value for each SQ1 to reduce multipole lock points in the total composed

response. The “fast” SQ2 flux bias switches as the multiplexer cycles through rows. This comes with a potential cost of increased settling times since the SQ2 feedback values must also slew.

2.9.5.11 SQUID tuning

A semi-automated SQUID tuning algorithm has been developed in IDL²⁶ by the UBC team [3]. The user must first select SQUID biases, but then an automated tuning of all three SQUID stages is made in the following way. All biases, flux feedback, multiplexing and FLL servoing are initially turned off. The SSA biases are turned on and the SSA flux feedback is swept across the range of the output DAC. The error signal plotted against the flux feedback gives the $V - \Phi$ response of the SSAs. From these curves a lock point is selected at the midpoint between the peak to trough response along the SSA flux axis. This is done since the modulation depth of the SQ2s is large compared to the SSA flux quanta for the BICEP2 MUXs.

The SQ2 biases are then turned on and the SQ2 flux feedback is swept. A software servo is used to maintain the lock points of the SSAs as the SQ2 feedback is swept so that only the $V - \Phi$ response of the SSAs is acquired. Flux is fed back to the SSAs to keep the error signal at the value corresponding to the SSA feedback selected in the previous step. The lock points for the SQ2s are selected to be the midpoint between the peak and trough along the SQ2 flux axis on the longer baseline of the shallow slope, since the modulation depths of the SQ1s are large compare to the SQ2 flux quanta. This choice comes at the expense of increased settling times but at the benefit of a large reduction in multiple lock points. Since turning on the SQ2s changes the input flux to the SSAs the final SSA flux bias is selected from these new lock points and is then applied.

The same procedure is repeated for selecting the SQ1 lock points. Row multiplexing is turned on but the FLL servo is left off. The SQ1 biases are turned on and the SQ2 flux bias is servoed to maintain the SQ2 lock point as the SQ1 flux feedback is swept. The SQ1 $V - \Phi$ response curves are the SQ2 flux bias versus the SQ1 flux sweep. The SQ1 lock points are selected at the midpoint of the peak to peak $V - \Phi$ response to symmetrize their flux projection onto the SQ2s. For the previous steps there were only 16 lock points selected for the SSAs and SQ2s. There are 528 lock points for the SQ1s with 33 different SQ1 lock points for each SQ2. With fast SQ2 flux bias enabled the SQ2 flux bias is selected from those unique lock points. When it is disabled the SQ2 flux bias is selected from the mean of the lock points of the SQ1s in the column. The selected SQ2 flux bias is now applied.

At this point, the system has selected and applied the SSA and SQ2 flux bias values that are

²⁶Interactive Data Language: www.exelisvis.com

needed to achieve the selected lock points on each stage. The SQ1 flux feedback is swept again, but this time without servoing the SQ2s. This yields the total composed $V - \Phi$ response for each channel (the so-called SQ1 ramp) in ADC output versus SQ1 flux sweep. The desired open-loop lock point for each channel is selected from these curves. In order to minimize multiple lock points the midpoint between the wrapping extrema is selected. This ADC value is nulled in the differential amplifier and a final open-loop SQ1 flux feedback sweep is made to confirm the lock points. The total open-loop gain — the differential SQ1 flux feedback to ADC output — is determined from the slope at the lock point on these curves as well as the presence of multiple lock points. The steeper open-loop slope corresponds to the steeper slope of the SQ1 $V - \Phi$ response, which is selected through the sign of the servo gain. The FLL servo mode is now turned on so that the selected ADC values are acquired and all input flux to the SQ1s is nulled with servoed SQ1 flux feedback. At this point the system is tuned and ready for data acquisition. Alternatively, the tuning files produced by this algorithm can be used to put the MCE into a specific tuning configuration.

Chapter 3

Observations

3.1 overview

BICEP2 was deployed to the South Pole station in December 2009. After a short period of integration and optimization, scientific observations began in January 2010. The instrument has been continuously observing since then, apart from 1–2 month interruptions during the austral summers for calibration measurements (see § 5). BICEP2 is currently in its third observing year, and will conclude observations in late 2012. At that point a final round of characterization measurements will be acquired and the telescope will be decommissioned. The cryostat has been kept cold continuously through this time period, and no mechanical changes have been made to the telescope apart from periodic adjustment of the azimuthal tilt of the mount.

3.2 observing site

BICEP2 is located at the geographic South Pole near the Amundsen-Scott South Pole Station, a permanent US facility administered by the National Science Foundation. The instrument is located in the Dark Sector, a radio-quiet region ~ 1 km from the station, on the second floor of the Dark Sector Laboratory (DSL) building, which it shares with the 10 m South Pole Telescope (SPT). The combination of high altitude, low temperature and dry atmosphere makes the South Pole one of the premier millimeter-wave observing sites in the world. The South Pole is located on the high polar plateau at 2835 m elevation. Due to a constant high-pressure zone the effective altitude is even higher, with a mean pressure altitude of 3200 m or 681 mb [38]. The temperature is an average of 49°C and the lowest recorded temperature was -82.8°C . The precipitable water vapor (PWV) at the South Pole is measured daily using balloon-borne radiosondes, and is less than 0.5 mm over



Figure 3.1: An image of the Dark Sector Laboratory, housing BICEP2 on the top of the building to the right and the South Pole Telescope on the left.

half of the time. The next driest premier observing site in the Atacama desert has higher PWV in the lowest 25% quartile in the winter than the 75% quartile in the summer at the South Pole, as shown in figure 3.2. The optical depth (τ) at the South Pole is also extremely low with a value of $\tau \sim 0.03$ at 150 GHz measured by BICEP1 [61]. The opacity is actually lower for the best days in the Atacama, but it is low for a larger fraction of the days at the South Pole [50].

At a latitude of 90° there is only one long day of approximately 5.5 months and a long night of equal length, with approximately two weeks of sunrise and sunset in between. The lack of diurnal heating and cooling makes for very stable observing conditions at the South Pole, especially during the austral winter. This, along with katabatic wind patterns above the South Pole, give it a very stable atmosphere. One of the most important traits of the site is the polar latitude, which keeps an observing field fixed in elevation and continuously visible. This, along with weather that's extremely stable even during the summer, allows BICEP2 to continuously integrate its CMB field year round.

The Amundsen-Scott South Pole Station is populated by 250 people during a short summer season lasting from the beginning of November until mid-February. Outside of this window the temperatures are too extreme for substantial aircraft to service the station, and so the station becomes effectively closed to the outside world. During the remainder of the year the winter population drops to between 30 and 40 scientists and logistic staff. The lack of direct access to the telescope during the winter months generally ensures a continuous period of observation uninterrupted by calibration measurements and optimization. It also requires a winter-over scientist to maintain the telescope, a job that has been expertly performed by Steffen Richter for the full three years of observation. The

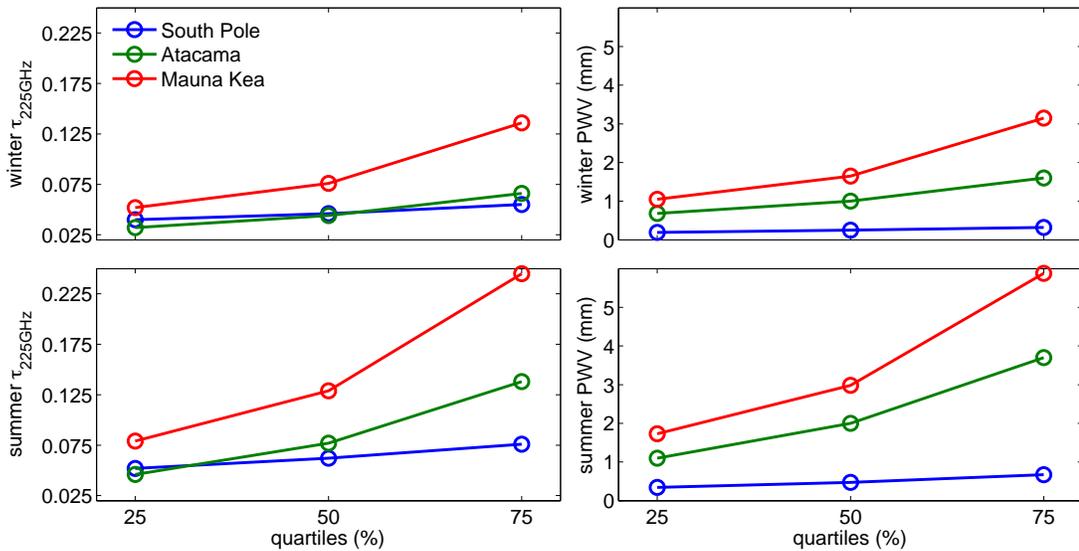


Figure 3.2: Optical depth at 225 GHz (left) and precipitable water vapor (right) at the leading millimeter-wave observing sites versus % quartile. The winter months at the observing sites (top) provide better observing conditions than the summer months (bottom). The South Pole and Atacama desert are the two best operational mm-wave observing sites in the world. At the $\sim 40\%$ quartile there is a crossing for the two sites. The best days at the Atacama have lower τ than the best days at the South Pole. However, the South Pole enjoys a larger fraction of good days. The PWV plots show that the South Pole is the driest of the three sites.

telescope can be indirectly accessed during several periods a day through low-bandwidth connections to the TDRS and GOES satellites. The science data are transmitted to United States daily over a dedicated satellite uplink.

3.3 scan strategy

BICEP2 observes by scanning in azimuth at fixed elevation. The elevation is stepped 0.25° over a range of 55° – 59.75° . At each step the telescope scans to the right and back 53 times over a 64.4° range. The 53 scans, or 106 left- and right-going half-scans, form a scan set. During the ~ 50 min scan set the telescope is continuously scanning in azimuth. As the center of the field drifts by during the scan set, the center of the scan pattern does not track the drift. The azimuth range is held fixed during the scan set and adjusted to the updated field center at the next elevation step. Thus, sources of contamination — such as ground pickup, magnetic pickup, and scan-synchronous thermal or microphonic pickup — remain fixed in azimuth. This contamination can be removed by projecting out a scan-synchronous template for each azimuth-fixed half-scan in the final analysis. This removes CMB signal on large angular scales, but the mode loss for $\ell > 40$ does not pose an

issue.

To give complete polarization coverage the boresight is stepped between four angles separated by 45° every 72 sidereal-hour observing block. The angles (68° , 113° , 248° , 293°) were chosen for their thermal stability. They give two independent and complete measurements of Q and U and allow for a jackknife test between angle sets. Rotation of the boresight has the added benefit of reducing systematic contamination.

A scan speed of $2.8^\circ/\text{s}$ was selected to project the target multipole range ($\ell \sim 30\text{--}300$) to a frequency range of 0.1–1 Hz at $\sim 57^\circ$ elevation. This modulates the signal above most of the $1/f$ noise of the atmosphere and well above that of the detectors. The bandwidth of the detectors and multiplexer servo is greater than 100 Hz and does not significantly roll-off the detector signal at this scan speed. This rapid scan rate is achieved without inducing motion-driven microphonic pickup in the detectors. The portion of the scan used in analysis is restricted to a region of linear velocity that comprises 77.7% of the scan, limiting the azimuth range to $\pm 25^\circ$.

Raster scanning at the South Pole gives no cross-linking since there is no sky rotation. This means that certain spatial modes in the CMB are not observed. However, the success of previous experiments operating with the same scan strategy at the South Pole indicates that this hasn't been a problem. BICEP1 used the same observing strategy and currently has the best direct constraint on the scalar-to-tensor ratio [13]. Furthermore, in the case of a noise-dominated experiment in the presence of $1/f$ noise raster scanning provides optimal sensitivity [15].

3.4 observing fields

The goal of detecting or constraining primordial B-mode polarization requires careful choice of the observing region in order to avoid confusion from polarized foregrounds. As described in § 1.3 gravitational lensing, synchrotron radiation and thermal dust emission produce B-mode polarization. The size of the region is a balance between lensing confusion at small angular scales and Galactic foreground confusion at large angular scales. The gravitational lensing signal rises with ℓ , with the ratio of lensing to primordial signal growing rapidly above $l \sim 100$, which calls for larger coverage to avoid. The Galactic foregrounds power spectra scale as $\sim l^{-0.6}$. This, along with the especially high contrast of dust emission, calls for smaller coverage. The coverage area also affects the sample variance in the detection of the CMB signal. A larger patch gives a greater number of modes and less variance, but increases the exposure to Galactic foregrounds. The minimum field size was chosen to achieve 10% sample variance. The optimization leads to a range of $\ell = 50\text{--}100$, or angular scales

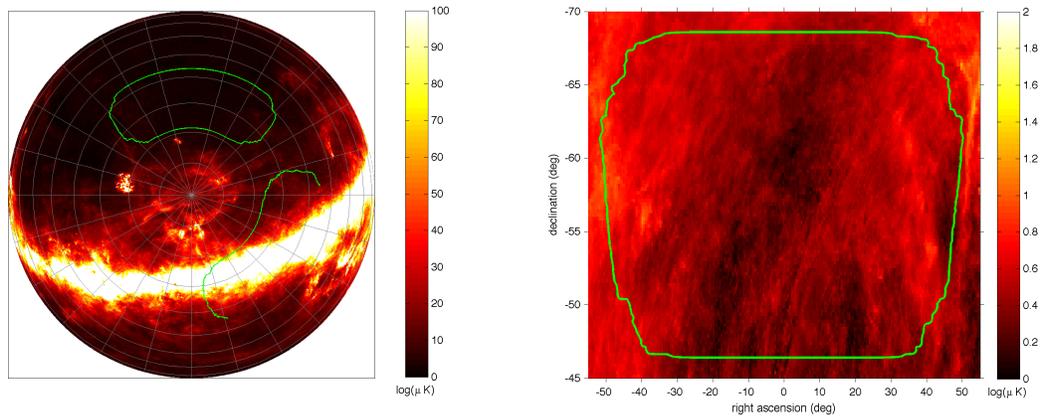


Figure 3.3: The BICEP2 observing fields and the Galactic dust emission at 150 GHz. *Left:* The FDS Model 8 prediction of the dust emission has been scaled to 150 GHz and projected on the sphere in celestial coordinates. The CMB field is outlined on the top and the bright Galactic field is indicated on the bottom right. *Right:* The dust emission is plotted in equatorial coordinates and overlay with the CMB field.

4° – 2° , where a primordial signal may be detected. This CMB field size gives sensitivity to multipoles as low as $\ell \sim 30$. The upper end of the multipole sensitivity ($\ell \sim 300$) is set by the 0.5° beam scale of the instrument.

A region called the “Southern Hole” offers a window to the CMB that is uniquely clean of foreground contamination. In this region the power spectrum of dust emission is lower by over a factor of 100 than the median over the entire sky. Based on the FDS¹ model 8 template for unpolarized dust emission, and assuming 5% fractional polarization, the B-mode signal from dust emission should be equivalent to $r \sim 0.02$ at 150 GHz [21][47]. The level of synchrotron signal should not contribute much more than this. In this field the total level of foreground confusion should be less than $r \sim 0.03$ over a range of $\ell = 50$ – 150 .

CMB observations are the primary science goal of BICEP2 and 60 hours out of every 72 sidereal-hour observing cycle are dedicated to the CMB field. The field is centered on ($RA = 0^{hr}$, $dec = -57.5^\circ$) and fills roughly half of the Southern Hole. It is constantly observable from the South Pole, allowing BICEP2 to concentrate its sensitivity on the small region and integrate deeply. The integration time is concentrated on a $\sim 800 \text{ deg}^2$ patch in the CMB field that comprises the cleanest 2% of the sky. The remaining observation time, six hours out of every 72, is spent observing a bright region of the Galactic plane centered at ($RA = 15^{hr} 42^m$, $dec \simeq -54^\circ$). Observation of this secondary field allows for a sensitive characterization of polarized Galactic dust emission and will provide a polarization angle calibration source for experiments such as Planck that require absolute

¹Finkbeiner-Davis-Schlegel

	phase	duration (sid-hr)	start LST (sid-hr)	region	start elevation (deg)
day 1	A	6	23	maintenance	90.0
	B	9	05	CMB upper	57.5
	C	9	14	CMB lower	55.0
day 2	D	6	23	Galactic	51.5, 53.25, 55.0, 56.75
	E	9	05	CMB lower	55.0
	F	9	14	CMB upper	57.5
day 3	G	6	23	CMB	55.0, 56.75, 58.5, 56.75
	H	9	05	CMB (lower/upper)	55.0, 57.5, 55.0, 57.5
	I	9	23	CMB (upper/lower)	57.5, 55.0, 57.5, 55.0

Table 3.1: Phases of the BICEP2 observing schedule. The schedule is based on a 72 sidereal-hour cycle and divided into nine phases. Each day has one 6-hour and two 9-hour phases, with at least 18 hours per day dedicated to the CMB. The pattern for visiting the upper and lower elevation portions of the CMB field are alternated each day. Each 72-hour block is observed at one of four fixed boresight angles. Phases D and G-I rotate through a pattern of four different starting elevations every block.

angle calibration $\leq 1^\circ$ [6][42].

3.5 observing schedule

BICEP2 observations occur on a 72 sidereal-hour cycle with the boresight fixed at one of the four angles listed above. The block is divided into nine phases: a maintenance phase (A) and eight observing phases (B-I). One of the observing phases (D) is dedicated to the bright Galactic region, while the other seven are spent observing the main CMB field. The maintenance phase is used to cycle the fridge, replenish the liquid helium, and several other tasks explained in § 3.6. Six of the CMB phases (B, C, E, F, H, I) last nine hours and include 10 scan sets, the remaining CMB phase (G), galaxy phase, and maintenance phase last six hours and include seven scan sets. The pattern of coverage of the observing field is shown in figure 3.5. Each day includes a 6-hour phase and two 9-hour phases. Two phases (C, E) always cover the lower half of the CMB field, two (B, F) always cover the upper half of the elevation, and one (G) covers the middle. The remaining two CMB phases (H, I) alternate between upper and lower coverage out of phase with one another. The order in which the upper and lower portions of the field are visited alternates every day, which allows for a jackknife test for any nonstationary azimuth-fixed contamination. At the conclusion of the 72-hour cycle the boresight angle is rotated 45° and the cycle is repeated. After all four angles have been visited the beginning elevation offsets are changed for some of the phases (D, F-I) as shown in table 3.5. There are four elevation offset patterns in total that are cycled through.

Calibration measurements are performed at the beginning and end of the observing phases and

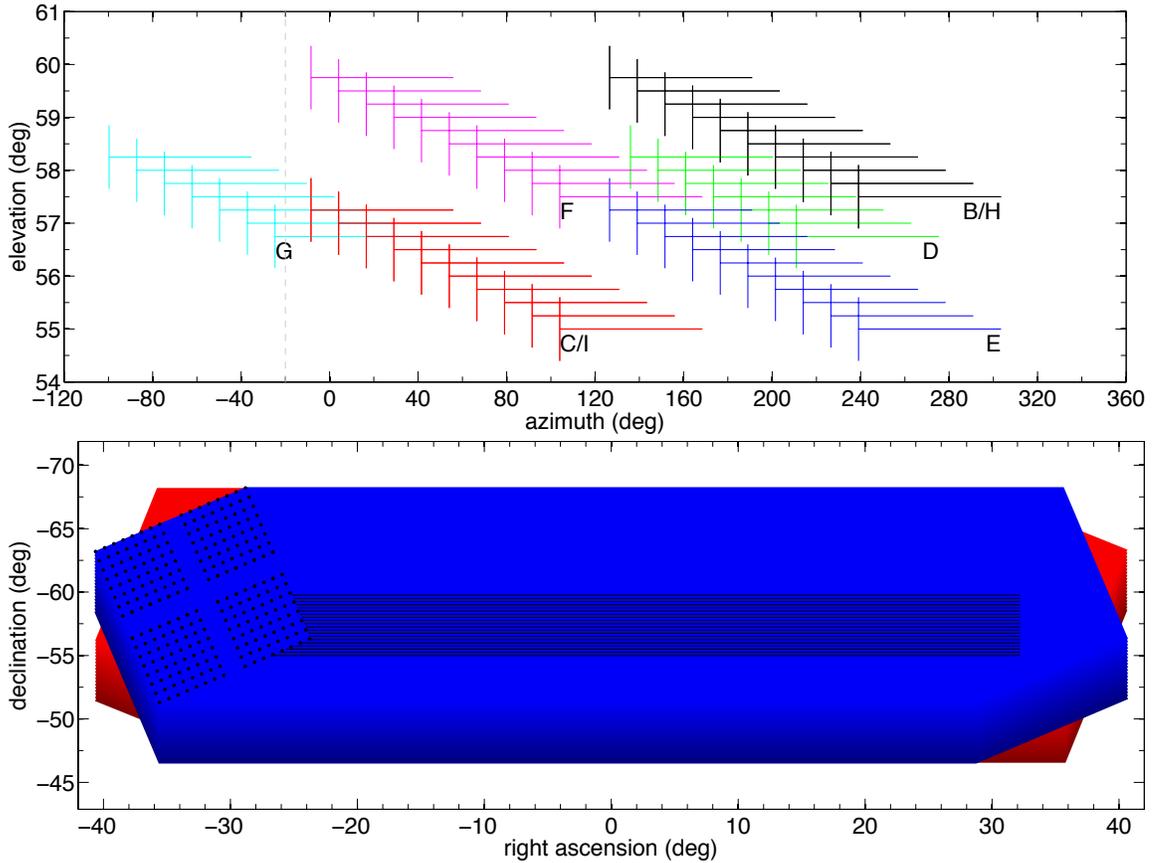


Figure 3.4: Scan pattern of the observing fields in rectangular and equatorial coordinates. *Top:* The scan pattern for each of the eight observing phases is shown in rectangular coordinates for one “start elevation” configuration listed in table 3.5. The phase letter indicates the beginning elevation of the scan set, which initiates the 7 – 10 scan set phase. Each scan set is observed at constant elevation over a fixed range of azimuth and then stepped up in elevation for the next scan set. *Bottom:* The scan pattern is now shown in equatorial coordinates for two non-overlapping boresight angles. This figure illustrates the resulting detector and polarization coverage of the field. The beams are shown as black dots, and the elevation steps for each scan set are shown as black horizontal lines. This plot does not include the increase in coverage in right ascension caused by the drifting of the field during the scan set.

scan sets. The calibration measurements, which are explained in § 3.7, consist of a sky dip, full and partial load curves, and elevation nods. Before the 9-hour CMB phases a skydip and full load curve are executed and then another full load curve is acquired following the scan sets. Each scan set is bracketed by partial load curves and elevation nods.

At the beginning of each observing phase a routine is run to configure the system for data acquisition. The pointing offsets for the mount are zeroed. The heat switches connected to the pumps are turned on and the focal plane temperature control servo is initiated. The sync box and MCE are configured with the observing settings, including the date type and low-pass filter, flux jump correction settings, frame rate and decimation factor. The SQUIDs are configured with a fixed tuning that is updated only after changes in the magnetic environment. The TESs are driven normal so that they can be operated on the superconducting transition and the biases are set to the standard observation settings. The low-pass filter and down sample factor used by GCP for detector data acquisition are set. After a pause for thermal settling the system is ready for data acquisition.

3.6 maintenance

A 6-hour phase is dedicated to maintenance of the telescope. During this time period the winter-over scientist performs routine tasks including a fridge cycle, liquid helium transfer and star pointing. To increase observation efficiency the maintenance tasks are all completed at the same time, though they could be staggered to reduce liquid helium consumption. The need for a maintenance phase every 72 hours is dictated by the hold time of the fridge.

In addition to the cryogenic and star pointing tasks described below, the following maintenance tasks are performed. The SQUIDs are tuned (§ 2.9.5.11) following the fridge cycle when temperatures have equilibrated. A single tuning is used to configure the telescope as long as possible, so the new tunings are generally not implemented. Accumulated snow is cleaned from the ground shield, boot, forebaffle and window. The membrane pressure is carefully monitored to ensure that it is taut. In the past, data have been contaminated by a common-mode signal introduced by motion of the membrane when it is not properly pressurized.

3.6.1 cryogenics

Under the loading conditions typical of standard observations the UC stage is able to cool the focal plane to the ~ 280 mK operational temperature for ~ 80 hours. The UC (^3He) still is able to reach

~ 250 mK, but the focal plane is operated at ~ 280 mK. The IC ${}^4\text{He}$ still reaches ~ 1 K for 36 hours until it runs out and warms to ~ 2.6 K. The focal plane is buffered from temperature changes of the IC stills by its temperature control servo. A double fridge cycle is performed in order to ensure reliable cycling of the fridge. The fridge cycle lasts approximately 4 hours. During the cycle the cryostat is pointed at zenith to attain the best cycle performance.

The cryostat can hold 100 L of liquid helium and uses ~ 22 L per day during normal observations. The liquid helium is refilled during the fridge cycle from a 250 L storage dewar. The storage dewars allow for two transfers and are refilled from the $\sim 10,000$ L station supply. With adequate management and efficient transfers the station supply lasts the nine months when the station is closed, before reserves can be flown in at station opening. The cryostat could be filled every ~ 110 hours for an increase in transfer efficiency, but is coordinated with the fridge cycle for improved observation efficiency and convenience for the winter-over.

3.6.2 star pointing

A pointing model is used to determine the telescope or boresight pointing from the raw mount encoder values. The model includes 10 dynamic parameters: the magnitude and direction of the tilt of the azimuth axis, the magnitude of the elevation axis tilt, the three encoder zero positions, two telescope flexure terms, and the magnitude and direction of the collimation error of the star pointing camera. To constrain the 10-parameter model the boresight and optical telescope should be pointed at 20 stars or more.

A star pointing schedule is run at each maintenance phase if the weather permits it, and before and after every leveling of the mount. The star pointing is a semi-automated schedule that requires the user to point the telescope at as many stars as possible. Stars are brought into view of the optical camera by the schedule, the user selects the centroid, and the next star is found. The elevation and azimuth offsets are recorded so they can be fit to the pointing model. The procedure is repeated for three deck angles: -45° , 45° and 135° . The procedure typically requires 1.5 hours in the summer when up to 12 stars can be found. In the winter up to 24 stars are found, requiring around three hours. The star pointing is run near the end of the fridge cycle to prevent the elevation tilt from interfering with the cycle efficiency.

In conjunction with the star pointing, the tilt of the azimuth axis is measured using orthogonal tilt meters mounted on the azimuth stage. A drift of $\sim 1'$ /mo has been observed and attributed to settling of the building. The telescope mount is leveled during the maintenance phase whenever the

tilt reaches $1'$. Star pointing schedules are run before and after every mount leveling to maintain accuracy of the pointing model.

3.7 calibration measurements

Calibration measurements bracket every scan set and the 9-hour CMB phases, as described in § 3.5. They consist of elevation nods, sky dips, and full and partial load curves. The elevation nods are used for relative gain calibration of the detectors explained below. A sky dip is executed before each 9-hour CMB phase to provide a profile of the atmospheric conditions. During a sky dip the telescope steps in elevation from 55° to 90° in 5° steps, with a pause of 20 s at each elevation step for a total duration of ~ 6 min. The sky dips are not currently processed during data quality checks or used in analysis.

Load curves acquire the I-V response of the TES detectors to help monitor the detector loading and gain, and linearity of the closed-loop SQUID amplifier. A full load curve ramps the the detector bias from high in the normal region of the detectors down to zero applied bias. This acquires the normal, superconducting transition and superconducting branches of the TES I-V response. Since the bias has been ramped to zero and the detectors have latched superconducting, the biasing steps explained in § 3.5 must be repeated. The resulting thermal transient requires several minutes to recover from, which is too long of a wait between scan sets. Instead, a partial load curve is acquired to trace out the normal region and a portion of the superconducting transition. The bias is ramped from a high (normal) value down to the bias applied to each column. Thus, the detectors remain in the superconducting transition and only a small thermal transient from the bias ramp must be recovered from. Acquisition of a partial load curve takes ~ 10 s and the full load curve takes ~ 30 s. Currently the full load curves are not processed or used in analysis.

3.7.1 relative gains

Elevation nods are executed before and after each scan set to determine the relative gains of the detectors for each scan set. The response of the detectors is fit to a model for the change in loading with the change in line-of-sight air mass, $T_{atm} \propto \csc(el)$. This gives a gain for each detector in analog-to-digital units (ADUs) per air mass, which is then normalized across the array. This gain calibration step serves to normalize the detector response in a pair. Detector pairs are differenced to produce a polarization measurement and are summed for a temperature measurement. Gain

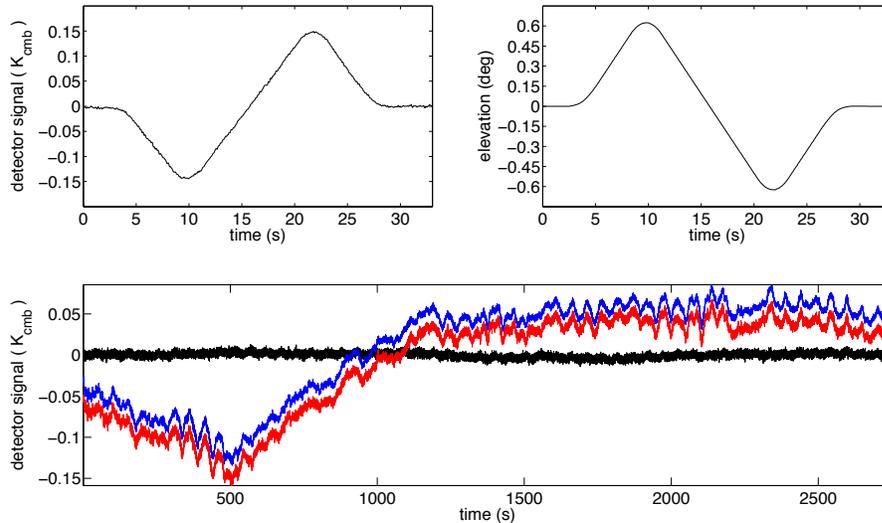


Figure 3.5: Elevation nods and individual and pair-differenced timestreams calibrated with the relative gain. *Top:* The detector response (*upper left*) to the change in loading during a $\pm 0.6^\circ$ leading (up-down-return) elevation nod (*upper right*) is shown. The signal and elevation offsets have been removed from these plots. The detector response goes down at higher elevations due to the lower line-of-sight atmospheric loading. *Bottom:* The timestreams of two detectors (red and blue, offset for clarity) in a polarization pair have been normalized with the relative gains derived from their elevation nods. The difference of the gain-calibrated timestreams is shown in black. The response of the detectors has been calibrated to antenna temperature using an absolute calibration derived from BICEP2 maps cross-correlated with WMAP maps.

calibration also normalizes the detectors to one another so that their measurements of temperature, Q, and U may be coadded during map making.

The elevation is scanned $\pm 0.6^\circ$ from the value used in the current fixed-elevation scan set. The typical response of the detectors is ± 150 mK during good weather, as shown in figure 3.7.1. The nod is executed as rapidly as possible without introducing significant motion-induced thermal contamination in the timestreams. This reduces the reliance on measurement of the transfer functions of the detectors to very low frequencies. Thermal stability of the BICEP2 focal plane allows the elevation nods to be acquired in 27 s, compared to the 51 s elevation nods used in BICEP1. To reduce a systematic bias from the motion-induced thermal response of the detectors, the pattern of the leading elevation nod (up-down-return) is reversed for the trailing nod (down-up-return) and the gains are averaged.

The relative gains within a polarization pair are corrected so that their timestreams may be differenced to produce polarization measurements. This rejects both atmospheric signal and CMB temperature. It also significantly reduces common-mode signals such as atmospheric fluctuations and scan-synchronous fluctuations. Assuming that the atmospheric emission is unpolarized this

procedure ideally results in a measurement of only the CMB polarization. However, because the in-band spectral shape of the atmospheric emission differs from the CMB temperature derivative, this method is sensitive to mismatches in the passbands of the paired detectors (see § 5.6.3).

Each detector timestream is normalized by the ratio of the median gain across the well-behaved detectors in the array to its own gain. This preserves the timestreams in units of ADUs. An absolute calibration to refer ADUs to antenna temperature (μK_{cmb}) is derived from cross-correlations between the coadded maps and WMAP maps.

3.8 initial data reduction

Standard observation data are recorded at 5 Hz for a total of 6.2 GB per day. The raw data are transferred daily from the South Pole to servers in the United States by satellite. Once per week an initial data reduction step is performed in order to check the data quality, as explained in the next section. The resulting time ordered data (TOD) files also serve as the low-level reduced-data product used in the analysis pipeline. The timestreams are analyzed and saved for each scan set, which is also referred to as a tag. The naming convention used for tags is *YYYYMMDDPSS.dkXXX*, where *YYYYMMDD* represents the year, month and day of acquisition, *P* is the phase, *SS* is the scan set number and *XXX* is the boresight (*dk*) angle. The analysis pipeline, including the initial reduction step, is implemented in MATLAB².

The first step in the initial data reduction is to load in the raw data saved in the arc files. Then the transfer function is computed from the GCP and MCE filter settings and measured detector transfer functions. This yields a finite impulse response (FIR) deconvolution kernel that removes the response of the GCP and MCE filters, and detectors. The time-domain deconvolution kernel is convolved with the detector timestreams and invalid regions at the beginning and end are replaced with NaNs. The next steps involve repairing regions of skipped and glitchy data. A de-skipping step identifies samples in the antenna layer that were not properly recorded by the BLASTbus PCI card and interpolates the data from adjoining samples. A de-glitching step identifies samples where a spike or step occurs in the detector timestreams. The former may occur due to a cosmic ray hit and the latter is usually caused by an uncorrected jump of the feedback flux at the edges of the DAC output range. Data around these time periods are replaced with NaNs to account for the ringing of the FIR deconvolution kernel. Next the elevation nods and partial load curves that bracket the scan set are reduced and recorded. The detector timestreams are normalized using the relative gains and

²MATrix LABoratory by MathWorks: www.mathworks.com

the calibrated TOD are saved to disk.

The instantaneous telescope pointing is determined from the mount encoder values recorded with the detector timestreams and the dynamic pointing model. The boresight pointing, along with the radio pointing that describes each detector beam relative to the boresight, describes the location of the detector beams on the sky. Analysis of the star pointing schedules that are acquired every three days give the eight (currently the flexure terms are not used) parameters used in the inverse³ pointing model along with the encoder values. The parameters used to analyze a particular scan set are interpolated from the star pointing schedules that bracket it. The star pointing schedules are analyzed during the weekly initial data reduction to ensure an up-to-date pointing model is applied and saved with the TOD.

3.9 data quality checks

During the weekly initial data reduction a suite of plots are made to monitor the data quality. The duty of generating and examining the plots is rotated among the members of the collaboration. Tags that show strange features in the data or particularly poor weather are noted to verify that they are properly selected by the cut infrastructure in the analysis pipeline. Four types of plots are generated and displayed in an HTML browser: a browser, daily examiner, cryo examiner and the tuning plots.

The scan set browser shows a variety of diagnostic plots that are organized by scan set, phase and 72-hour observing block. The plots that are most commonly viewed include a plot of the pair-sum and pair-difference timestreams during the azimuth scans and elevation nods, maps of the pair-sum and pair-difference detector covariance matrices and maps of T, Q, and U formed from the phase. A plot of the calibration quantities over the phase shows the gains from the elevation nods, the fluctuations of the light and dark channels and the changes in the Joule power and fractional resistance of the detectors. These latter two quantities are useful for monitoring the detectors for signs of saturation during periods of poor weather.

The daily examiner displays housekeeping data, diagnostics for GCP data acquisition, and some motion control data. The housekeeping data include plots of the readout of the NTD thermistors and the heaters and are useful for identifying unusual or even systematic thermal events. A plot of the fridge cycle is monitored to verify proper cycling. The cryo examiner plots display the volume, temperature and pressure of the three cryo-cooled storage dewars that contain the winter's liquid

³The inverse pointing model takes the encoder values to ideal topocentric coordinates, while the forward pointing model does the opposite.

helium supply.

The results of the SQUID tunings that are performed during each maintenance phase are closely monitored. The servoed and open-loop SQUID $V - \Phi$ response curves for the three stages are plotted along with the lock points selected by the tuning algorithm. The new tuning is plotted against the current fixed tuning used for the observation configuration. These plots are used to decide if the fixed tuning is still appropriate or if the tuning should be updated. A change in the magnetic environment of the SQUIDs — through flux trapping, for example — will cause phase shifts in the $V - \Phi$ curves along the abscissa. In this case the DC flux biases applied to the SSAs or SQ2s no longer correspond to the desired lock point and may lead to reduced gain and wrapping of the composed response curves.

3.10 observation efficiency

BICEP2 enjoys a modest increase in observation efficiency over BICEP1. One contribution comes from the increase in the hold time of the fridge, which needs to be cycled every three days instead of two. Given a 6-hour maintenance phase for both experiments, this means up to 92% of the time can be used for observations in BICEP2 as opposed to 88% for BICEP1. In addition, optimization of the scan profile has decreased the portion of the scan cut from analysis from 26% to 22%. When such things as scan calibrations and mount movement are taken into account, BICEP2 can ideally spend $\sim 79\%$ of its time on source. When the summer calibration period is neglected, this ideal was nearly being met in 2010. After dedicated observing began in mid-February, nearly 75% of the calendar year was spent on source [45]. Accounting for the fraction of the scan cut in analysis, the ideal observing efficiency is around 61% for BICEP2, while only 54% for BICEP1 [62].

Chapter 4

Detector and Readout Characterization

4.1 overview

In order to achieve the best instrument sensitivity and mapping speed the detectors and readout have been optimized for the best per-detector sensitivity and highest possible yield. This must be done while maintaining low systematic artifacts, such as crosstalk, that could introduce spurious polarization. The detectors and readout system must be well characterized in order to optimally configure them.

The bias of the detectors was optimized for the highest sensitivity, and the multiplexer was optimized for the highest readout rate and yield. The BICEP2 detectors show a significant high-frequency excess noise plateau, which causes degradation of the noise in the science band through aliasing. This calls for the highest multiplexer rate possible. However, certain types of multiplexer crosstalk are expected to increase with the readout rate. The need for decreased aliased noise must be balanced with the need for low crosstalk and multiplexer artifacts.

Conservative choices were made for the detector biases and readout rate when the instrument was deployed. These choices were made to ensure stability of the responsivity of the detectors and prevent crosstalk that could produce systematic contamination. Following careful characterization of the crosstalk the multiplexer rate was increased. After careful study of the stability of the responsivity of the detectors the biases were optimized. A wiring error was discovered part way through the first observing season that produced high levels of crosstalk. The SQUID biases were changed in order to mitigate that effect, which resulted in a large increase in detector yield and mapping speed.

The detector noise has been studied to understand the components that contribute to the overall

Table 4.1: Designed and measured detector performance

parameter	G_c	T_c	R_{tes}	R_n	β	η_{opt}
	pW/K	mK	$m\Omega$	$m\Omega$		
target	80	450	30	60	2.5	0.30
median measured	100	520	61	74	1.9	0.45

noise and sensitivity. Those studies required measurement of many of the TES parameters, which were used as inputs to the noise model. To understand how noise power on the bolometer, or the noise-equivalent power (NEP), is referred to the antenna to give the noise-equivalent temperature (NET), the optical efficiency of the detectors was characterized. All NETs in this chapter are given in units of $\mu K\sqrt{s}$, where the temperature is that of a CMB blackbody, $K \equiv K_{\text{CMB}}$.

4.1.1 device parameters

The device parameters of a TES are measured from its current-voltage characteristic or load curve. A load curve is acquired by recording the output current of the TES while varying its bias so that it transitions from the normal to superconducting state. First, the detector must be put into the normal state. This is done by either biasing it with enough current for it to self-heat above the critical temperature or by raising the base temperature of the array above the critical temperature, while maintaining a sufficient bias to prevent it from latching to the superconducting state. The bias is then slowly ramped to zero, tracing out a load curve like the one shown in figure 4.1.

The load curve is divided into the normal branch, superconducting transition and superconducting branch. The slope of the normal branch gives the normal-state resistance (R_n) of the detector, while the slope of the superconducting branch gives the input coil to feedback coil current gain (M_{in-fb}) of the first stage in FLL. This value is needed to refer the applied feedback current to the current flowing through the input coil and TES. In the superconducting state all of the detector bias current flows through the TES and input coil branch of the bias circuit, assuming that the parasitic resistance of that branch is negligible compared to the shunt resistance (R_{sh}). A schematic of the TES bias circuit and first-stage of the SQUID readout is shown in figure 2.1. The shunt resistance was measured to be $3.0 m\Omega$ in a separate 4-wire measurement at ~ 300 mK. With this value, along with a measurement of M_{in-fb} for each circuit¹, the normal resistance (R_n) can be measured from the slope of normal branch.

To properly deduce the output current and the applied detector bias current shown on the

¹in practice these were similar enough to assume a constant value of 15.5 was used for each circuit

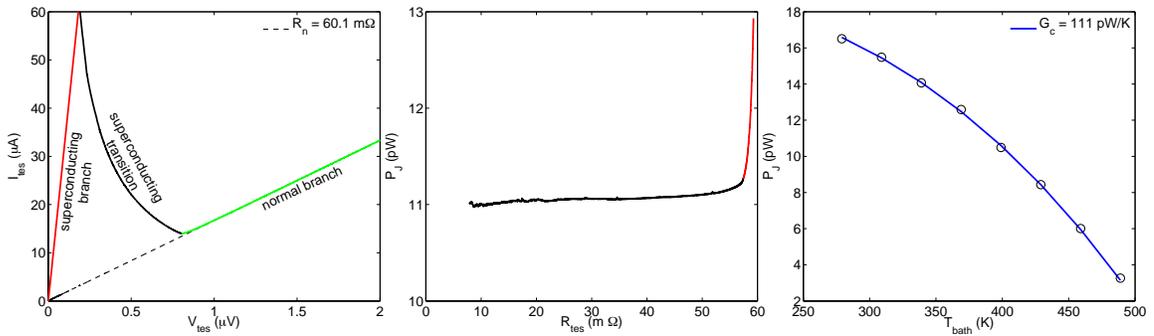


Figure 4.1: TES load curve analysis. *Left*: An I-V curve for a typical TES showing the normal branch, superconducting transition and superconducting branches. The normal resistance is fit from the slope of the normal branch. The input to feedback current gain is derived from the slope of the superconducting branch. *Middle*: The Joule power in the superconducting transition (black) is fairly constant due to strong ETF. *Right*: The parameters G_c , β and T_c can be determined from a fit of equation 2.1 to the Joule power from load curves acquired a range of base temperature.

ordinates of figure 4.1, the cryostat line resistances and internal resistances and voltage outputs of the MCE must be accurately measured for the feedback and detector bias lines. Typical values are shown in table 2.3. In addition to R_n the load curves give a measurement of I_{tes} and R_{tes} through the superconducting transition. Due to strong electrothermal feedback, the power dissipated by the detector $P_J = I_{tes}^2 R_{tes}$ remains constant through the transition, as shown in the figure. The average of this flat region through the transition gives the Joule power of the detector, which changes in direct response to the optical power. In order to determine the additional detector parameters (G_c , T_c , and β) load curves must be acquired at a range of base temperatures. The power in the transition at each base temperature is fit to equation 2.1 as shown in figure 4.1.

The detector parameters were measured at the South Pole, and during engineering runs. These measurements provided feedback to the fabrication group that was aiming for the targets shown in table 4.1. The parameters are used as input in the detector noise model in § 4.1.6.1, which can be compared to the measured noise to assess the performance of the detectors. In addition, partial load curves are acquired before and after every scan set to dynamically measure P_J , R_{tes} and I_{tes} during observations. The Joule power and operational resistance change in response to optical loading, which can be used to indicate periods of poor weather.

The measured detector parameters characteristic of the TESs during standard observations are summarized in table 4.1 along with the targeted value. The detectors were designed for an optical load of ~ 3.6 pW and a safety factor of ~ 2.5 for $P_{sat} = 9$ pW. For this saturation power $G_c = 80$ pW/K for $T_c = 450$ mK and $T_{bath} = 250$ mK. The actual measured conductance was in the range of $G_c \simeq 80$ –155 pW/K. Tiles 3 and 4 were found to have a mean $G_c \sim 80$ pW/K, in agreement

with the design target. Tiles 1 and 2 have higher and more dispersive conductances ranging from 115–155 pW/K. These lower conductance tiles were added prior to deployment to replace an earlier pair of tiles and came from a different fabrication run. The median measured value of 100 pW/K is presented in the table and used in the noise model below. The thermal conductance exponent was measured to be $\beta \simeq 1.9$, lower than the anticipated 2.5.

The transition temperatures were found to be better matched, ranging from 500–530 mK, but still above the targeted value of 450 mK. The median value of 520 mK is presented in the table and is used in the noise model. The operational resistance was designed for $R_{tes} = 30 \text{ m}\Omega$, a factor of 10 higher than R_{sh} to voltage bias the TES. The detectors were anticipated to be operated at the midpoint of the transition, so $R_n = 60 \text{ m}\Omega$. Based on the partial load curves acquired during observations the normal resistance is slightly higher than the target, with a mean value $R_n = 74 \text{ m}\Omega$. Due to the impact of aliased excess noise, discussed below, the detectors are operated high in the transition at a mean resistance of $R_{tes} = 61 \text{ m}\Omega$, corresponding to a fractional resistance $R_{\%} = \frac{R_{tes}}{R_{norm}} = 82\%$. A mean operation current of $I_{tes} = 13.5 \text{ }\mu\text{A}$ was measured from load curves.

Additional detector parameters — the loop gain (\mathcal{L}_I), time constants τ_{\pm} , logarithmic temperature sensitivity at constant current ($\alpha_I \equiv \frac{\partial \log R}{\partial \log T}$), logarithmic current sensitivity at constant temperature ($\beta_I \equiv \frac{\partial \log R}{\partial \log I}$), and heat capacity (C) — can be determined from complex impedance measurements. The complex impedance can be extracted from the response of the detectors to a small square-wave signal applied to the TES bias lines. Measurements of the response of the BICEP2 detectors to bias square waves was made prior to deployment. Those measurements were only used to determine τ_{\pm} , as explained below. A full complex impedance analysis was not performed. The temperature sensitivity used in the analysis below has been approximated from the steepness of resistance versus temperature profiles. They give $\alpha_I \sim 100$, which is consistent with the design value. The heat capacity assumed in the analysis below is the design $C = 1.2 \text{ pJ/K}$. The logarithmic current sensitivity is assumed to be $\beta_I = 0$.

4.1.2 time constants and stability

The effective thermal time constant of the detectors must provide sufficient bandwidth so that signals within the science-band are not rolled-off. The detectors must also be designed for and operated so that they are stable and have no thermal oscillations over a wide range of bias and loading conditions. The first requirement means $\frac{1}{2\pi\tau_-} \gg 1 \text{ Hz}$, so that the thermal response is much faster than the high-end of the science band. The second requirement means $\tau_- \geq \tau_+$, which are the conditions for

an over or critically damped TES. All voltage-biased TESs are stable unless they are underdamped (§ 2.1).

For strong ETF, well-separated τ_{\pm} , $\beta_I = 0$, and $R_{sh} \ll R_{tes}$ the time constant τ_- can be expressed as the effective thermal time constant $\tau_{eff} = \frac{C/G}{1+\mathcal{L}_i}$. This is just the natural thermal time constant enhanced by the loop gain of ETF, which increases the speed of the detectors' response to thermal signals on the bolometer. Given the measured parameters above, the loop gain for the median detector is $\mathcal{L}_i = \frac{P_I \alpha_I}{G_c T_c} \simeq 21$. For the median conductance $G_c \simeq 100$ pW/K and designed heat capacity $C = 1.2$ pJ/K the effective thermal time constant is $\tau_{eff} \simeq 0.5$ ms with bandwidth $f_{3dB} \simeq 300$ Hz. This thermal bandwidth is far larger than the 0.1–1 Hz science band, so there should be no roll-off of signals. This estimate is consistent with the results for τ_- measured from the response of the detectors to small steps in the applied bias. The time constants have been measured by introducing a square-wave signal to the detector bias lines. The timestreams show a fast electrical response to the change in applied bias and a slow thermal response to the heat introduced on the island. The effective thermal time constants have been estimated from fits to the stepped response. The time constants are < 1 ms even low in the transition. The faster electrical time constants are more difficult to measure, but are consistent with $\tau_{elec} = \frac{L_{tot}}{R_{tes}} \simeq 26$ μ s for $L_{tot} = L_{nyq} + L_{in} = 1.6$ μ H and $R_{tes} = 61$ m Ω .

The large separation of the two time constants ensures that the detectors do not oscillate for our choice of operating biases. Although the detectors have been slowed down by the added heat capacity of the gold, they are still sufficiently fast for the scan speed and their transfer functions are flat within the science band.

4.1.3 loading

The TESs are only able to function when the incident optical power is lower than what can be removed from the island by the thermal conductance. Large momentary signals created by the amplified sources during beam mapping and increased loading during periods of bad weather can cause the detectors to saturate. For a voltage-biased TES with a narrow transition region, the saturation power can be written as $P_{sat} = (1 - R\%) P_{bath}$ [30]. The fractional resistance tracks changes in optical loading in accordance with the power balance in equation 2.2. As the optical load (P_{opt}) increases the Joule power ($P_j = V_{tes}^2/R_{tes}$) decreases and the operational resistance of the voltage-biased TES increases. When P_{opt} reaches P_{bath} and the Joule power reaches zero, the detector becomes unresponsive to optical signals. Detectors biased in the midpoint of the transition

will have a saturation margin of half the bath power.

The titanium TESs were designed for a total saturation power of $P_{bath} = 9$ pW for an assumed load of 3.6 pW and safety factor of 2.5. The actual total saturation powers were estimated from the measured detector parameters using equation 2.1. Tiles 1 and 2, with higher G_c , have a median $P_{bath} \simeq 18.5$ pW, while tiles 3 and 4 have a median $P_{bath} \simeq 12$ pW. Load curves were acquired at the $\sim 57^\circ$ observing elevation during the summer to determine P_J . These gave a median $P_J \simeq 14$ pW for tiles 1 and 2, and a median $P_J \simeq 7$ pW for tiles 3 and 4. Using the estimated P_{bath} , measured P_J , and the steady state power balance $P_{bath} = P_J + P_{opt}$, the total loading during the measurement was estimated at $P_{opt} \simeq 5$ pW. This include the atmospheric loading and internal loading from the receiver. The loading in the summer can be as high as 6 pW during periods of weather suitable for observing. The atmospheric loading is lower during winter observing conditions, with a total optical load of around 4 pW. The loading during winter is comparable to the 3.6 pW load the detectors were designed for. The detectors have higher G_c and T_c than the target, so they have a higher than expected loading margin or safety factor.

The fractional resistance and Joule power are measured for every scan set by the bracketing partial load curves. As the detectors near $\sim 90\%$ they begin to exhibit reduced response to elevation nods. This generally occurs during periods of poor weather with higher atmospheric loading. The fractional resistance can be used as an additional indication of poor weather and as a data cut.

4.1.4 optical efficiency

High end-to-end optical efficiency is important to couple the detectors to as much of the CMB signal as possible to give the best sensitivity. The optical efficiency (η_{opt}) effectively refers the noise power at the bolometer to antenna temperature. The optical efficiency includes reductions in efficiency due to the optics, such as reflections and in-band absorption, as well as the detector passband and the inherent efficiency of the antennas. The optical efficiency was measured by acquiring TES load curves with the detectors subjected to two extreme antenna temperatures. A cone with a diameter slightly larger than cryostat window was lined with microwave absorber. The cone was placed over the window while at ambient temperature (300 K) and load curves were acquired for the detectors. The cone was then saturated with liquid nitrogen (77 K), placed over the window and load curves were acquired again. The cone was found to stay isothermal at 77 K for a period of ~ 1 minute, so the load curves were acquired within this time frame. The power in the transition region was found for both sets of load curves for each detector. The optical efficiencies were then determined

Table 4.2: Measured and predicted crosstalk mechanisms

type	level	linear	affecting
inductive nearest-neighbor	0.3% (-25 dB)	yes	adjacent rows in column
inductive SQ2 to next-nearest SQ1	0.09% (-31 dB)	yes	entire column
inductive SQ1 to feedback	0.1% (-30 dB)	yes	entire column
settling	0.03% (-36 dB)	yes	next row in column
common ADC, induced steps	~ 1 ADU	no	same row other columns
common ADC, unservoed TES	0.01% (-20 dB)	no	same row other columns
threshold	10–25 ADU	no	same column other rows
common TES bias	< -37 dB	yes	same column other rows
common TES substrate	< -72 dB	yes	same array tile
common SQ1 bias	0.03% (-35 dB)	yes	same row other columns
ghost beams	0.01% (-40 dB)	yes	single channel

from the ratio of the change in power to the change in temperature, with an average of $\frac{\Delta P_j}{\Delta T} = 0.23$ pW/K_{RJ} \simeq 0.13 pW/K_{CMB}. Given the beam size and assumed throughput of the telescope, this is equivalent to $\eta_{opt} \simeq 0.45$ end-to-end optical efficiency [46]. The detectors were designed for $\eta_{opt} = 0.30$ optical efficiency. This improvement means they will have increased sensitivity; the same bolometer power will correspond to a lower antenna temperature.

4.1.5 crosstalk

Crosstalk creates additional systematic effects that must be well understood to assess their impact on the polarization fidelity of the measurements. The use of a multiplexer in BICEP2 presents numerous potential sources of crosstalk that did not exist in the single-channel readout architecture used in BICEP1. The multiplexer crosstalk arises due to the use of common components for rows or columns of detectors in the multiplexer to reduce wiring count, and close proximity of magnetically sensitive components. All 32 detectors in a multiplexer column share a common pair of bias lines located on the Nyquist chips. They also share common first-stage (SQ1) flux feedback lines, summing coil and second- (SQ2) and third-stage (SSA) SQUIDs and ADC. Additionally, the 16 ADCs and output DACs share a common ground in the MCE. The 16 SQ1s in a multiplexer row share common biasing-address lines. In addition, the close packing and common substrate of the detectors creates a potential for thermal crosstalk of the detectors.

Crosstalk has been assessed through maps acquired by rastering bright sources such as the moon and the BNS (§ 5.2). The signal to noise in these maps is generally good enough to probe crosstalk to a level of around -40 dB. The fractional crosstalk levels are derived from the ratio of the integrated signal at the pickup channel to the integrated signal at the source channel. The fractional crosstalk is then converted to dB using the standard $10\log_{10}$, which is appropriate to measuring power levels

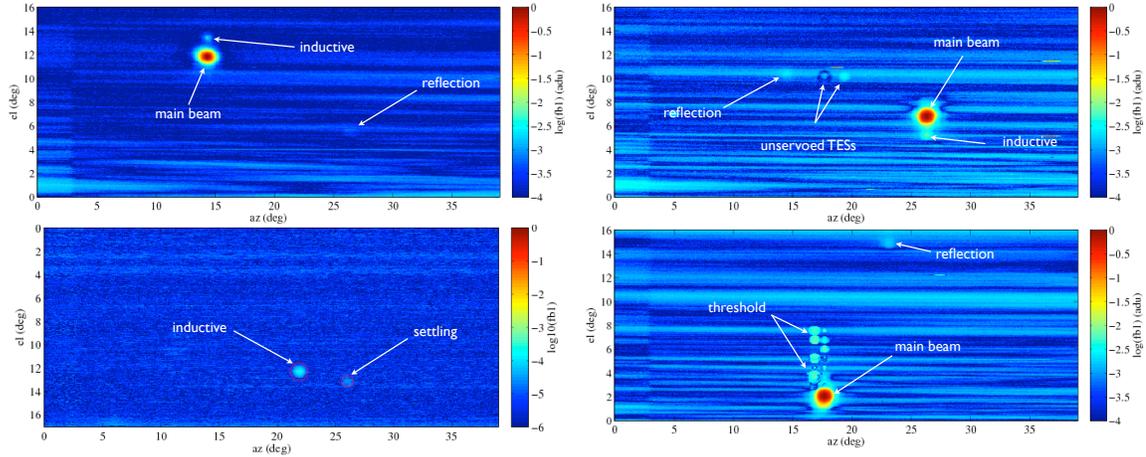


Figure 4.2: The various types of crosstalk observed in bright source rasters. This is a compilation of four beam maps of the moon taken during the austral summer with the far-field flat mirror. The typical noise floor in an optically coupled channel is -35 dB. The map in the bottom left is for a dark SQ1 and has a much lower noise floor (-55 dB), which is useful for measuring settling crosstalk.

of signal on the sky.

4.1.5.1 inductive

The close proximity of the magnetically sensitive SQ1s and their input coils on the multiplexer chips leads to inductive nearest-neighbor crosstalk. Changes in the current through the input coil of a detector results in flux at the SQ1 it's coupled with, but also a certain coupling of flux to the adjacent SQ1s. The gradiometric design of the SQ1s and input coils will not reject the large gradient produced by the nearby inductive coil at the SQ1. The crosstalk only shows up in the pickup channel when its SQ1 is biased on. The signal is clear in beam maps, where a faint neighboring beam appears in the maps as shown in figure 4.2. Only a single inductive crosstalk beam appears in the maps since one of the adjacent rows has a crosstalk beam that is colocated with the main beam. Inductive crosstalk is also apparent in detector load curves. The magnitude of the crosstalk has been measured from the ratio of the integrated power of the crosstalk response to the source response at 0.3%, or -25 dB. This level is consistent with the value² 0.25% reported by NIST-Boulder for an earlier version of the multiplexer [16]. They also report additional crosstalk from the next nearest-neighbor at a level of 0.025% (-36 dB). This is near or below the level of signal-to-noise in the beam maps and has not been detected.

Inductive crosstalk also occurs between the input coils of the SQ2 and nearest SQ1s, and is largest for the SQ1 located next to the SQ2. It was estimated by NIST to be 0.25% (-26 dB) between the

²The reference reports dB as $20\log_{10}$.

SQ2 and nearest SQ1 and 0.09% (-30.5 dB) for the next SQ1 for an older version of the multiplexer chip [16]. For this reason, the SQ1 nearest to the SQ2 is chosen as the dark SQ1 that is used to monitor common-mode signals. If a TES were coupled to the input coil of this nearest SQ1 (row 1) then its signal would be coupled to every other channel in the column through the common SQ2. The effect is predicted to be small enough for the next SQ1 (row 2) that its input can be safely coupled to a TES.

Coupling between the input and feedback coils and the common feedback line causes crosstalk between an on-state SQ1 and the other 32 off-state SQ1s in the column. Flux feedback is applied to the on-state SQ1 to null the input coil flux, but that same flux is coupled to all the other input coils. Screening currents in the superconducting input transformers then couple to the bias circuits of the TESs. Those circuits respond to the change in flux with a time constant governed by the total inductance and dynamic resistance of the TES. The residual signal that remains in that channels circuit when it is eventually visited and sampled depends on the coupling between the feedback and input coils k_{fb-in} . An additional crosstalk contribution arises from the power dissipated in the TES by the change in flux in its bias circuit. These two contributions have been mitigated by the use of a dummy SQ1 to significantly reduce the input-feedback coupling (§ 2.5.2). The introduction of the dummy SQ1 in the second-generation multiplexer chips reduced k_{fb-in} from 0.6 to 0.016. The level of crosstalk from the residual flux is worst for the next row visited after the channel originating the signal. For this channel, under standard observing conditions, the crosstalk should be 0.1% (-29.7 dB) at 15.5 kHz and only 0.07% (-31.7 dB) at a 25 kHz detector sampling rate. The crosstalk from the power dissipated in the TESs should be negligible, at only -36 dB for common-mode signals and -51 dB for signals originating in a single channel.

4.1.5.2 settling

There are three switching time constants of interest in the multiplexer. The multiplexer rate is limited by the time constants for switching the input and feedback flux to the SQ1s, τ_{in} and τ_{fb} . If the multiplexer is run any faster than this then changes in the input flux can't be compensated and measured. These have been measured to be $\sim 0.2 \mu\text{s}$ in BICEP2. The other two time constants are the rise and fall time for switching the SQ1s on and off. The rise time τ_{rise} is governed by the open-loop bandwidth of the amplifier chain. In BICEP2 it is believed to be dominated by the dynamic resistance of the SQ2, and stray cable and SSA input inductance in the SQ2-SSA circuit. The fall time τ_{fall} is longer since superconducting SQ1s lower the resistance and increase the L/R time

constant in the SQ1 bias circuit. While τ_{fb1} and τ_{in} set the upper limit for the rate of multiplexing, τ_{rise} and τ_{fall} affect the settling time crosstalk levels. For a fixed crosstalk budget, they too can limit the multiplexer rate.

As the multiplexer switches from row to row the amplifier settles from the previous value to the next. The amplifier switches between the different zero-error signal levels for the rows in the column. The zero-error signals are the ADC values at which the flux in the input coil of the SQ1 is nulled by the flux feedback. Changes in the input flux $\delta\Phi_{in}$ from the zero-error signal level can be transferred to the next row in the multiplexer sequence due to the finite settling time of the amplifier. The situation is further complicated by transients that occur when switching between rows, which can increase the slew distance. The switching transients are caused by the difference in the voltage levels and switching time constants of on-state and off-state SQ1s. For the choice of three-turn input coil SSAs, SQ2 biases and lock points, and cabling, $\tau_{fall} \simeq 0.4 \mu s$. During the initial setup a conservative row dwell time $\tau_{dwell} = 1.96 \mu s$ was chosen to allow the amplifier to settle $\tau_{settle} = 1.76 \mu s$ before being sampled for $0.2 \mu s$ by the ADC. This gives a safety factor of around 4.4 in the settling time before sampling. The frame rate, or detector readout rate, corresponding to τ_{dwell} for a 33 row multiplexer is 15,462 Hz, given by equation 2.8.

The upper limit for the open-loop fraction of settling time crosstalk was estimated at only -35 dB for a multiplexer operating at double this rate [16]. So, this was expected to be a very conservative initial choice for the readout rate. The median value measured in BICEP2 for 15.5 kHz^3 from the ratio of integrated crosstalk to source power was 0.025%, or -36 dB. When the multiplexer rate was increased to 25 kHz^4 the median value increased only marginally to 0.037%, or -34 dB. This is far lower than the contribution of inductive crosstalk and expected to make a negligible systematic contribution. These low levels of crosstalk are near the limit of the signal to noise in light channels. These values were measured from the maps made from the response of the dark SQ1, which have a noise floor around -60 dB.

4.1.5.3 common ADC ground

Two sources of crosstalk occur due to the use of a common ground for the 16 ADCs that sample the SQUID amplifiers. The ground is also shared with the DACs that output the SQ1 flux feedback in the readout cards. When the output for a channel nears the edge of the DAC range it is automatically brought into range by the multiplexer, which is implemented at the next visit of that channel's row

³ $\tau_{dwell} = 1.96 \mu s$ and $\tau_{settle} = 1.76 \mu s$

⁴ $\tau_{dwell} = 1.2 \mu s$ and $\tau_{settle} = 1.0 \mu s$

(§ 2.5.2). The large swing in the DAC signal is accompanied by a large current change that causes the common ground to shift. This means that the zero-error ADC values for channels in the same row shift and the flux feedback values calculated from the previous visit no longer zero the ADC input. Thus, at the next visit of that row an additional flux feedback offset is applied. The flux jump and step in the output only affect channels in the same row on the different columns. The resulting step size is usually ~ 1 ADU. An extreme example of this occurs when a channel loses lock and the DAC output fluctuates wildly. The impact on the data is discussed in the next section that covers a related effect.

The same issue occurs when a readout channel that is coupled to an optically active TES is turned off in the dead detector mask. This mask informs the MCE to turn off the FLL servo for that channel by setting the PID gains to zero. Due to the common SQ1 bias lines the SQ1 for an unservoed row is still turned on. A change in the TES signal and corresponding change in input flux are not nulled by feedback flux. The resulting current change at the ADC causes a shift in the common ground, creating a source of crosstalk in the same row of the different columns. The ground shift tracks the input coil current, but since there is no FLL to linearize the SQUID amplifier the crosstalk is nonlinear. The magnitude of the effect is determined in part by the open-loop gain of the amplifier chain. The crosstalk has been measured in beam maps to be as high as 0.01% (-20 dB). It is easily remedied by not turning off optically coupled TESs in the dead detector mask.

4.1.5.4 threshold

A source of crosstalk not previously reported in the literature was discovered in the BICEP2 beam maps, which is referred to as threshold crosstalk. The crosstalk manifests itself in beam maps, such as the one shown in figure 4.3, as flat-bottomed, disk-shaped wells in several channels within the same multiplexer column. The crosstalk has the opposite sign as the source signal and an amplitude that is somewhat fixed for a given channel around 10–25 ADU. Timestreams of the detector response during a fixed-elevation azimuth scan are also shown in the figure. The onset of crosstalk is a step response in the pickup channel to the source channel crossing a fixed feedback DAC output threshold. When a source channel crosses a positive DAC threshold the pickup channel steps down by 10–25 ADU. If the source channel nears the edge of the DAC output range and its output is corrected downward the pickup channel's response steps back up to the value prior to the crosstalk onset. Evidence of this effect have not been found in the dark SQ1s that have open input coil circuits. Examples of the same effect have since been located in the ACT data, which uses the same version of multiplexer

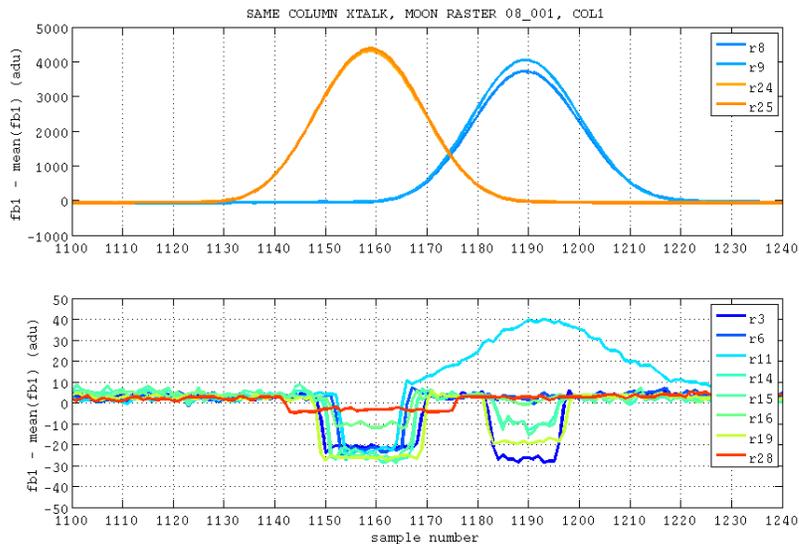


Figure 4.3: Timestreams of channels showing crosstalk when the source channels exceed a DAC output threshold. Top: Four potential crosstalk source channels show a large response when slewing across a bright, ground-fixed source. Bottom: Eight channels show a fixed offset crosstalk response when the source crosses some individual DAC threshold set by the pickup channel. The negative offset disappears when the source drops below the threshold.

chips used in BICEP2 [23][44].

This type of quantized crosstalk suggests a flux quantization effect in the SQ1s. Data has been taken to test this theory but has not yet been analyzed. The impact of this effect has not been well-studied, but it doesn't appear to be wide spread in the observation data. It should only occur when detector signals drift greatly in response to changes in loading and cross the feedback DAC thresholds. There is evidence of this in timestreams when the DAC output flux jumps back in range for a source channel. Several other channels in the same column have corresponding steps in their timestreams of the opposite sign. The effect is similar to the effect on the timestreams as the crosstalk caused by flux jumps due to the common ground of the ADCs. However, the opposite sign and 10–25 ADU magnitude of the step differentiate them from the other mechanism. In both cases, the main impact will be on the noise model, which is not designed to handle step functions in the data. It is not a constant crosstalk source like inductive crosstalk, so can be handled in analysis by locating and correcting the steps where possible. The most notable effect of threshold crosstalk is in the quality of calibration beams maps. The bright calibration sources cause the DAC threshold to be crossed frequently, so there are lots of threshold crosstalk features in these maps. This complicates the study of smaller stationary crosstalk sources such as settling and inductive crosstalk, since the

widespread threshold crosstalk overwhelms their crosstalk signatures.

4.1.5.5 detectors

Crosstalk can arise due to the common bias lines and common substrate of the detectors. As described in § 2.19, the 32 TESs are each shunted by $3.0\text{ m}\Omega$ resistors connected in series to a common pair of bias lines. The crosstalk from this biasing scheme is calculated at a level of -52 dB for crosstalk originating from a single channel and as much as -37 dB for common-mode signal. Generally the variations of the TES output current are small enough that this linear crosstalk effect is not problematic. However, if a TES begins to oscillate the pickup on the other channels can be substantial. This can be avoided by careful selection of the TES biases.

The effect of the common bath or substrate of the detectors can induce crosstalk from one detector to the nearest neighbors as the power dissipated in the island flows to the substrate and raises its temperature. This effect should be extremely small due to the high internal conductivity of the arrays and their heat sinking to the isothermal base plate. It is estimated for the standard BICEP2 observing conditions at < -72 dB.

4.1.5.6 other

Several other sources of crosstalk in various categories include crosstalk from wiring issues, common SQ1 bias lines, and ghost beams. Severe crosstalk between several multiplexer rows was discovered in beam maps taken during the initial set up of BICEP2. Single-channel maps showed multiple beams with amplitudes comparable to the expected main beam. The crosstalk corresponded to strange shapes in the $V - \Phi$ response curves of many channels in those rows, as shown in figure 4.8. The problem was traced to wiring shorts between the bias lines of several rows, which caused SQ1s to be inadvertently biased on and read out when there row was not being visited. This crosstalk was mitigated as explained in § 4.1.8.2.

Crosstalk is expected from the use of a single pair of bias lines to address all SQ1s in a given row. A signal at one SQ1 bias circuit changes the loading of the lines coupled to all other SQ1s in the row. It is estimated at a level of 0.03% (-35 dB) for signals originating from a single source. The effect is linear and can sum as high as 0.45% (-23.5 dB) for common-mode signals on the other 15 detectors in the row. There is no evidence yet for this effect in the BICEP2 data.

Ghost beams are a type of crosstalk that BICEP2 has in common with BICEP1 [12]. They appear in beam maps as the image of the source rotated 180° about the boresight of the telescope. Ghosting

Table 4.3: Noise budget for the BICEP2 detectors from the noise model. The design NEP was calculated from the design parameters in table 4.1. The median NEP was calculated from the median measured parameters in the same table. The photon noise was modeled in both cases for a 4 pW load. The Johnson noise includes the low-frequency contributions from both the shunt and TES. The aliased (detector and amplifier) noise and total noise are shown for three readout rates. At 15.5 kHz the total aliased noise exceeds the photon noise. Above 20 kHz the photon noise is expected to be larger than the total aliased noise.

	NEP design	NEP median
	aW/ $\sqrt{\text{Hz}}$	aW/ $\sqrt{\text{Hz}}$
photon	36	41
phonon	20	27
Johnson	2.6	4.0
aliased, total	9.6	46/36/22
total	42	68/62/55
readout rate (kHz)	15.5	15.5/20/30

is caused by reflections of the source image off the focal plane then off of the nearest filter and back to the opposite side of the focal plane. Due to the use of monochromatic antireflection coatings on the lenses and filters of BICEP2, the effect was expected to decrease from the -22 dB level measured in BICEP1 for the 150 GHz channels [12]. The level of crosstalk taken from the ratio of the integrated power of the crosstalk signal to the source in the beam maps is 0.01% (-40 dB) in BICEP2.

4.1.6 noise characterization

The detectors were designed for near-background-limited performance given the expected loading from the cryostat and atmosphere during standard observing conditions at the South Pole (§ 2.1). The expected contribution from each of the noise components is shown in table 4.3. The BICEP2 detectors were designed for a total NEP of 42 aW/ $\sqrt{\text{Hz}}$. For a designed optical responsivity of 30% and 25% fractional bandwidth this NEP corresponds to a per-detector NET of 330 $\mu\text{K}\sqrt{\text{s}}$. In this section, the analysis of the individual contributions to the total detector noise are presented for data taken with the early 2010 bias settings. This analysis informed the decision to increase the readout rate to reduce the significant contribution of the aliased noise to the total noise.

The noise contributions were derived from expected or measured detector parameters and the frequency-dependent noise model taken from the reference [30]. The components to the NEP predicted by the noise model are given in equations 4.3, 4.4, 4.2. The contributions considered in the total are shown in equation 4.1. The amplifier and aliased excess noise contributions, which must be measured from noise spectra, are included in the noise model.

A characteristic and well-measured detector was first analyzed to assess the agreement between the measured and modeled noise spectra. The median measured detector parameters from table 4.1

were then used as inputs to the noise model to estimate the median NEP and noise-equivalent current (NEI). The median modeled total noise was then compared to the measured distribution of the NEP and NEI.

4.1.6.1 noise model

An analysis of the contributions to the total noise of the BICEP2 detectors is presented. The noise components include photon, phonon (thermal fluctuation), Johnson, excess and amplifier noise. In order to determine the components of the total noise the detectors must be sampled above the standard readout rate. This allows one to assess the high-frequency noise contributions and prevent obscuration from aliasing. Special data taking modes of the MCE were used to sample the detectors at 400 kHz and 50 MHz, with downsampling and filtering turned off. Noise spectra were produced from the timestreams, and aliasing calculations were performed to assess the impact of the noise above the Nyquist frequency in the science band. All equations for the modeled noise components are taken from the reference [30]. The total NEP of a detector is given by the quadrature sum of the noise components given above as:

$$NEP_{total}^2 = NEP_{photon}^2 + NEP_{tfn}^2 + NEP_{Johnson}^2 + NEP_{amp}^2 + NEP_{aliased,total}^2 \quad (4.1)$$

The contributions are the photon noise, phonon or thermal fluctuation noise, the Johnson noise of the TES and shunt, the unaliased in-band amplifier noise, and the total aliased noise, respectively. This final term combines the major aliased contributions of the out-of-band amplifier and excess detector noise. The DC contributions to the noise are given in the equations below. The effect of ETF on the transfer functions of the noise contributions can be seen in figure 4.5.

Photon noise, from shot noise and Bose noise, is the largest contribution to the unaliased noise levels in the science band. The contribution is estimated from:

$$NEP_{photon}^2 = NEP_{shot}^2 + NEP_{bose}^2 = 2h\nu Q_{load} + \frac{2Q_{load}^2}{\nu \frac{\Delta\nu}{\nu}} \quad (4.2)$$

where ν is the band center, $\frac{\Delta\nu}{\nu}$ is the fractional bandwidth, and Q_{load} is the optical load on the detectors.

The thermal fluctuation, or phonon noise, is attributed to thermodynamic fluctuations of the detector's thermal impedance. After photon noise, this is expected to make the largest contribution

to the unaliased noise level in the science band for the target conductance. It is given by:

$$NEP_{ifn}^2 = 4k_B T_c^2 G_c F(T_c, T_{bath}) \quad (4.3)$$

where $F(T_c, T_{bath})$ is a factor that accounts for nonlinear thermal conductance and is estimated at 0.5.

Johnson noise from the detectors and shunts contributes little to the noise within the science band. The detector Johnson noise could make a significant contribution but is suppressed at low frequency by the loop gain in electrothermal feedback. The DC contribution from these sources are estimated from:

$$NEP_{Johnson}^2 = NEP_{tes}^2 + NEP_{sh}^2 = 4k_B T_c R_{tes} I_{tes}^2 \frac{1}{\mathcal{L}_i^2} + 4k_B T_{sh} R_{sh} I_{tes}^2 \frac{(\mathcal{L}_i - 1)^2}{\mathcal{L}_i^2} \quad (4.4)$$

The amplifier noise contributed by the cold and warm electronics has been estimated by sampling each channel at 50 MHz with the detectors superconducting and the multiplexer turned off, or open-loop. The noise spectra of the TES-coupled channels are rolled off by the Nyquist filters at ~ 300 Hz, down to the unaliased noise level of the the amplifier around 200 MHz. The spectra are then rolled off around 1 MHz by antialiasing filters in the MCE. Using the open-loop gain of the SQUID amplifiers, the noise was referred to input to give the NEI of the amplifier (§ 2.9.5.11). The mean of all the channels is $NEI_{amp} \simeq 2$ pA/ $\sqrt{\text{Hz}}$, in agreement with the design budget for the unaliased amplifier noise. The 50 MHz spectra were then downsampled to 15.5 kHz to estimate the aliasing. This gave $NEI_{amp/aliased} \simeq 25$ pA/ $\sqrt{\text{Hz}}$, or $NEP_{amp/aliased} \simeq 18$ aW/ $\sqrt{\text{Hz}}$, for the mean aliased amplifier noise. This is included in the total aliased noise given below that is attributed mostly to excess noise. So the aliased amplifier noise is not included in the final noise total to avoid over-counting.

Excess noise is not included in the TES noise model so it must be estimated from the detectors. It has been estimated in the characteristic detector analyzed below from the quadrature difference of the total predicted and the measured noise spectra. As shown in Figure 4.5, the in-band, unaliased NEI of the excess noise at low frequency is negligible, but the aliased contribution is significant. The high-frequency noise in the detector is dominated by excess noise, and thus so is the aliased noise. The in-band aliased noise is determined by downsampling the 400 kHz measured spectrum to a 15.5 kHz sampling rate.

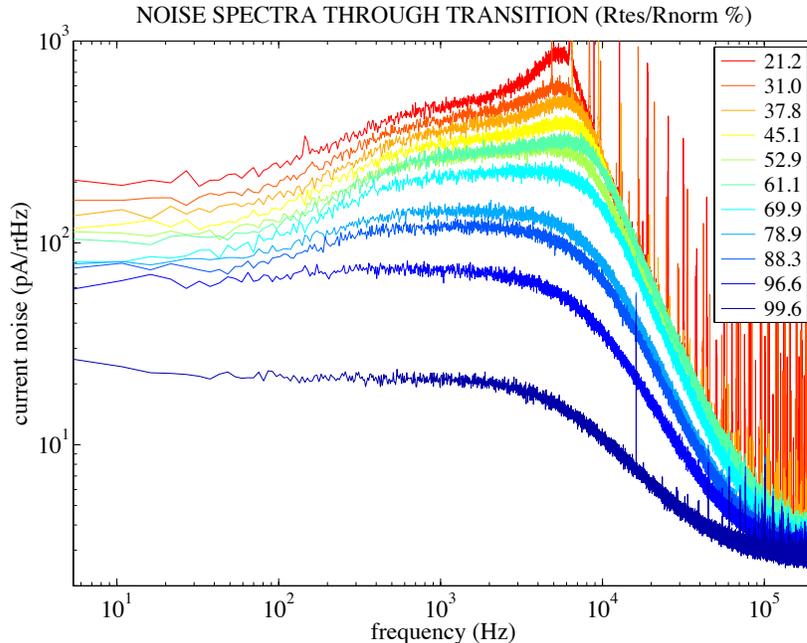


Figure 4.4: Noise spectra through the transition for a characteristic detector sampled at 400 kHz. As the detector is biased lower in the transition an excess noise plateau develops and eventually the detector begins to oscillate. For the standard observation bias this detector is operated at 86% fractional resistance.

4.1.6.2 individual detector

In this section the measured spectrum of a characteristic detector is compared to the full, frequency-dependent noise model derived from its measured parameters. Noise spectra of a characteristic, optically coupled detector are shown in figure 4.4. The spectra were acquired at 400 kHz for a range of detector biases through the transition. The fractional resistance ranged from very high in the transition ($R_{\%} = 99.6\%$) where the detector is normal to very low in the transition ($R_{\%} = 21.1\%$). Moving down through the transition, the low-frequency noise level increases and an excess noise plateau develops. As the detectors are biased lower into the transition the excess noise plateau rises. The noise power above the Nyquist frequency increases, increasing the level of aliased noise. Very low in the transition the detector begins to oscillate, with a resonance in the spectra forming around 6 kHz.

The noise spectrum of the same detector at the standard observation bias used before mid-September 2010 is shown in figure 4.5. For the standard observing bias this detector was operated at $R_{\%} = 86\%$. As seen in figure 4.4, this detector was operated far from the onset of thermal oscillations low in the transition. The various noise contributions predicted from the detector's

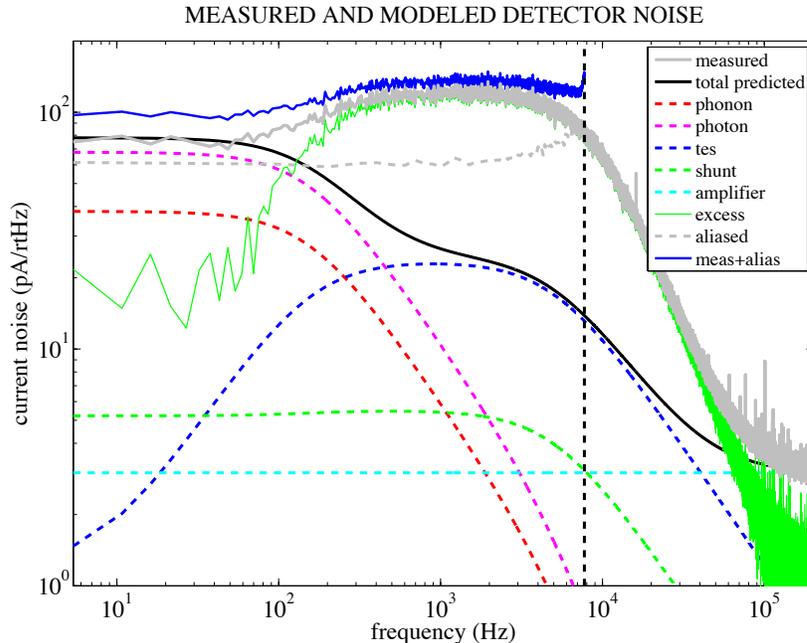


Figure 4.5: Measured and modeled current noise for the characteristic detector at the observation bias. The aliased noise level is calculated from the measured noise spectrum above the Nyquist frequency, shown as the dashed vertical line. The total predicted noise is calculated from the quadrature sum of the noise components modeled using the measured detector parameters.

measured parameters have been overlaid with the measured spectrum. The major low-frequency noise contributions are made by the phonon, photon and total aliased noise. The measured noise agrees with the total predicted noise at low frequency when modeled with $Q_{load} \simeq 6$ pW. This value is consistent with the range of loading measured during the austral summer (§ 4.1.3). The aliasing calculation gave ~ 60 pA/ $\sqrt{\text{Hz}}$ total aliased noise at low frequency for this detector.

The quadrature difference was taken of the measured and total predicted noise spectra. This provides a rough estimate of the transfer function of the excess noise spectrum, which roughly follows that of the Johnson noise. In addition, the quadrature sum of the measured spectrum and estimated aliased spectrum was taken. This value is shown as the “meas+alias” spectrum in the figure. The rms of the spectrum is ~ 100 pA/ $\sqrt{\text{Hz}}$ in the science band. This is consistent with the value measured for the detector in the standard closed-loop multiplexing mode. This confirms that the open-loop gain used to refer the open-loop 400 kHz spectra to the input coil were properly measured.

The same analysis has been performed on the noise spectra of dark detectors for which the photon noise is negligible. The low-frequency agreement of the measured and modeled noise is also

quite good. This suggests that the detector bias parameters have been accurately measured, and the responsivity and scaling of the NEI are approximately correct. For comparison with general estimates of noise below, the significant low-frequency modeled noise components for this characteristic detector are: $NEI_{tfn} \simeq 38 \text{ pA}/\sqrt{\text{Hz}}$, $NEI_{photon} \simeq 68 \text{ pA}/\sqrt{\text{Hz}}$, $NEI_{aliased} \simeq 60 \text{ pA}/\sqrt{\text{Hz}}$, for a $NEI_{total} \simeq 98 \text{ pA}/\sqrt{\text{Hz}}$ including aliased noise.

4.1.6.3 median detector parameters

The noise fitting procedure has been validated for a few characteristic light and dark channels. In the remainder of this section the expected levels of each noise contribution in the science band has been estimated from the median measured detector parameters from standard observing conditions. The estimates can be compared to the noise modeled for the characteristic detector. A DC power-to-current responsivity of $s_{I,dc} = -1.38 \text{ pA/aW}$ has been used to convert between NEP and NEI referred to the input coil of the first-stage SQUID. This value has been derived from the median detector bias parameters (R_{tes} , I_{tes}) during observation for the case of strong electro-thermal feedback given by equation 2.4.

For the median detector parameters the $NEP_{tfn} \simeq 27 \text{ aW}/\sqrt{\text{Hz}}$ and $NEI_{tfn} \simeq 38 \text{ pA}/\sqrt{\text{Hz}}$. Using the mean operational parameters of the detectors: $R_{tes} \simeq 61 \text{ m}\Omega$, $I_{tes} \simeq 13.5 \text{ }\mu\text{A}$, $\mathcal{L}_i \simeq 21.6$, along with $R_{sh} = 3.0 \text{ m}\Omega$ and $T_{sh} \simeq 500 \text{ mK}$, the combined Johnson noise of the TES and shunt is estimated at $NEP_{Johnson} \simeq 4 \text{ aW}\sqrt{\text{Hz}}$ and $NEI_{Johnson} \simeq 5 \text{ pA}\sqrt{\text{Hz}}$. The total aliased noise levels for the detectors were calculated from their 400 kHz spectra. The median value is $NEI_{excess/aliased} \simeq 64 \text{ pA}/\sqrt{\text{Hz}}$ and $NEP_{excess/aliased} \simeq 46 \text{ aW}/\sqrt{\text{Hz}}$, for the 15.5 kHz readout rate of the 2010 observing season. The median level agrees with the aliased noise in the characteristic detector of $60 \text{ pA}/\sqrt{\text{Hz}}$. The unaliased and aliased amplifier levels of $NEI_{amp/aliased} \simeq 25 \text{ pA}/\sqrt{\text{Hz}}$ or $NEP_{amp/aliased} \simeq 18 \text{ aW}/\sqrt{\text{Hz}}$ were given above. For a range of optical loading $Q_{load} \simeq 4\text{-}6 \text{ pW}$ the expected photon noise is $NEP_{photon} \simeq 41 - 56 \text{ aW}\sqrt{\text{Hz}}$ or $NEI_{photon} \simeq 56 - 77 \text{ pA}/\sqrt{\text{Hz}}$. For this level of optical loading, the BICEP2 detectors are approaching background-limited performance if aliased noise is excluded. However, the photon noise can be dominated during periods of good weather and low optical load by the $\sim 64 \text{ pA}/\sqrt{\text{Hz}}$ aliased excess noise.

The total expected noise given the range of loading is $NEP_{total} \simeq 68\text{-}78 \text{ aW}/\sqrt{\text{Hz}}$ or $NEI_{total} \simeq 93 - 107 \text{ pA}/\sqrt{\text{Hz}}$. The characteristic detector considered earlier has a total noise with aliasing of $NEI_{total} \simeq 98 \text{ pA}/\sqrt{\text{Hz}}$, which agrees with this range. The noise components from modeling the median device parameters is compared to the design components in table 4.3. In general the

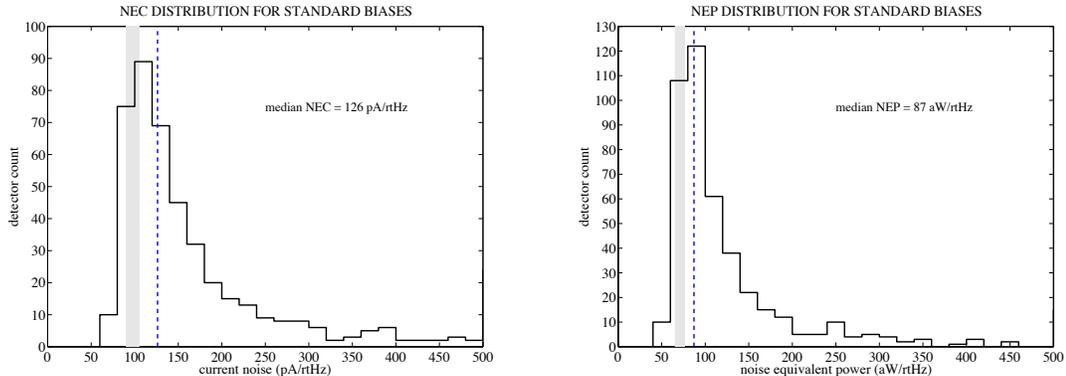


Figure 4.6: Histograms of the measured NEIs and NEPs for the operational detectors at the standard observation biases, based on 60 s timestreams acquired during the austral summer. The medians of the distributions are indicated by the dashed lines. The shaded regions show the range of NEIs and NEPs expected from the noise model based on median measured detector parameters.

contributions to the NEP are higher than the design. The high-frequency excess noise plateau is not predicted by the model and was not included in the design NEP. Thus, the total aliased noise is far higher than the design value.

Histograms of the measured NEIs and NEPs for the operational detectors at the standard biases are shown in figure 4.6. Noise spectra were derived from 60 s timestreams acquired at the mean elevation of the observation field during the austral summer. The NEI was taken from the rms of the current noise spectra within the science band. The NEPs were scaled from the measured NEIs using the individual power-to-current responsivity of the detectors, calculated from V_{tes} for strong ETF. The median measured $NEI = 126 \text{ pA}/\sqrt{\text{Hz}}$ and the median measured $NEP = 87 \text{ aW}/\sqrt{\text{Hz}}$. This median measured NEP is more than double the $NEP_{total,design} = 42 \text{ aW}/\sqrt{\text{Hz}}$.

The ranges of NEP and NEI calculated from the median detector parameters, shown as the shaded regions in figure 4.6, are lower than the median of the actual measured values shown as the vertical dashed lines. The calculated range still falls within the bulk of the measured distribution. The agreement would likely be better if the noise was modeled from the actual distribution of measured detector parameters.

4.1.7 aliased noise

Aliased excess noise dominates the noise of a significant fraction of the BICEP2 detectors at a 15.5 kHz readout rate. Based on estimates of noise from the median detector parameters, the total NEP without aliasing should be in the range of $49 - 62 \text{ aW}/\sqrt{\text{Hz}}$ for 4–6 pW of loading. When the median aliased contribution of $46 \text{ aW}/\sqrt{\text{Hz}}$ is included the expected range increases to

68 – 78 aW/ $\sqrt{\text{Hz}}$. In the absence of aliased excess noise there would still be aliased amplifier noise, which was estimated at $NEP_{amp/aliased} \simeq 18 \text{ aW}/\sqrt{\text{Hz}}$. If this contribution is included instead then the total range increases only a little to 52 – 65 aW/ $\sqrt{\text{Hz}}$. This value is approaching but still not as small as the design goal of $NEP_{total,design} = 42 \text{ aW}/\sqrt{\text{Hz}}$.

The aliased excess noise seen in 2010 could have been addressed by changing out the focal plane for NYQ chips with higher inductance or TESs with less excess noise. The former would require additional heat capacity to maintain the same margin on stability. Since the origin of the excess noise is not well-understood, the latter option would be difficult to achieve. Changing out the focal plane comes with additional risk and would require extensive recharacterization of a new focal plane.

In the end, the best option was to pursue a higher readout rate for the 2011–2012 observing seasons. Based on aliasing calculations of the 400 kHz noise spectra, increasing the readout rate to 20 kHz would reduce the median aliased noise level to 36 aW/ $\sqrt{\text{Hz}}$, and increasing it to 30 kHz would reduce it to 22 aW/ $\sqrt{\text{Hz}}$. Above a 20 kHz readout rate the aliased noise is expected to fall below the 36 aW/ $\sqrt{\text{Hz}}$ expected from photon noise for 4 pW of loading. As shown in table 4.3 the total NEP was expected to fall to 55 aW/ $\sqrt{\text{Hz}}$ going to a 30 kHz readout rate. This was beginning to approach the design goal. When the additional gain from the higher-than-expected optical efficiency was included, reaching the design NET was feasible.

4.1.8 optimization

A significant campaign was undertaken in the fall of 2010 to improve per-detector NETs and multiplexer yield to realize higher array sensitivity and mapping speed. The TES biases were optimized to improve the per-detector sensitivity. The SQ1 biases were overhauled to mitigate the crosstalk introduced by the wiring issues. Finally the tuning algorithm was improved to stabilize the choice of SQUID lock points and reduce multilock points in the open-loop composed $V - \Phi$ response. These changes were implemented beginning September 14, 2010 and have been used for the remaining observations. During 2010 the readout rate was fixed at the conservative value of 15.5 kHz to ensure the data were not contaminated by settling crosstalk. Studies indicated that an increase to 25 kHz would significantly improve the mapping speed of the instrument without introducing substantial crosstalk. In January 2011 the multiplexer rate was increased to 25 kHz and the remaining observations have been made at this rate. The details of those optimizations are given in this section and the impact on the instrument sensitivity is given in § 4.1.9. To summarize, BICEP2 has had three unique instrument configurations. It ran from January 2010 to September 14, 2010 with the

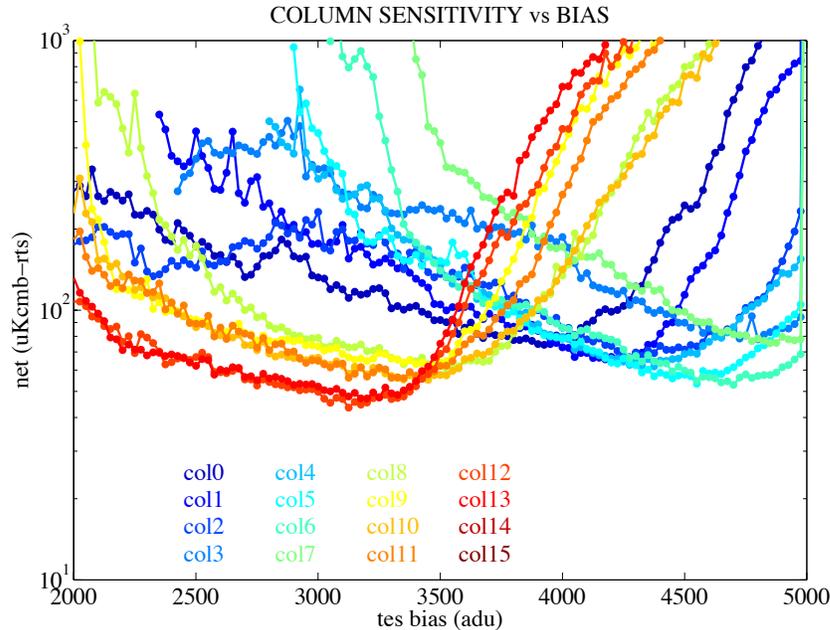


Figure 4.7: Sensitivity profiles for the 16 detector columns used to select the optimal detector biases. The total NET for a column is determined by inverse variance weighting the NETs for the 32 detectors biased in that column. The sensitivities were estimated from an assumed zenith temperature and are not used for absolute sensitivity estimates, only the relative values are used for bias selection.

initial detector and readout configuration. It then ran until January 2011 with new SQUID and TES biases that improved the yield and sensitivity. From January 2011 onward it has been running with these optimized biases at 25 kHz with a significant improvement in mapping speed.

4.1.8.1 detectors

Detector biases have been selected in order to minimize the instrument NET, while ensuring stability of the detector responsivities. The 32 TESs within a multiplexer column share a pair of bias lines and there are a total of 16 bias lines for the columns. The biases selected for a column are a compromise for the TESs within the column. TESs that are better matched will be more optimally biased due to the overlap of the individual sensitivities at a given bias for the column. Due to the wiring of the shunt resistors in series on the NYQ chips, the biases must also be selected so that detectors do not oscillate and cause crosstalk in the other channels of the column. This occurs only low in the transition and there was no evidence of oscillating TESs for the biases chosen for standard observations.

The biases for each column were determined from per-column NET versus TES bias profiles

shown in figure 4.7. The profiles were created by combining measurements of the current noise of each detector with an estimate of its optical responsivity at each bias step. The profiles reflect the fact that as the detectors are biased lower in the transition their current noise decreases and their responsivity increases. The profile hits a minimum and then begins slowly rising due to aliasing of out-of-band excess noise that grows lower in the transition, as shown in figure 4.4. The spread in the minima for the NET profiles is due to the difference in thermal conductances for the four detector tiles. Tiles 3 and 4 correspond to columns 8 through 15, whose minima are tightly grouped and require lower applied bias. Detectors on these tiles have lower and more consistent thermal conductances, so less electrical power is needed to maintain them in the transition and there is less spread in the optimal biases.

The per-column NETs were derived from noise spectra and elevation nods acquired at each bias step. The TESs were driven normal and the detector bias was stepped down finely through the transition. The telescope was pointed at an elevation of 55° , which is representative of the loading during normal observations. Sixty second timestreams were acquired for all channels at each bias step to estimate the current noise. The noise level was taken from the average across the 0.1–1 Hz science band. The current noise was referred to a pseudo-NET by scaling the gain from the elevation nod using an assumed zenith temperature of $12.3 \text{ K}_{\text{CMB}}$ typical of summer observing conditions. The per-column sensitivity at each bias step was then calculated by inverse-variance weighting the individual NETs of the functioning detectors within a column. This method was only used to optimize the detectors and not to determine the actual per-detector NETs, due to uncertainty in the choice of zenith temperature. To accurately characterize the sensitivity an absolute calibration must be applied, as described in § 4.1.9.

Conservative biases were chosen from figure 4.7 for the initial configuration of the instrument during deployment. The biases were selected lower in the transition than the minima, far from the regions of high NET slope near the minima at high bias. This was done to ensure the detectors were well into their transitions and had stable optical responsivities. However, using these lower biases gave an instrument NET that was not optimal. In September 2010 a number of optimizations were made including a change of the detector biases. Studies of the stability of the optical responsivity showed that detectors that were biased close to their minima and near the high NET slope did not show signs of instability. So, the biases for most of the columns were increased to the minima of the profiles. The improvement in the per-detector sensitivities was expected to be $\sim 10\%$.

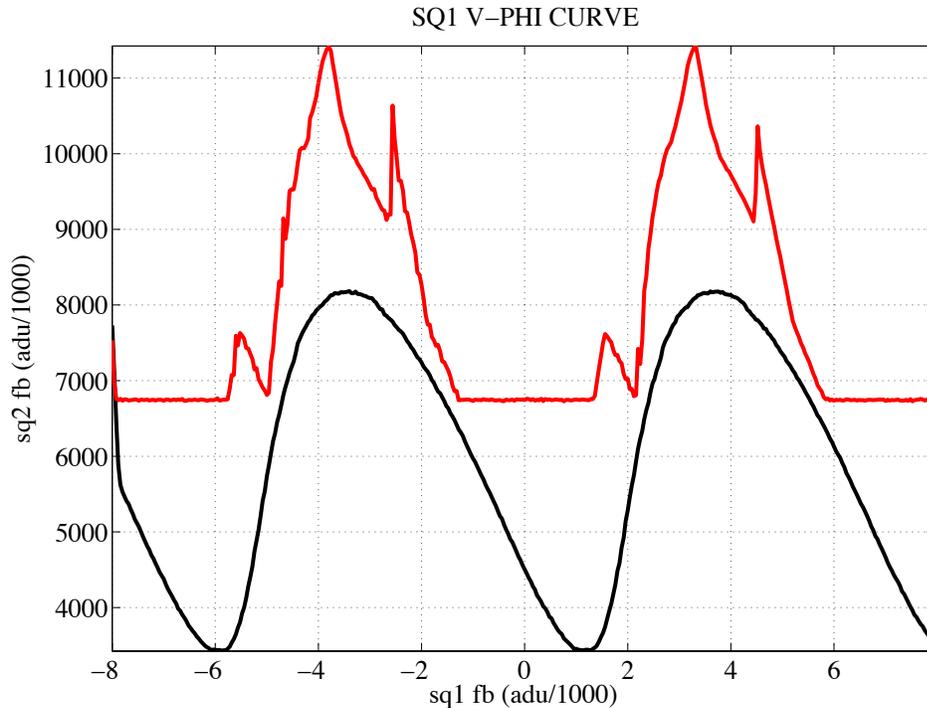


Figure 4.8: SQ1 $V - \Phi$ response before and after an overhaul of the SQUID biases. The top (red) curve was the $V - \Phi$ response for the initial choice of SQ1 bias. The SQ1 is under-biased and three additional SQ1s are being turned on in the column at the same time due to shorting current to other rows. The bottom (black) curve is the response after a reoptimization of the SQ1 biases to remove the multiple SQ1 response. The curves have been offset for clarity. They are plotted as the SQ2 flux feedback that must be applied in software to linearize the amplifier in response to a ramp of the SQ1 flux feedback.

4.1.8.2 multiplexer

The multiplexer has been optimized for maximum yield and readout rate, while minimizing crosstalk. The SQUID tuning parameters have been chosen for highest yield at the expense of some of the end-to-end amplifier gain. The amplifier noise referred to input is thus larger, but as will be discussed below, it is subdominant to other noise sources. The optimization of the multiplexer included an overhaul of the SQUID biases and changes to the default tuning algorithm.

Initial biases were selected to address an issue with resonances seen in many of the SQ1s. These resonances are due to the proximity of the input and feedback coils deposited on top of the planar SQUIDs and are usually eliminated by intracoil damping resistors (§ 2.5.2). However, the damping was insufficient in the MUX chips deployed in BICEP2 and small resonance bumps appear near the middle of many of the SQ1 $V - \Phi$ curves. If the total composed lock point is selected on one of these bumps then the amplifier gain will be low and nonstationary, which affects the level and

stationary state of the input-referred amplifier noise. To avoid these resonances the SQ1s were initially biased below I_c^{max} as shown in figure 4.8. This causes flat regions in the response curve where superconducting current flows with no generation of voltage. Underbiasing the SQUIDs shifts the resonances away from the middle of the $V - \Phi$ curve so that they are unlikely to be locked on. Small resonances are barely visible near the peaks of the black $V - \Phi$ curve in the figure. Underbiasing SQUIDs also has the effect of giving them higher gain but lower dynamic range. Note the increased steepness but lower modulation depth for the red (underbiased) curve compared to the black curve in figure 4.8. The initial SQ1 biases were also selected so that a small bias was applied to off-state SQ1s even when they were not visited. The off-state bias was chosen to be well below I_c^{min} . This was done as a precaution to reduce transients caused by switching SQ1s on and off that was believed to lead to longer settling times and higher crosstalk (§ 4.1.5.2).

The biases were reoptimized in September 2010 to address a few problems found while studying beam maps acquired during deployment. Maps for some of the channels showed the main beam and additional beams of equivalent power (§ 4.1.5.6). The issue was eventually traced back to a wiring short between the SQ1 bias lines for several rows (3, 6 and 8) somewhere in the interior of the cryostat. The channels that displayed this crosstalk had SQ1 $V - \Phi$ curves like the one shown in red in figure 4.8. Three additional SQ1s are biased on in the column in addition to the desired SQ1. Channels with this level of crosstalk must be cut from analysis. Evidence of multiple SQ1s showed up in $V - \Phi$ curves for 96 light channels in 64 detector pairs.

This problem was mitigated by selecting biases for the shorting rows that would turn the row of interest on without providing enough current to exceed the I_c^{min} of the row it's shorted to. There are 16 SQ1s with various critical currents that are controlled by a single row bias, so the bias selection required optimization to maximize the number of recovered channels. The off-state SQ1 bias was set to zero since it contributes to the total bias current. Raw, 50 MHz, multiplexed timestreams were acquired from the ADC on each column to assess the impact on the switching transients. No evidence was seen for increased switching transients by setting the off-state biases to zero. Studies of settling crosstalk also showed no evidence for increased contamination (§ 4.1.5.2). After the bias reoptimization, $V - \Phi$ curves like the black curve in figure 4.8 showed no evidence for multiple SQ1s being turned on. Beam maps acquired during the 2010–2011 austral summer showed the parasitic beams had been eliminated. This reoptimization recovered 87 light channels in 55 pairs that would have been cut from analysis.

The largest improvement in sensitivity was made by increasing the multiplexer rate. As described

in § 4.1.6 the per-detector sensitivity was significantly degraded at a readout rate of 15.5 kHz due to aliasing of out-of-band excess detector noise. A higher readout rate reduces the aliasing but may also increase crosstalk, mainly from increased settling crosstalk from a shorter row dwell time (§ 4.1.5.2). The sensitivity versus bias study in § 4.1.8.1 was repeated at higher readout rates. As expected from aliasing simulations, the sensitivity greatly improved as the readout rate was increased. A rate above 25 kHz could not be used because of difficulties in reliably tuning the SQUIDs above that rate. The study showed that the increase to a 25 kHz readout rate would give a factor of 0.84 reduction in instrument NET or a factor of 1.38 increase in mapping speed. A study of the settling crosstalk in beam maps showed that increasing to this rate did not significantly increase crosstalk levels (§ 4.1.5.2). So, beginning in January 2011 a readout rate of 25 kHz has been used.

4.1.8.3 yield

Of the 512 TESs on the BICEP2 focal plane, 441 are optically coupled and operational (86%). Based on electrical testing of the arrays prior to integration in the focal plane the fabrication yield was 99% with 507 operational detectors. The microstrip coupling the antenna to the TES has been disconnected for 12 detectors to create dark detectors used to assess common-mode magnetic and thermal pickup in the TESs. The MUX chips are cryogenically screened prior to integration, which revealed a loss of six channels. After the multiplexers were integrated and the SQUIDs were properly biased and tuned, various issues with the first-stage SQUIDs have further reduced the number of operational channels by 32. In addition, an electrical connection issue on one multiplexer column (column 14) has eliminated the ability to read out its 32 optically coupled channels. Finally, a few detectors are unusually noisy and have been flagged and removed from analysis. Many of these issues overlap, but the net result is that 71 antenna-coupled TESs in 60 detector pairs are nonoperational. Due to pair differencing 196 detector pairs will be used in analysis, for a total of 392 operational detectors. The ideal yield focal plane yield would be 500 optically coupled channels or 250 pairs. These yield issues result in a factor $\sqrt{\frac{196}{250}} \simeq 0.89$ degradation in the instrument sensitivity from the ideal. Assuming all 64 pairs that showed evidence of multiple SQ1s on would have to be cut from analysis, there are 132 useful pairs before the bias changes and 187 pairs after. The additional SQ1 shorts then give a factor of 0.73 degradation in the instrument sensitivity from ideal. After the bias changes this improved to only a factor of 0.86 degradation. The bias changes thus gave a factor of 0.84 lower instrument NET or a factor of 1.42 increase in mapping speed. Those changes were not implemented until late 2010, so they mostly improve the mapping speeds of the 2011 and 2012

seasons. When the improvement of the readout rate and bias changes are considered together there should be a factor 1.96 increase in mapping speed between the 2010 and 2011–2012 seasons.

4.1.9 sensitivity characterization

NETs have been estimated using two methods: directly from timestreams, and from maps coadded from those timestreams. Data were acquired during ordinary CMB observations and processed into maps in the standard way. The sensitivity was estimated using a subset of four days of data from each of the observing seasons. To confirm that using a small subset of data to represent months-long observing conditions is valid, the analysis was repeated on a larger data set spanning six months. In addition, the sensitivity of BICEP1 was derived from four days of data with the same procedure to compare with the published BICEP1 sensitivity numbers and the BICEP2 NETs reported here.

Detector timestreams were processed by first deconvolving the effects of antialiasing filters applied to the data and by removing glitches (§ 3.8). Next, each detector timestream was normalized across the focal plane from a fit of its response during an elevation nod to a small change in air mass. The timestreams of detectors within a polarization pair were then summed and differenced. Since detectors within a pair are sensitive to orthogonal polarization, the difference is used to measure Stokes Q and U. The sum gives a measurement of temperature. Common-mode atmospheric signal is well-removed from the pair-differenced data, but is still present in the sum. Additional atmospheric signal was removed by fitting out a third-order polynomial from each half (left- or right-going) of the fixed elevation azimuth scans. The variance was calculated for each half scan and detector, and the inverse was used to weight the coadded signal in map making. The current noise is referred to NET by using absolute calibrations derived from cross-correlations with WMAP maps. Absolute calibration numbers were derived from maps that included only the four days of data used to estimate the NETs. The signal to noise in these maps is adequate to give reliable absolute calibrations.

For clarity, only the analysis of two instrument configurations are presented here. The majority of the data acquired in 2010 had the initial SQUID and TES bias settings, since those biases were not reoptimized until mid-September 2010. Those early 2010 observing settings are referred to as the 2010 settings below. The other configuration is for 2011, which had improved biases and a 25 kHz readout rate. The settings for the current 2012 observing season are the same as the 2011 season. The final mapping speed comparison presented in table 4.4 does not include any channel cuts for such things as crosstalk. So, the numbers presented for the 2010 data should be an optimistic estimate of the instrument sensitivity and mapping speed since some of the channels with multiple

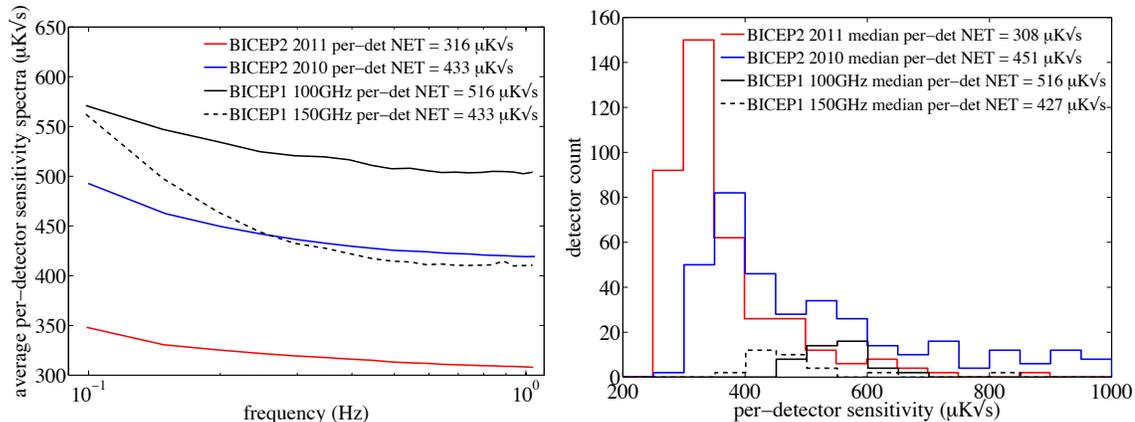


Figure 4.9: Per-detector sensitivity spectra averaged over all detectors and the distribution of per-detector NETs derived from average sensitivity spectra for each detector. The NETs are calculated from the rms of the sensitivity spectra over the 0.1–1 Hz science band. The median per-detector NETs agree with the values derived from the averaged spectra within 4%.

SQ1s on will have to be cut. Fewer cuts are expected for the 2011–2012 seasons since there has been no evidence for multiple SQ1 crosstalk from beam maps after the SQUID bias changes.

4.1.9.1 timestream-based sensitivity estimates

The first method for estimation of the per-detector and instrument NETs was to derive them directly from processed timestreams. The one-sided power spectral density (PSD) and variance were calculated for each half scan and all pair-differenced detector timestreams. The half-scan variance was then used to create an inverse-variance weighted average PSD for each of the differenced detector pairs. The inverse-variance-weighted average PSDs for each pair were then inverse-variance weight averaged for each spectral bin to form an average per-detector sensitivity spectrum. No anomalously high weights were included in these averages. The PSDs were then scaled from units of ADU^2/Hz to units of NET ($\mu\text{K}\sqrt{\text{s}}$) by dividing the square root of the PSDs by $\sqrt{2}$ and applying the absolute calibration. During timestream processing, detector signals within a pair are differenced and divided by 2, which gives a per-pair NET. The PSDs were multiplied by $\sqrt{2}$ to scale them to per-detector NETs.

The per-detector sensitivity spectra averaged over all detectors are shown for each instrument configuration in figure 4.9. The per-detector sensitivities given in the figure are taken from the rms over the 0.1–1 Hz science band. These values are included in table 4.4 for comparison with the NETs derived from the maps. The table also includes an estimate of the instrument NET for comparison of sensitivity and mapping speed. The instrument sensitivity was estimated by taking the inverse

Table 4.4: Sensitivity and mapping speed comparison for BICEP1 and BICEP2 from the map-based and timestream-based methods. The instrument sensitivity shown for BICEP1 accounts for both frequency bands. The map-based and timestream-based methods agree within 2%.

	light detector count	per-detector NET $\mu\text{K}\sqrt{2}$		instrument NET $\mu\text{K}\sqrt{2}$		mapping speed vs BICEP1		mapping speed vs BICEP2 2010	
		time	map	time	map	time	map	time	map
BICEP1	–	–	–	54.3	53.5	–	–	–	–
100 GHz	44	516	505	–	–	–	–	–	–
150 GHz	32	433	426	–	–	–	–	–	–
BICEP2 2010	384	433	422	21.8	21.5	6.20	6.19	–	–
BICEP2 2011	394	316	313	15.9	15.8	11.7	11.5	1.88	1.85

square root of the sum of the inverse of the individually averaged PSDs. The rms was then taken over 0.1–1 Hz of the resulting instrument sensitivity spectrum to give the instrument NET.

The per-detector NETs were determined from the rms over the science band for each of the individually averaged, per-detector sensitivity spectra. The distributions of the per-detector NETs for each instrument configuration are shown in the histogram of figure 4.9 along with the median of each distribution. The median per-detector NETs are within $\sim 4\%$ of the values derived from the sensitivity spectra averaged over all detectors. The count in each bin of the histogram has been doubled to account for the two detectors that were initially averaged together in the pair-differenced timestream. The units have been properly scaled to per-detector NET.

The per-detector NETs derived for BICEP1 are $516 \mu\text{K}\sqrt{\text{s}}$ and $433 \mu\text{K}\sqrt{\text{s}}$ for the 100 GHz and 150 GHz channels, respectively. These numbers are less than 4% lower than the published NETs of $530 \mu\text{K}\sqrt{\text{s}}$ and $450 \mu\text{K}\sqrt{\text{s}}$ [59]. The difference can be attributed to analyzing four days of particularly good weather, while the published numbers are averaged over the entire first two years of observations. To check the validity of representing entire observing seasons by 4-day subsets of data, an additional estimate of the BICEP2 2010 NETs was made from data acquired between April and September 2010. A per-detector NET of $448 \mu\text{K}\sqrt{\text{s}}$ has been estimated from the larger data set and is only $\sim 3\%$ higher than the $433 \mu\text{K}\sqrt{\text{s}}$ derived from the small data set. The close agreement validates the representation of an entire observing season with only four days of data. It also confirms the discrepancy between the published BICEP1 sensitivity and the number presented here.

4.1.9.2 map-based sensitivity estimates

To confirm the map noise is integrating down with observation time as expected, it’s important to derive sensitivity numbers from the maps as well. The processed, pair-differenced timestreams were

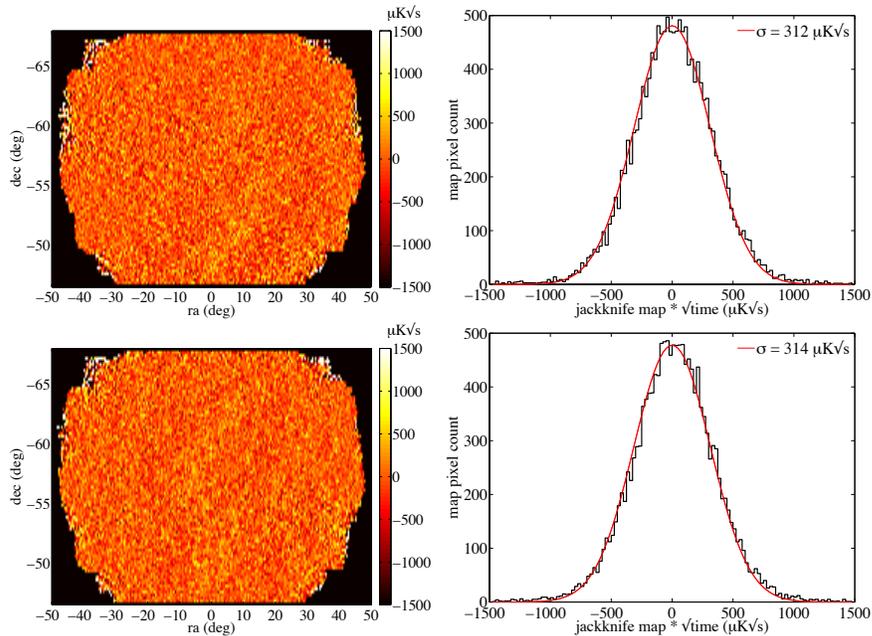


Figure 4.10: Maps and histograms used to derived the map-based instrument sensitivity for the 2011 observing configuration. *Left:* The product of the scan-direction jackknife map and the square root of the integration time map for Stokes Q (top) and U (bottom). *Right:* The distribution of the values in the sensitivity maps fit with a Gaussian (red curve). The resulting σ is an estimate of the per-detector NET for Q (top) and U (bottom). The average σ from Q and U is used in table 4.4.

binned into maps with a resolution of 0.25° per pixel. The inverse of the half-scan variance is used to weight the detector data accumulated in each pixel. For each instrument configuration, Stokes Q and U maps were constructed for the left- and right-going half scans. The temperature maps were not used for sensitivity estimates because they are created from pair-summed timestreams and have too much atmospheric $1/f$ noise. A jackknife map was then formed by differencing the two scan-direction maps to remove signal from the maps while preserving the noise.

The jackknife maps are multiplied by the square root of the map of integration time, to create the maps shown in the left panels of figure 4.10. The map of integration time is made by binning $1/f_{rate}$ seconds for each detector sample, where $f_{rate} \simeq 100$ is the final downsampled data rate. A Gaussian is fit to the distribution across the map pixels. The resulting σ gives an estimate of per-detector NET as shown in the right panels of figure 4.10. The Q and U maps are formed from pair-differenced data that averages two detector measurements. The integration time per map pixel is summed over all detector pairs, so the product of the map and square root of time will represent a per-pair NET averaged over all detector pairs. However, since the pair-differenced measurements are divided between the Q and U maps, they both acquire a factor of $\sqrt{2}$ increase in NET, which is the same factor that must be applied to scale them to per-detector sensitivities. The per-detector

Table 4.5: Designed and measured sensitivity. The measured values are the average between the two measurement methods in table 4.4. The per-pair NEQ assumes the current upper limit on the cross-polarization response $\epsilon < 0.005$. The light detectors do not include final analysis cuts. The BICEP2 2012 observing season has the same observing configuration and sensitivity as the 2011 season.

	light detector count	per-detector NET $\mu\text{K}\sqrt{2}$	instrument NET $\mu\text{K}\sqrt{2}$	per-pair NEQ $\mu\text{K}\sqrt{2}$	instrument NEQ $\mu\text{K}\sqrt{2}$
design	496	330	14.8	233	14.8
BICEP2 2010	384	428	21.7	304	21.8
BICEP2 2011	394	315	15.9	224	16.0

NETs from the Q and U maps are averaged and given in table 4.4. The instrument NETs shown in the table were derived by dividing the per-detector NETs by $\sqrt{N_{det}}$. These values do not take into account additional channel cuts that occur in the final analysis. Such cuts will likely remove many of the channels that show multiple SQ1s. Changes in yield will affect the instrument NETs, but the per-detector NETs given in the table are still accurate.

The map-based per-detector NETs agree with the timestream-based estimates within 2%. The instrument NETs are only 0.6–2% higher when derived from the timestreams than from the maps. The close agreement of the map-based numbers with the timestream-based numbers indicates that uncorrelated noise in the detector timestreams is averaging down as expected.

4.1.9.3 sensitivity and mapping speed

The results of the sensitivity of timestream-based and map-based sensitivity study are summarized in table 4.4. A comparison of results — averaged for the two methods — with the design goal is shown in table 4.5. The per-detector sensitivity of $428 \mu\text{K}\sqrt{\text{s}}$ for the 2010 observing season does not meet the design target of $330 \mu\text{K}\sqrt{\text{s}}$. The degradation of the sensitivity in 2010 was due to the aliasing of out-of-band excess noise. When the readout rate was increased from 15.5 kHz to 25 kHz and the biases were optimized the 2011 per-detector NET became better than the design target at $315 \mu\text{K}\sqrt{\text{s}}$. The same sensitivity and mapping speed is expected for the 2012 as the 2011 season since the same readout rate and biases are used.

The fivefold increase in detectors and better per-detector NETs gave BICEP2 a factor of 6.2 increase in mapping speed over BICEP1 for the 2010 season. The 2011 season was even better, with a factor 11.6 increase in mapping speed over BICEP1. The optimizations made for the 2011 season gave a factor of 1.87 total increase in mapping speed over the 2010 season. Due to yield issues, though, the instrument NET for 2011–2012 is still a bit worse at $15.9 \mu\text{K}\sqrt{\text{s}}$ than the target of

14.8 $\mu\text{K}\sqrt{\text{s}}$. The 2011–2012 mapping speed is only a factor of 0.87 lower than the design goal.

Additional analysis was performed to discriminate the contributions of the readout rate increase and bias changes to the sensitivity. A 4-day subset of observing data acquired between mid-September and 2011 was analyzed. The instrument NET derived from this late 2010 data set is 19.1 $\mu\text{K}\sqrt{\text{s}}$. This corresponds to a factor 0.88 decrease in instrument NET or a factor 1.30 increase in mapping speed given by the bias changes. The total improvement between the early 2010 and 2011 seasons was a factor 1.87 increase in mapping speed. The readout rate increase was responsible for a factor of 1.44 additional increase in mapping speed.

The sensitivity to polarization fluctuations is given by the noise-equivalent Stokes Q (NEQ). It is derived from the per-pair NET, since a measurement of polarization requires a pair of detectors. The NEQ also accounts for the reduction in sensitivity from reduced polarization efficiency from cross-polar response as:

$$NEQ_{pair} = \frac{1}{1 - \epsilon} \frac{NET_{det}}{\sqrt{2}} \quad (4.5)$$

The per-pair and instrument NEQs in table 4.5 were derived from the upper limit of $\epsilon < 0.005$ from § 5.4 and the NET averages in the table. Due to the high polarization efficiency of the detectors the instrument NEQ is nearly the same as the instrument NET, or nearly ideal.

Chapter 5

Instrument Characterization

5.1 overview

In order to detect a faint B-mode polarization signal an instrument must have high sensitivity and careful control of systematics. The characterization of the instrument is presented with an assessment of the systematic errors that will limit the detection of a primordial B-mode signal. In order to reach the map depth necessary to detect such a small signal an instrument must integrate those maps over long periods of time with high sensitivity (§ 4).

There are several advantages to the characterization of the small, 26 cm aperture of BICEP2. The far-field is only ~ 70 m away, which means sources can be placed at a reasonable distance from the receiver and will produce a large response. The compact aperture also allows the entire beam to be filled with calibration sources during near-field characterization.

In the flat sky approximation the timestream response $d(t)$ of each detector to sky radiation with time is modeled by the equation:

$$d(t) = K_t \cdot \{n(t) + g \int A_e F(\nu) d\nu \int P(\Omega) d\Omega \times [T + \frac{1-\epsilon}{1+\epsilon}(Q\cos 2\psi + U\sin 2\psi)]\} \quad (5.1)$$

where K_t is the time-domain transfer function of the detector and electronics filters, $n(t)$ is the noise contribution, g is the detector's responsivity, A_e is the effective antenna area and is proportional to λ^2 , $F(\nu)$ is the total spectral response, $P(\Omega)$ is the antenna response as a function of position Ω , $\frac{1-\epsilon}{1+\epsilon}$ is the effective polarization efficiency given a cross-polarization response ϵ , and ψ is the absolute polarization angle. The temperature and polarization of the underlying sky signal, given

by the Stokes parameters T, Q, and U, can be recovered with precise measurements of these input parameters.

Uncertainty on the measurements of these parameters can create false B-mode polarization when Q and U are recovered from equation 5.1, or propagate to uncertainty in the polarization power spectra. The former will create false polarization that could exceed the signal level of the primordial B-mode polarization BICEP2 was designed to pursue. Furthermore, even if certain parameters are precisely calibrated systematic mismatch in a detector pair will cause contamination of the polarization maps.

Monte Carlo simulations are run with the measured parameters and their uncertainties to assess the impact on the polarization fidelity. In some cases the impact can be calculated analytically. In general, simulations must be run because the effect is largely dependent on the distribution of the parameters in the focal plane, and on the scan strategy. The benchmark for how well a parameter must be measured, or how much mismatch can be tolerated, is set by the level at which it introduces false signal in the B-mode polarization power spectrum at $\ell \sim 100$. The raw sensitivity for one year of integration with BICEP2 should easily allow for the detection of a B-mode signal at a level of $r \sim 0.1$. At $r = 0.1$ the B-mode polarization power spectrum ($\frac{\ell(\ell+1)}{2\pi}C_\ell^{BB}$) peaks at $\ell \sim 90$ at a level $0.007 \mu\text{K}^2$ (§ 1.2). Benchmark values listed in this chapter are the precision and limits on mismatch that must be reached to prevent spurious B-mode at this level.

If mismatch is measured at a level that introduces negligible polarization at the multipoles of interest, then there is no issue. However, if it produces B-mode in excess of the primordial signal being measured then the contamination must be removed as described in § 5.7.2. To accurately remove the signal the properties sourcing the spurious B-mode polarization must be measured precisely in order to keep from introducing large uncertainty in the cleaned polarization power spectra.

Each austral summer — the beginning of November through mid-February — the BICEP2 team deploys to the South Pole to make precision characterization measurements of the instrument. The instrument is kept cold and operational during the austral summer calibration runs, and the instrument has never been warmed up over its three-year observing campaign. CMB observations are routinely interspersed with calibration measurements to exercise the minimal changes made to the system. Although the summer CMB observations are not currently used in analysis, the quality of the data is high due to the superb observing conditions of the South Pole.

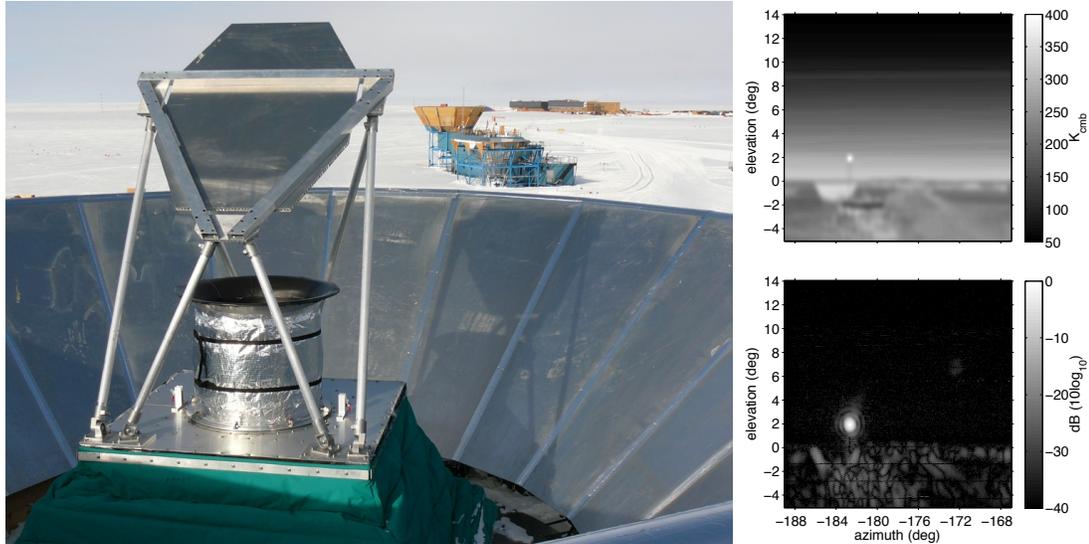


Figure 5.1: The setup and result of mapping a ground-fixed calibration source. *Left:* The observing configuration for mapping a calibration source mounted on a mast on the roof of MAPO. The far-field flat mirror is mounted to the boot platform and moves with the azimuth and elevation motion of the mount. In the background a mast can be seen extending just above the brown, plywood ground shield of Keck. *Right top:* A single-channel map of a bright, modulated source on the mast at $az = -182.5^\circ$, $el = 2^\circ$. The map has not been demodulated, and shows the response of the detectors to change in loading with elevation and from ground-fixed structures. *Right bottom:* The same map has been demodulated and plotted on a log scale to show the high signal-to-noise of the modulated source.

5.2 measurement techniques

Many aspects of the instrument characterization, including the beams and polarization, are measured by mapping ground-fixed sources and the moon. The ground-fixed calibration sources are mounted on a mast on the roof top of the Martin A. Pomerantz Observatory (MAPO) building. It is located 200 m from BICEP2 in the far field of the telescope. The moon never rises above $\sim 21^\circ$ elevation at the South Pole and the top of the mast is located at 2° elevation. Mapping these low elevations requires the use of a far-field flat mirror to clear the edge of the ground shield, as shown in figure 5.1. The mirror is rigidly fixed to the top of the boot platform and moves with the azimuth and elevation motion of the mount. The mirror is large enough to fully cover the beams emerging from the telescope. A pointing model that includes the additional parameters introduced by the mirror is used in the analysis of the beam and polarization measurements. The raw timestreams are deconvolved to remove a group delay introduced by the MCE filter that would otherwise cause a lag of the detector data behind the pointing data. The increased loading at low elevation requires the detectors to be biased on the aluminum transition. There should be no difference in the measurement of the beam

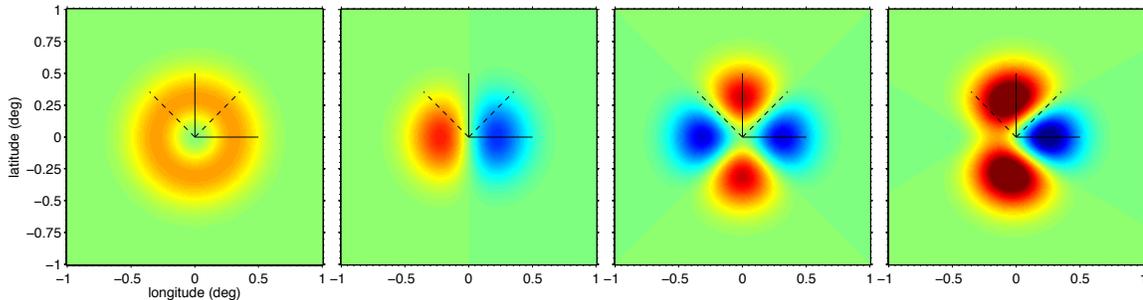


Figure 5.2: Decomposition of differential beam systematics. The difference of the beams has been taken as A-B. From left to right the components of the basis are: the monopole term from differential beam width; the dipole term from differential pointing; and the quadrupole term from differential ellipticity. The combination of the three differential effects is shown in the last panel. The solid (dashed) axes are the component of the effect relative to focal plane Stokes Q (U), for a complete decomposition basis.

and polarization properties using either the titanium or aluminum transitions, since the antenna is common to both TESs.

The ground-fixed sources include an amplified broadband noise source (BNS) and a thermal blackbody source. The BNS is an amplified thermal source with a 140–160 GHz passband. It’s high output is useful for the investigation of low-level effects, such as crosstalk and off-beam response. When placed on the mount the source produces a ~ 500 K_{CMB} signal at the telescope. The output amplitude is modulated using a PIN switch to improve signal to noise against the high-contrast horizon. The polarization of the emission is controlled by an output horn, which can be switched between circularly and linearly polarized, and a wire grid mounted in front of the source aperture. The thermal blackbody source is a box with a 26 cm aperture that contains a mirror and chopper blades covered with microwave absorber.¹ The output is chopped between ambient temperature and a flat mirror that directs the beams to the ~ 12 K_{CMB} sky to produce a ~ 5 K_{CMB} modulated signal.

5.3 beam characterization

The BICEP2 optics were designed to form circular Gaussian beams of 0.47° FWHM. The optics, focal plane hardware and the antennas cause departures from this ideality, forming beams that are well fit by elliptical Gaussians. The elliptical Gaussian beams are described by the parameters: r and θ that specify the location of each beam relative to the boresight in polar coordinates, the beam

¹The diameter was made as close to filling the ~ 1 m beams at the mast as possible.

widths² σ_a and σ_b along the major and minor axes, and α that specifies the rotation of the major axis relative to the radial vector to the boresight. The average beam width σ for a detector is defined as the mean of σ_a and σ_b .

Relative asymmetries of beams within a detector pair generates polarization contamination when they are differenced, leaking unpolarized CMB temperature fluctuations into Q and U. For elliptical Gaussian beams the differential mismatch can be decomposed into a basis of the monopole, dipole and quadrupole terms [54]. Monopole leakage from the beams arises from the differential beam size of two circular Gaussian beams, parametrized by $(\sigma_A - \sigma_B)/\sigma$ for detectors (A, B) with mean beam width σ . Dipole mismatch (or beam squint) is due to differential pointing of two circular Gaussian beams and is parametrized as $(r_A - r_B)/\sigma$. Quadrupole mismatch (or beam squash) is due to differential ellipticity of the two beams, requiring relative asymmetry and rotation of the the major axes of the two beams to produce leakage. It's parametrized as $(e_A - e_B)/2$, where $e = (\sigma_a - \sigma_b)/(\sigma_a + \sigma_b)$ is evaluated for each detector (A, B).

The differenced patterns of these three effects are shown in the first three panels of figure 5.2, along with their combined effect in the last panel. The monopole and quadrupole terms couple to the second spatial derivative of the CMB temperature field and the dipole term couples to the first spatial derivative (temperature gradient). For very small beams the CMB temperature field appears smooth across the beam scale and the contamination is mitigated. However, for the $\sim 1^\circ$ beams of BICEP2, the leakage of the bright CMB temperature fields can be substantial. The quadrupole term creates false polarization in the same way the instrument measures it. It can not be removed or averaged down through boresight rotation as the monopole and dipole leakage can. However, if the quadrupole term is stable it will produce only rotationally symmetric polarization contamination and will not produce spurious B-mode polarization [26]. If these effects are stable and precisely measured, the polarization leakage can be removed using a template for the true temperature field on the sky as described in § 5.7.2. Instability or uncertainty in their measurement, however, will create false polarization and limit the certainty in their removal from the polarization maps.

The beam parameters are measured by rastering the moon and sources fixed to the calibration mast on MAPO using the far-field flat mirror. The measurements made with the bright moon and BNS are checked for nonlinearity by repeating the measurements with the lower-amplitude thermal source. The effects of the timestream filtering are deconvolved from the signal and the maps are binned into topocentric coordinates. The beams are fit with two-dimensional elliptical Gaussians

$${}^2\sigma = \frac{FWHM}{\sqrt{8\ln(2)}}$$

to determine the beam parameters given above. The individual beam centers are determined from correlations with WMAP maps as explained in § 5.5. There is some disagreement between the centroids derived from the WMAP correlations and from the sources. The discrepancy is being resolved, but is likely due to a combination of uncertainty in the pointing model using the mirror and correction of parallax error. While there is disagreement between the absolute pair centroids the differential pointing derived from source maps agree well with the results derived from the CMB (§ 5.5.1).

The analysis of the beam parameters is ongoing and most measurements have not yet achieved the precision necessary to assess their systematic impact. The beam width has been precisely measured to $\sigma = 0.221^\circ \pm 0.006^\circ$ or $FWHM = 31.2' \pm 0.8'$. This value is slightly larger than the design $\sigma = 0.201^\circ$ calculated from physical optics propagation software. The differential beam width was $(\sigma_A - \sigma_B)/\sigma < 0.007$ in simulations, which is consistent with the measured value 0.003 ± 0.002 over the array. The beam ellipticity was $e < 0.003$ in simulations and the current measured value is weakly constrained to $e = 0.03 \pm 0.03$. The differential ellipticity was simulated to be $(e_A - e_B)/2 < 0.001$ and has been measured to be 0.012 ± 0.009 .

5.4 polarization characterization

The polarization orientation angle (ψ) and cross-polarization response (ϵ) must be accurately measured for each detector pair. These quantities are used to refer pair-differenced timestreams into measurements of polarization on the sky. Miscalibration of the absolute orientation angle mixes measurements of Q and U and thus E- and B-mode polarization. So, uncertainty in the relative orientation angles must be carefully controlled. Cross-polarization response and imperfect orthogonality ($\psi_A \not\perp \psi_B$) of detectors (A, B) within a pair affects the measurement of Q and U equally, which reduces the polarization efficiency but does not cause mixing, or rotation. Uncertainty in the relative orientation angles ($\Delta(\psi_A - \psi_B)$) and the effective polarization efficiency from cross-polarization response ($\frac{1-\epsilon}{1+\epsilon}$) leads to uncertainty in the polarization power spectra amplitudes $C_\ell^{EE, BB}$.

The orientation angles and cross-polarization response have been measured using several different techniques. The analysis has not been completed, but predictions for the levels of uncertainty are discussed below. The measurements below do not immediately give ψ , which gives the orientation of polarization on the sky. Instead these measurements give polarization orientations with respect to the focal plane that are then referred to the sky. The use of the far-field flat mirror introduces additional complications into the analysis of the orientations. The measurements were performed

with the dielectric sheet calibrator (DSC) created for BICEP1 and the linearly polarized BNS placed on the calibration mast [59]. The basic measurement technique is to introduce linearly polarized light into the cryostat and vary the angle of the incoming polarization relative to the detectors by rotating the boresight of the telescope. The resulting sinusoidal response versus angle is fit for each detector pair, with the amplitude giving ϵ and the phase giving an angle used to derive ψ .

The DSC introduces polarized light to the cryostat using a thin polypropylene film oriented 45° to the boresight axis. The film acts a beam splitter, sending a portion of the beam into an ambient-temperature load on the wall of the DSC and a portion to the cold sky. The DSC is mounted to the boot platform and the boresight is slowly rotated while pointing at zenith. The polarization signal is modulated by ~ 250 mK but the beams remain fixed on the sky under representative loading conditions. Using different film thicknesses, the measurement can be run on either the titanium or aluminum transition. Detector pairs near the optical axis show response in close agreement with a sinusoid, while off-axis detectors have a complicated, but well understood, response. This technique is convenient enough for the winter-over to perform during the austral winter to check for variation of the angles with time. Uncertainty in the orientation angles determined from this measurement is set by the determination of the orientation of the dielectric sheet relative to the focal plane. Preliminary analysis of the data show that they should be able to constrain the absolute and relative orientation angles as well as the precision achieved by BICEP1. The DSC should give an uncertainty on absolute and relative orientations of $\Delta\psi \simeq 0.1^\circ$ and $\Delta(\psi_A - \psi_B) \simeq 0.1^\circ$ rms. For BICEP1 the uncertainty on the absolute angles was increased to a conservative estimate of $\Delta\psi < 0.7^\circ$ rms due to a possible discrepancy between data taken before and after the focal plane was serviced [59]. There should be no change in the absolute orientations in BICEP2 since no mechanical adjustment has been made to the telescope since observations began.

An independent measurement of ϵ has been made using the modulated BNS mounted on the MAPO calibration mast. The BNS is fitted with a rectangular horn and placed in a housing containing a rotating wire grid that illuminates the cryostat with linearly polarized light. The combination of horn and wire grid has been measured to have $< 0.03\%$ cross-polar emission. The orientation of the wire grid can be controlled remotely and be determined to within $< 0.01^\circ$. The boresight angle was selected so rows or column of detectors in the focal plane are projected at a constant elevation. The elevation is finely stepped to fully cover each row of beams and the azimuth is scanned to sweep the entire focal plane across the source. These small elevation rasters are repeated for source polarization orientations spanning 360° at 15° intervals. After each row of beams has been scanned

in the array, the boresight is rotated by 90° — swapping focal plane rows and columns — and the schedule is repeated.

A Gaussian is fit to the demodulated response of the detectors. The peak response versus the polarization orientation of the source is fit with a sine wave as $R = A \sin 2(\varphi) + R_{offset}$, where φ is the orientation angle in focal plane coordinates. The cross-polarization response is the ratio of the minimum to maximum response $\epsilon = (A - R_{offset}) / (A + R_{offset})$. The orientation angles will be used as a cross-check for the angles measured using the DSC. Uncertainty in the orientation angles should be dominated by the uncertainty of the absolute source orientation relative to the focal plane. A preliminary conservative upper limit on the cross-polarization response has been placed at $\epsilon < 0.005$, corresponding to a polarization efficiency of $> 99\%$. This is consistent with the cross-polar leakage being dominated by a 0.3% nearest-neighbor inductive crosstalk described in § 4.1.5.1. Also, preliminary analysis of the rotating polarized source data indicate a rotation of $\sim 0.5^\circ$ in the absolute orientations relative to the design.

The polarization power spectra are proportional to polarization efficiency as $(1 - \epsilon)^2 / (1 + \epsilon)^2$ according to equation 5.1. So, uncertainty in the polarization power spectra scale as the square of the uncertainty in the measurement of the polarization efficiency. Thus, the uncertainty on the cross-polarization response $\Delta\epsilon$ must be ± 0.026 to achieve 10% uncertainty in the polarization power spectra.

The polarization orientation of detectors can be non-orthogonal and can differ from the designed absolute alignment. Uncertainty on the absolute orientations mixes the measurement of the E-mode power spectrum into the B-mode power spectrum $C_\ell^{BB} = \sin^2(2\Delta\psi) C_\ell^{EE}$.³ Assuming a conservative value for the E-mode spectrum at $\ell \sim 100$ of $1 \mu\text{K}^2$, this requires $\Delta\psi < 2.4^\circ$ to prevent contamination in the B-mode spectrum at the benchmark level of $0.007 \mu\text{K}^2$. Uncertainty in the relative polarization orientations only affects the amplitudes of the polarization power spectra as $C_\ell^{EE, BB} \propto \cos^2(2\Delta(\psi_A - \psi_B))$. This requires the uncertainty on the relative orientations to be $\Delta(\psi_A - \psi_B) < 9^\circ$ for a limit of 10% uncertainty on the polarization power spectra.

5.5 pointing

A pointing model is used to refer the raw mount encoder values to positions of each beam on the sky and transform Q and U from the instrument basis to the sky. The boresight pointing is used to take the mount encoder values measured along with the timestreams to the position of the boresight

³See [54] for alternative derivations leading to the same result.

on the sky. The radio pointing takes the boresight pointing and angle and locates the beam of each detector pair on the sky. In the absence of differential pointing beam centers are colocated for detectors in a pair, but if differential pointing is present it is the average of the two beam centers.

Error in the pointing at 1% of the beam width is predicted to contaminate the B-mode power spectrum at a level of $r \sim 0.0001$ [26]. The false B-mode polarization scales as σ^2 so the pointing error is limited to 4.2' for the 31.2' average FWHM beam width. This level of pointing accuracy is achieved by the use of a stable observing platform and mount, routine star observations and monitoring and adjustment of the azimuthal tilt. As described in § 5.5, analysis of the star observations taken nearly every three days shows 0.2' rms residual error in the telescope pointing.

The radio pointing is determined by correlating the pair-sum, or temperature, maps of each detector pair with Q, V, and W-band WMAP maps.⁴ The input WMAP maps are first processed using the synfast Healpix utility to remove the average of the beam transfer function over the differencing assemblies for the band, and are then smoothed to the BICEP2 31.2' beam scale. The analysis pipeline is used to convert the smoothed WMAP maps into detector timestreams by interpolating the trajectories of the ideal radio pointing of the beams. The resulting timestreams are then processed into full coadded maps in the same way the real data are processed into maps, including scan-synchronous subtraction and polynomial filtering.

Pair-sum maps for each optically coupled detector pair are correlated with the WMAP maps to find the offset of the common beam centroid relative to the boresight pointing. The pair-sum maps are translated by 0.05° steps in longitude and latitude over $\sim \pm 1^\circ$ and the correlations are calculated. The offsets that maximize the correlation give the offsets of the common beam center for the pair relative to the ideal values calculated from the focal plane solid model and plate scale. The procedure was repeated for the entire 2010 data set, separating the pair maps by deck angle to determine the rotation and offset of the focal plane and mount axes. The residual scatter on the initial analysis including only the 2010 data is $\sim 2.1'$, which should improve by including additional data and refining the analysis procedure.

5.5.1 differential pointing

The same procedure was applied to individual detector intensity maps in order to determine the differential pointing of the beams in a polarization pair. The results are beam offsets for the A and B channels within a pair relative to the common radio pointing determined above. The differential

⁴WMAP data products are available at: <http://lambda.gsfc.nasa.gov>.

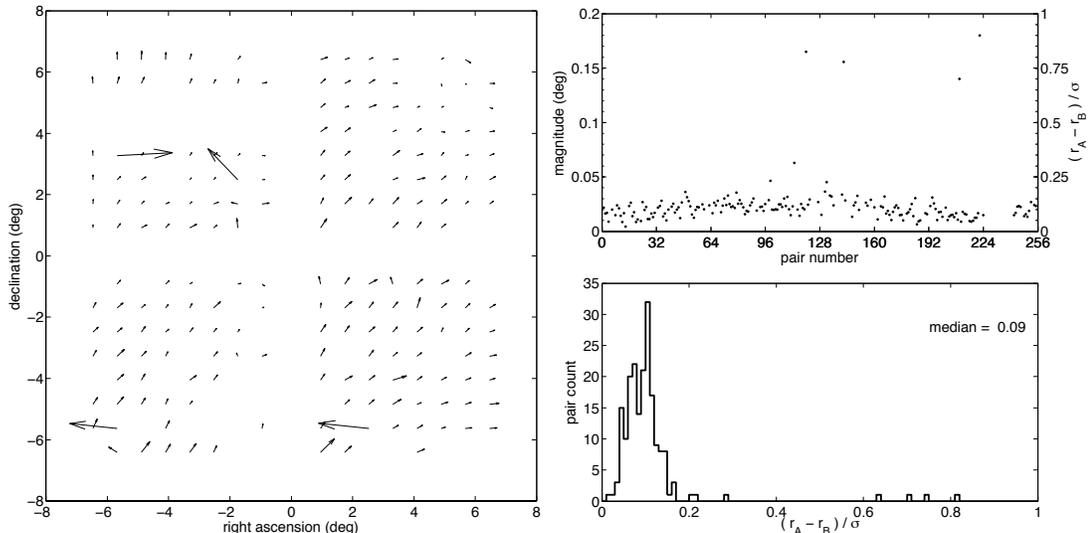


Figure 5.3: Measured differential pointing. Left: The direction and magnitude of the A-B offset of the beams is shown. The magnitude has been exaggerated by a factor of 10 for clarity. The arrows originate from the common centroid of the beams. The focal plane has the same orientation as in figure 2.6. The direction of the offset does not rotate with the tile orientation. Right top: The magnitude (in degrees) or fractional difference of the offset with pair number. Right bottom: The distribution of the fractional difference. The median magnitude of the differential pointing is 0.02° or $1.2'$. This corresponds to a median $(r_A - r_B)/\sigma = 0.093$ assuming the mean beam width of the array $\sigma = 0.22^\circ$.

pointing has been precisely measured at $1.2' \pm 0.4'$. This corresponds to $(r_A - r_B)/\sigma = 9.3\%$ with σ set to the average beam width of the focal plane. These results agree with the analysis of moon rasters, with a median difference in magnitude of $0.2'$ and no systematic difference in angle. Since they are differential measurements, the agreement is much better than the discrepancy between the fit for the beam centroids.

This is by far the largest beam mismatch effect in the BICEP2 focal plane and will require the mitigation technique described in § 5.7.2. The differential pointing in BICEP1 was reported at a mean level of $(r_1 - r_2)/\sigma = 1.3\%$ for detectors (1, 2) within a PSB pair [59]. The individual beam centroids have been measured to high enough precision to use as input in the regression analysis and simulations to assess the contamination of the B-mode power spectrum. The differential pointing has an impact not only on polarization contamination through dipole leakage, but also on the calibration of relative gains and on the common-mode rejection of atmospheric fluctuations.

5.6 relative detector gains

The timestreams of detectors within a pair are differenced to measure polarization. The DC and frequency-dependent gains of the detectors and filters must be deconvolved from the timestreams before differencing. They must be well-characterized to do this without introducing contamination in the polarization measurement. The gains are corrected by first removing the frequency-dependent temporal transfer function from the raw detector timestreams. The DC gain differences are then corrected using the response of the detector to a change in loading during an elevation nod (§ 3.7.1). The gains are normalized by the response of the detectors to atmospheric emission, which differs from the spectrum of the CMB anisotropy. This method will properly reject the unpolarized emission of the atmosphere, but can lead to imperfect rejection of the CMB anisotropy if there are discrepancies between the transmission spectra of detectors in a polarization pair. The significant pointing offset of detectors within a pair also leads to gain mismatch as described below.

Gain error acts in a similar way as the monopole beam term with sensitivity to the temperature field and its second spatial derivative. The contamination of the polarized power spectra scales with the gain error g as $C_\ell^{EE, BB} \propto g^2$. The level of false polarization contamination is sensitive to the distribution of relative gain errors across the array as well as the scan strategy. For stationary gain error observing at several boresight angles averages down the contamination since the sign of the leakage is flipped under a 180° boresight rotation. Larger errors at the periphery of the array lead to larger contamination due to the reduced overlapping coverage. To accurately quantify the level of polarization contamination the residual relative gain errors are measured from the coadded maps. The measured distribution of the errors — or best estimate of the measured gain uncertainties — is then used in simulations to assess the impact.

5.6.1 temporal transfer functions

The temporal transfer functions are the amplitude and phase of a detector’s response to signals at different temporal frequencies. As the telescope is scanned at $2.8^\circ/\text{s}$ the spatial scales of interest ($\ell = 30 - 300$) are projected to a frequency range of 0.1–1 Hz. The transfer function of each channel must be removed, or deconvolved, to recover the true signal before the detectors are differenced. Since the amplitude of the relative gains are normalized by the gains derived from the ~ 21 s elevation nods, the transfer functions must be measured to at least 0.05 Hz. Any low-frequency roll-off of the transfer functions between the science band and the DC gain pinned by the elevation nods must be accounted for. In BICEP1 significant low-frequency roll-off occurred in a number of the

polarization-sensitive bolometers (PSBs) [59][61]. The longer duration of the BICEP1 elevation nods (~ 50 s) meant the transfer functions needed to be measured down to at least 0.02 Hz to account for this, which required a factor of 2.5 increase in measurement time.

Due to strong electrothermal feedback, the bolometers used in BICEP2 are extremely fast. Based on the response of the detectors to small signal TES bias square waves, the median detector time constant is ~ 0.5 ms. Assuming the detector transfer functions are dominated by electrothermal feedback and not the response of the antenna, the inherent transfer functions of the detectors should have a bandwidth of ~ 300 Hz. The bandwidth of the FLL servo of the multiplexer is expected to be ~ 1 kHz. The combined bandwidth of the detectors and multiplexer should have only a small impact on the temporal transfer functions within the science band.

The decimation filters applied by the MCE and GCP have the largest effect on the detector timestreams. Currently, only the filter response is deconvolved from the timestreams during initial data reduction (see § 3.8). In the future the measured transfer functions of each readout channel will be incorporated in the deconvolution. The time-domain kernel used for deconvolution is calculated directly from the GCP and MCE settings used during data acquisition. These settings can differ for data acquired during CMB observations and calibration measurements. Beam mapping data must be deconvolved to account for the group delay of the MCE filter, which causes a one sample lag between the detector timestreams and pointing information. This creates a two sample shift between the left- and right-going scans, elongating the beams in the azimuth direction unless properly accounted for.

Measurement and analysis of the transfer functions has been adapted from the techniques used by BICEP1 [62][59]. The BNS was used to load the detectors with a square wave signal. The switching time of the BNS was only a few nanoseconds and does not limit the bandwidth of the measurement. The detectors and multiplexer were configured in the standard observation mode. The source was placed on a short mast that was rigidly fixed to the boot platform and out of the line of sight of the detectors to avoid additional loading. The signal was diffracted into the cryostat using a thin, low-emissivity foam slab embedded with circular washers. The mast allowed the elevation to be changed between zenith and $\sim 55^\circ$ at different boresight angles to test the effects of changes in loading on the transfer functions. The source amplitude was varied to test the effects of nonlinearity in the measurement. Finally, the transfer functions were measured with the different detector bias and tuning configurations with which the observing data have been acquired. The use of a broadband noise source rather than a narrow-band Gunn oscillator provides immunity from

near-field modulation of the detector response during the on-state of the measurement.

The data are analyzed by taking the ratio of the Fourier transforms of the reference square wave and detector response. The square wave modulation produces signal at the odd harmonics of the fundamental, so those frequencies are extracted from the ratio of the Fourier transforms. Preliminary analysis of the data confirm that the contributions of the detectors and FLL servo to the transfer functions can be ignored. Further analysis of the measured transfer functions will indicate the contribution of frequency-dependent gains to the uncertainty on the relative gains.

5.6.2 relative responsivities

In addition to correcting the frequency-dependent gain, the timestreams are normalized by the DC gain of the detector before differencing. The gains are derived from a pair of elevation nods acquired for every 50-minute scan set, as described in § 3.7.1. The gains are normalized across the focal plane to the median gain of the well-behaved channels. The correction of the relative gains within a detector pair allow them to be differenced, while the normalization across the focal plane allows all the timestreams to be coadded during mapmaking. The timestreams are coadded into maps in multiplexer readout units (ADU). An absolute calibration is measured from the cross-correlation of the final coadded maps with WMAP maps processed with the BICEP2 pipeline, as described in § 5.6.4. This absolute gain is used to scale the coadded map to CMB temperature units.

The relative gains are very stable over the 50-minute scan sets they bracket. This is quantified by the fractional change in the gain ratio derived from the leading and trailing elevation nods over the scan set, $\frac{((\frac{g_A}{g_B})_{trail} - (\frac{g_A}{g_B})_{lead})}{((\frac{g_A}{g_B})_{trail} + (\frac{g_A}{g_B})_{lead})/2}$. This quantity was measured for all the scan sets acquired in the 2010 observing season and found to be within $\pm 0.6\%$, including periods of poor weather.

There is some evidence of a systematic variation with loading. At the very top of the transition for $R_{\%} \gtrsim 90\%$ the responsivity begins to decrease. This is consistent with a lower α_I at the top of the transition. For fixed P_J , G_c and T_c then the loop gain is $\mathcal{L}_I \propto \alpha_I$. The lower α gives rise to lower \mathcal{L}_I , the ETF weakens and the responsivity can no longer be approximated by the reciprocal of the TES bias voltage, $s_{I,dc} \neq -\frac{1}{V_{tes}}$. The condition $R_{tes} \gg R_{sh}$ is still met and equation 37 of reference [30] reduces to the DC responsivity $s_{I,dc} = -\frac{\tau_{el} R_{tes} \mathcal{L}_I}{V_{tes} L_{tot}} = \frac{\tau_{el} V_{tes} \alpha_I}{L_{tot} G_c T_c} \propto \alpha_I$. The substitution of $\mathcal{L}_I = \frac{P_J \alpha}{G_c T_c}$ was made to remove R_{tes} , which varies in the transition. As α_I decreases high in the transition, so does the DC power-to-current responsivity. This occurs during periods of exceedingly poor weather so those data are generally removed by weather cuts. An additional cut on the fractional resistance can be implemented as well.

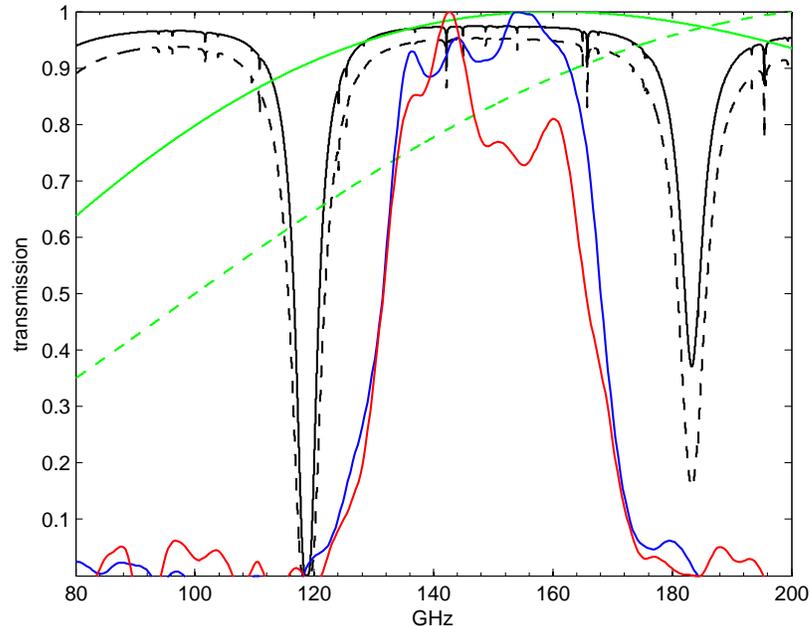


Figure 5.4: Measured spectral response for a pair (red and blue) of detectors. The spectra have been averaged over several measurements and normalized to unity. The simulated atmospheric transmission at the South Pole during winter at zenith (black) and at 57.5° elevation (black dashed) is shown. The spectrum of CMB (green) and its temperature derivative (green dashed) are also shown.

5.6.3 spectral response

Preliminary measurements of the end-to-end spectral response $F(\nu)$ of the detectors was made prior to deployment and in the austral summer of 2011 at the South Pole. This measurement includes the effects of the optics chain, low-pass edge and the band-defining microstrip filter. Spectra were measured using a Martin-Puplett Fourier transform spectrometer (FTS) with 0.3 GHz resolution. At the South Pole, beams were individually coupled to the FTS input port using the far-field flat mirror and a flat mirror placed at the edge of the ground shield. The output port was chopped between ambient temperature and microwave absorber cooled to 77 K with liquid nitrogen. Boresight rotation and translation of the FTS allowed for $\sim 60\%$ of the pairs to be measured. Characteristic spectra for detectors within a pair are shown in figure 5.4 for a preliminary analysis of the South Pole data. The spectra have been averaged over several measurements and peak normalized. The atmospheric transmission at the South Pole was simulated for typical winter observing conditions using the AM atmospheric modeling program.⁵ The atmospheric transmission is shown for zenith and the $\sim 57.5^\circ$ elevation of the CMB field. The line-of-sight air mass increases as the telescope is tilted down, which

⁵www.cfa.harvard.edu/~spaine/am/

reduces the transmission as shown. The spectrum of the CMB along with its temperature derivative are also shown, derived from the Planck function for a blackbody at 2.725 K.

The measured response $F(\nu)$ is not used in equation 5.1 to determine polarization from the detector response. Instead the spectra are measured to understand residual relative gain mismatch in the detector pairs. As shown in figure 5.4, the spectral shape of the atmospheric emission at the South Pole differs from the spectrum of the temperature anisotropy of the CMB. Since the relative gains are measured from atmospheric emission the unpolarized signal from the atmosphere is well-rejected by timestream differencing. If detectors within a pair have large enough discrepancies in their relative passbands then the CMB temperature fluctuations will not be properly rejected. Given high-quality spectra for detectors in a pair the amount of monopole leakage of the CMB anisotropy observed through the atmosphere can be estimated [6]. This can be compared to the monopole leakage expected for residual relative gain errors measured from the final maps or possibly used to mitigate the contamination.

Measurement of the spectra prior to deployment informed a decision to increase the band centers and revealed excess high-frequency response of the detectors. The passband center was increased from 145 GHz to 148 GHz to bring the lower edge of the pass band away from an oxygen line at 118 GHz. Overlap with this line increases loading of the detectors and reduces sensitivity and margin on saturation power. A water line at 183 GHz is far more problematic since the water is poorly mixed in the atmosphere. When the band centers were increased great care was taken to avoid overlap of the spectral response with the water line. As shown in figure 5.4, the spectra have no significant overlap with either line.

A high-frequency spectral leak was found during predeployment engineering runs. The response of the detectors to a modulated broadband thermal source was tested using thick-grille low-pass filters of various cutoff frequencies. The leaks are attributed to direct coupling of radiation to the detector islands that was partially mitigated by a redesign of the detector island. The remaining blue leak is filtered using a 8.3 cm^{-1} low-pass edge filter (§ 2.3.3).

5.6.4 absolute calibration

An absolute calibration to refer multiplexer readout units (ADU) to CMB temperature units (μK_{CMB}) has been made from cross-correlation with WMAP maps. The WMAP Q, V and W-band maps were first processed in synfast to remove the beam transfer function from the maps. The transfer function was taken from the average over the differencing assemblies in the band. The maps were

re-convolved with the beam transfer function for a $31.2'$ beam. The BICEP2 pipeline was used to create timestreams for each detector pair from the trajectory of its measured centroid across the WMAP map and its spatial derivatives. The timestreams were then filtered and binned into maps in the same way that real and simulated BICEP2 timestreams are processed into maps. Final coadded maps were made for the 2010 observing season and for the WMAP Q, V and W bands.

A Fourier-based cross-correlation analysis was then performed between the 2010 map and different combinations of the WMAP maps. The absolute calibration versus multipole moment was derived from the ratio of the cross-spectra from a flat-sky Fourier transform: $g_\ell = \frac{\text{reference} \times \text{calibration}}{\text{reference} \times \text{BICEP2}}$. The *reference* and *calibration* maps were taken from the set of maps for the different WMAP bands. Two different WMAP maps were used to prevent noise bias in the measurement. The gain was flat across the multipole range of interest $\ell=30\text{--}300$. A single gain is taken from the average across the multipole range. The gain is derived for all permutations of the calibration and reference maps to find the final absolute calibration. A gain of $2990 \mu\text{K}_{\text{CMB}}/\text{ADU}$ has been derived for the coadded map from the 2010 observing season. The standard deviation of the different permutations is $< 1\%$.

5.6.5 leakage assessment

An estimate of the monopole leakage from relative gain errors can be made from the maps. Relative gain error can be caused by spectral mismatch or transfer function miscalibration but can also be due to differential pointing. During an elevation nod the displaced beams in a detector pair see slightly different changes in air mass but are regressed against the elevation motion of the average centroid. Gain errors created in this manner depend on the boresight angle. In addition, displaced beams would sample slightly different atmospheric signals depending on the spatial scale of the atmospheric fluctuations. A campaign to measure the monopole leakage from maps is underway and the final results are not yet available.

To measure the effect individual channel intensity maps are cross-correlated with WMAP maps. These maps are coadded for each detector over the 2010 data. The channel maps have adequate signal to noise despite the lack of pair-differencing that removes common-mode fluctuations. The detector timestreams are deconvolved and filtered and binned in equatorial coordinates using the individual radio pointing of the detector, which includes its pointing offset. The channel maps are then cross-correlated in multipole space with WMAP maps processed with the pipeline. This is the same technique used to determine the absolute calibration for the full map coadded over the

Table 5.1: Systematic errors that could produce false B-mode polarization

property	parameter	benchmark	measured	notes
relative gain uncertainty	$\Delta(g_A/g_B)/(g_A/g_B)$	0.9%	–	BICEP1 benchmark, further analysis required for BICEP2
pol. orientation uncertainty	$\Delta\psi$	2.4°	0.1°	anticipated uncertainty based on BICEP1 experience
telescope pointing uncertainty	Δb	$4.2'$	$0.6'$	fit residual in optical star pointing calibration
differential pointing	$(r_A - r_B)/\sigma$	2.9%	9.3%	average, repeatedly characterized to 0.3% precision
differential beam size	$(\sigma_A - \sigma_B)/\sigma$	3.6%	–	BICEP1 benchmark, further analysis required for BICEP2
differential ellipticity	$(\epsilon_A - \epsilon_B)/2$	1.5%	–	BICEP1 benchmark, further analysis required for BICEP2

Table 5.2: Calibration uncertainties affecting the power spectrum amplitudes only

property	parameter	benchmark	measured	notes
absolute gain	$\Delta(g)/g$	5%	1%	variation between different WMAP maps
cross-polarization response	$\Delta\epsilon$	0.026	–	benchmark set by calculation not sim
relative polarization orientation	$\Delta(\psi_A - \psi_B)$	9°	0.1°	anticipated uncertainty based on BICEP1 experience

2010 data and detectors (§ 5.6.4). Instead of giving a single calibration number, this yields per-detector absolute calibrations. The ratio of the absolute calibration of detectors in a pair gives an estimate of the expected level of monopole leakage in the final map, which scales as g_A/g_B . Since the monopole leakage is averaged down by observing at different boresight angles, the maps are created by averaging over the angles to determine the final residual error in the maps. In addition, coadded maps are made for each of the four angles to find evidence of systematic relgain error caused by the differential pointing.

5.7 systematic impact

The impact of the uncertainty in the measurement of the above quantities can be divided into two categories. There are systematic errors that can produce false B-mode polarization as shown in table 5.1, or quantities whose uncertainty only affects the scaling of the polarization power spectra as shown in table 5.2. The benchmarks given in table 5.1 can be scaled to $r = 0.01$ by dividing them by $\sqrt{10}$.

For some of the systematic errors, the level of false polarization produced can be calculated analytically. For example, the contamination calculated for polarization orientation and telescope pointing uncertainty was given in the relevant sections above. The BICEP1 team confirmed the conservative analytic requirement that the telescope pointing uncertainty be $< 1\%$ of the beam width σ . Given the same mount and similar scan strategy this result should be applicable to BICEP2.

The effect of the systematic errors from beam asymmetries and relative gain error is sensitive to the distribution of those errors across the focal plane and the scan strategy. Simulations are

run through the pipeline to properly assess their impact on the polarization power spectra. The best estimate of the measured uncertainties are used in the simulations as explained below. The scaling of the contamination, though, can be determined analytically. The contamination scales as the square of the differential pointing, beam size and ellipticity [54]. The contamination from relative gain errors scales as g^2 . These analytical scaling relations were verified in simulations by the BICEP1 team [59]. The magnitude of these effects can be scaled to match different requirements on the power spectra contamination.

Precision measurements don't yet exist for some of the systematics listed in table 5.1. Those missing are the differential beam size, ellipticity, polarization uncertainty and relative gain uncertainty. The impact of the polarization orientation uncertainty can be calculated analytically and is independent of the scan strategy. The other quantities are dependent on the scan strategy and focal plane distribution and require simulation. They require a precision measurement or an upper limit so that simulations can be run to produce a benchmark. For now the benchmark for the differential beam size and ellipticity, and the relative gain uncertainty are approximated by the upper limits assigned from the BICEP1 simulations for $r = 0.1$ [59]. Due to the sensitivity to focal plane distribution and scan strategy, those can only be taken as rough estimates for the corresponding benchmarks for BICEP2.

The calibration uncertainties that only impact the amplitudes of the power spectra do not require simulations. The requirement on the power spectra amplitude uncertainty has been placed at 10%. The absolute gain calibration g is used to scale the maps to temperature units. The power spectra amplitudes scale as g^2 , so the absolute gain uncertainty must be 5% to limit the amplitude uncertainty to 10%. A conservative upper limit on the absolute gain uncertainty is $\Delta g/g < 1\%$ based on the scatter between cross-correlation with the various WMAP bands. The efficiency of the polarization measurement is set by the orthogonality of the detector polarization sensitivity axes and the level of cross-polarization response. The scaling of the spectra amplitudes with those quantities was described in § 5.4 and is shown in table 5.2. Only a preliminary measurement of the relative polarization orientations has been made. However, based on the experience of BICEP1 with the DSC, the uncertainty on that measurement should be 0.1° . This is well within the 9% limit on uncertainty required for 10% spectra amplitude uncertainty.

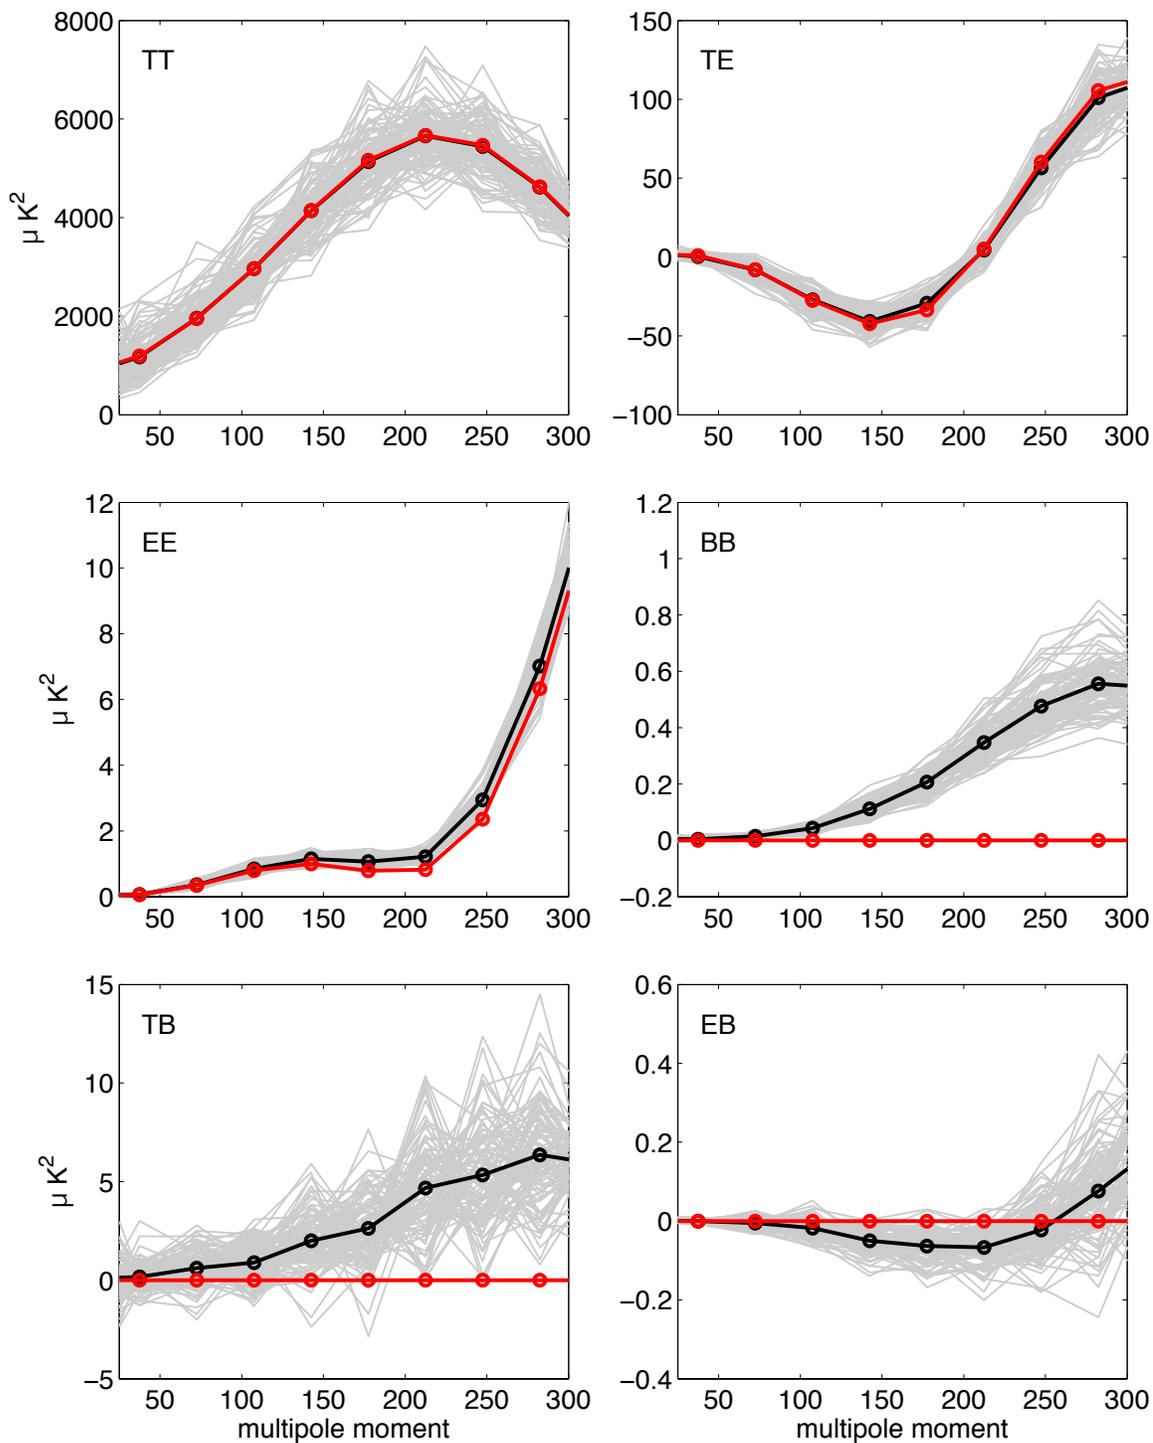


Figure 5.5: Simulated power spectra contamination from the measured differential pointing error. The gray traces are the spectra for the ensemble of 100 signal-only simulations. The black line is the mean of the 100 realizations, with a circle at the measured band powers. The red line is the expectation from the fiducial model with $r = 0$. The BB contamination rises with multipole moment. The contamination at $\ell \sim 90$ is $\frac{\ell(\ell+1)}{2\pi} C_\ell^{BB} = 0.04 \mu K^2$ in the mean and $0.07 \mu K^2$ for the worst fluctuation.

5.7.1 simulations

The goal of simulations is to assess the impact that measured or upper limits on systematic errors should have on the real power spectra. In the case of an upper limit for a systematic error a random distribution of the parameter with the measured rms over the array is used. In the case of a detected systematic error the best estimate of the error and measured distribution in the array is used. The value for the error used for each simulation realization is perturbed about the measured quantity by the estimated measurement uncertainty. In some cases the simulation results can be used to inform cuts of channels contributing the most to the contamination. The channel cuts are implemented in the simulated data and the improvement is gauged from the simulated power spectra.

The general strategy for simulating systematic error is as follows. An ensemble of realizations of the polarized CMB sky are created from model power spectra. For each realization of the signal-only simulation the map is read in and the trajectory of the detectors in the real timestreams are used to sample the input map. Depending on the systematic being explored, a perturbation is introduced to the timestream when it is sampling the map, such as a gain fluctuation or a pointing offset. The simulated timestreams are then filtered and coadded into maps in the same way the real timestreams would be processed into maps. The process is repeated for another realization, using a different input map and a different perturbation is introduced during the creation of the timestream. The power spectra are then taken of the ensemble of signal-only maps. The fluctuations in the power spectra over the ensemble represent the expected statistical spread due to the uncertainty in the measurement. An additional baseline set of 100 signal-only maps are created as well. These include no systematic effects and are used to calculate the transfer function of the filtering in the pipeline and to remove leakage into BB from the power spectrum estimation.

The differential pointing was measured to be $1.2' \pm 0.4'$ averaged across the focal plane. To simulate its effect the best estimate for the offset of each detector was used with additional scatter for the measurement uncertainty. To expedite the creation of the signal-only simulations only a subset of the data was used. The subset included eight complete phases of 10 scan sets each covering the four boresight angles twice. This ensured the scan strategy used in the simulation was representative of the boresight and map coverage expected for the entire data set. The pointing offset for each channel was included in the trajectory used to sample the input map. For each 8-phase realization an additional perturbation to that offset was introduced so that the rms scatter over the ensemble would be $0.4'$. When the timestreams were binned into maps the common centroid of the pair was used in the same way the real timestreams (that include the offset as well) are

binned. This process was repeated on the 8-phase subset of data to create an ensemble of 100 signal-only simulations. The outliers of the distribution in figure 5.3 with $\frac{(r_A - r_B)}{\sigma} > 0.2$ were not included in the coadded maps. The results of the simulations are shown in figure 5.5. The differential pointing contributes significant false power to BB that rises with ℓ . The mean of the 100-realization ensemble at $\ell \sim 90$ is $\frac{\ell(\ell+1)}{2\pi} C_\ell^{BB} \simeq 0.04 \mu\text{K}^2$, and the worst fluctuation is at a level of $0.07 \mu\text{K}^2$. The contamination of the power spectra scales as the square of the differential pointing offset. The necessary reduction in the differential pointing error can be determined from the desire to have contamination $\frac{\ell(\ell+1)}{2\pi} C_\ell^{BB} < 0.007 \mu\text{K}^2$ at $\ell \sim 90$ for $r = 0.1$. For the mean contamination level of $0.04 \mu\text{K}^2$ the required pointing offset is $\frac{(r_A - r_B)}{\sigma} < 0.039$, and for the worst fluctuation it is $\frac{(r_A - r_B)}{\sigma} < 0.029$. The requirement derived from the worst case is included in table 5.1. This can be compared to the benchmark set by BICEP1 of differential pointing $< 1.9\%$ needed for $r = 0.1$. The requirement for BICEP2 is more relaxed due to the improved averaging given by the focal plane.

5.7.2 mitigation

Despite the high-precision measurement of the differential pointing, this measurement can not be immediately used in equation 5.1 to remove the systematic effect. However, the monopole, dipole and quadrupole leakage can be removed by regressing out a template for the leakage in the final maps. This regression technique is currently being applied to BICEP1 three-year data and has shown that a large fraction of the BB contamination can be removed. The basis functions are similar to the ones shown in figure 5.2. They act as filters such that when they are convolved with the CMB temperature fields they create leakage that matches the spatial derivative they couple to. For example, when one convolves a single differential pointing beam with a WMAP temperature map the filtered map shows a clear sign of a first-order spatial derivative of the input map.

To remove the leakage from maps the differential basis must be measured to sufficient precision that the uncertainty in the regression doesn't create uncertainty in the polarization power spectra in excess of the signal being measured. The beams for the light pairs used in the final analysis are measured to high-precision. They are differenced and the basis functions are fit to the differenced beams. Each of the monopole, dipole and quadrupole moments have orthogonal basis functions to represent the orientation of the polarization sense of the pair with respect to the orientation of the beam mismatch. Leakage templates are then created in the time-domain or map-domain. The fitted basis functions are convolved with WMAP maps and their spatial derivatives to form the templates. The templates are then regressed against the pair-differenced maps or timestreams to determine the

correlation and the amount of leakage to remove from the final maps. The resulting power spectra accrue an additional uncertainty from this technique since the templates are created from noisy WMAP maps.

Chapter 6

Analysis

6.1 pipeline overview

The analysis of the first two years of BICEP1 data was performed using two pipelines [13]. The main pipeline used curved-sky power spectra estimators, while the alternate used flat-sky estimators. There was excellent agreement between the power spectra obtained from the two pipelines. The alternate BICEP1 pipeline was adapted from the pipeline used to analyze the second and third seasons of QuAD data [49]. All of these pipelines follow the MASTER formalism for power spectrum estimation and debiasing [24].

The BICEP2 pipeline has been adapted from the alternate, QuAD-derived BICEP1 pipeline. Based on the success of both experiments, and the close agreement with the results from the main BICEP1 pipeline, the BICEP2 pipeline should be adequate for pursuing a constraint on $r \sim 0.1$. Due to the higher sensitivity of BICEP2, the pipeline has required a few upgrades in order to pursue lower limits on r .

The pipeline begins with the low-level timestream processing explained in § 3.8. The DC and frequency-dependent gains of the timestreams are corrected. Then detectors within a polarization pair are summed to measure total intensity (T), and differenced to measure polarization (Q and U). The boresight angle is rotated every 72 sidereal-hours by 45° to exchange measurements of Q and U on the sky. The pair-sum and pair-difference timestreams are filtered and binned into maps using the common pointing of the pair as derived from the telescope pointing and measured pair offsets. The maps are then processed into angular power spectra using two-dimensional Fourier transforms under the flat-sky approximation. Monte Carlo simulations generate an ensemble of realizations of signal-only and noise-only timestreams, which are processed into maps in the same manner as the real data. The signal-only and noise-only maps are added together to create a corresponding

ensemble of signal-plus-noise simulations. The signal, noise and signal-plus-noise simulations are then processed into angular power spectra.

The mean of the ensemble of noise-only power spectra are then used to remove the noise bias from the real and signal-plus-noise power spectra. The signal-only power spectra are used to empirically determine the transfer function the instrument beam and timestream filtering imparts on the maps. In addition, these simulations are used to determine the $EE \rightarrow BB$ leakage caused by the finite coverage and pixelization of the maps. The mean of the BB spectrum in the ensemble of signal-only power spectra is then removed from the real and signal-plus-noise BB power spectra. Finally, the power spectra from the noise-only, signal-only and signal-plus-noise simulations and real data are corrected for the filter and beam suppression factor. Additional steps, discussed below, are then used to assess the quality of the real data.

6.2 cuts

A preliminary set of channel cuts has been selected for the analysis presented in this chapter. The data cuts have eliminated channels that contribute excessive leakage in the polarization maps. These cuts mainly removed channels that had high levels of multiplexer crosstalk, unusually high pointing offsets, or relative gain mismatch. No effort was made to differentiate these contributions to the contamination eliminated by the preliminary cuts. Detailed follow up work is currently underway to diagnose the sources of contamination.

The channel cuts were based on contamination in individual pair-difference maps for each detector pair. Two sets of maps were made for data acquired before and after the bias changes made in September 2010. In addition, scan-direction jackknife maps were created for each pair for the two observing configurations. A contamination metric was defined as the ratio of the the variance of each pair-difference map to the variance of its scan-direction jackknife pair-difference map. The variance maps were normalized by the noise in the jackknife maps to ensure that noisy channels were not erroneously cut.¹ Outlying channels in the distribution of the contamination metric were included in the channel cut.

The preliminary channel cuts have been applied to the maps and power spectra presented in the remainder of the chapter. The cuts have reduced the yield from 192 to 171 pairs for data acquired before the bias changes. For the late 2010 data the cuts have reduced the yield from 206 to 196 pairs. Due to the bias optimization there are more working pairs in late 2010 and fewer pairs have been

¹That said, excessively noisy channels are down-weighted during mapmaking and contribute little to the maps.

removed for contamination. The reduction in the contamination is due primarily to the removal of channels that have multiple SQ1s turned on at the same time.

6.3 mapmaking

The goal of mapmaking is to form T, Q, and U signal maps. In addition, variance maps and maps of the integration time are also produced. The variance maps are used to apodize the signal maps before calculating angular power spectra. The integration time is useful for extracting the sensitivity of a given map, as explained in section 4.1.9. The maps are binned in right ascension and declination in pixels that are 0.25° on a side. A “naive” mapmaking approach is adopted, one in which the noise is assumed to be white and the noise covariance matrix is assumed to be diagonal. To this end, the timestreams are filtered to remove correlated $1/f$ noise. The pair-differenced timestreams benefit from further $1/f$ rejection, so this assumption is more valid for the polarization maps. Since the goal of BICEP2 is a measurement of the polarization anisotropy, and not the temperature anisotropy of the CMB, this approach is justified.

The first step is to sum and difference the gain-adjusted timestreams ($d^{A,B}$) of light detectors within a pair and divide the result by two:

$$d_i^\pm = \frac{1}{2}(d_i^A \pm d_i^B) \quad (6.1)$$

where i indicates the sample number. To prevent the introduction of spurious polarization signals in the maps only light channels, in which both detectors are functioning, are used in mapmaking. The timestreams are divided into left- and right-going half scans, with a complete scan set having 106 half scans. The half scans are reduced to the sections where the azimuth motion is linear. Each pair-sum and pair-difference half scan is then filtered by a third-order polynomial to remove the bulk of the atmospheric $1/f$ noise.

The pair-sum and pair-difference left- and right-going half scans are then binned in azimuth. The mean over all of the left- or right-going half scans for a given channel is then subtracted from each left- or right-going half scan for that channel. Both of these filtering operations remove $1/f$ noise but also remove signal from the timestreams. The samples are inverse-variance weighted when accumulated into maps. The weight is the reciprocal of the variance calculated for each pair-sum

and pair-difference timestream across the scan set:

$$w_i = \frac{1}{\sigma_{\pm,i}^2} = \frac{1}{\sigma_{\pm,j}^2} \quad (6.2)$$

The single variance measurement for the scan set j is expanded to all samples i in the scan set.

Temperature maps are formed from the accumulation of all weighted, pair-sum samples i into the appropriate pointing bin as:

$$T = \frac{1}{w} \sum_i w_i d_{+,i} \quad (6.3)$$

where T is a matrix representing the pixels in the temperature map. The weighting is normalized by $w = \sum_i w_i$. The temperature variance map σ_T^2 is given by the weighted sum of the variance samples:

$$\sigma_T^2 = \frac{1}{w^2} \sum_i w_i^2 \sigma_{+,i}^2 \quad (6.4)$$

Each sample is binned in right ascension and declination based on the trajectory of the common centroid of the pair across the sky. This is determined from the boresight angle, measured pair offsets and the boresight pointing determined from the encoder samples and pointing model. In addition, a map of the total integration time for all samples is formed by binning $1/f_{sample} \simeq 0.05$ s for each pair sample i .

The Q and U maps are formed from the pair-difference timestreams. The Q and U variance maps are formed from the variance across the scan set of the pair-differenced timestreams. The product of the pair-differenced half scans and weights are multiplied by the sine and cosine of the detector angle projected on the sky. The detector angle is the measured orientation of the polarization sensitive axes of the pair referred to the sky for the given boresight angle. The relationship between the pair-differenced timestreams and the Q and U maps is given by:

$$\begin{pmatrix} \frac{1}{w} \sum_i w_i \cos(2\alpha_i) d_{-,i} \\ \frac{1}{w} \sum_i w_i \sin(2\alpha_i) d_{-,i} \end{pmatrix} = \begin{pmatrix} \frac{1}{w} \sum_i w_i \cos(2\alpha_i) \cos(2\alpha_i) & \frac{1}{w} \sum_i w_i \cos(2\alpha_i) \sin(2\alpha_i) \\ \frac{1}{w} \sum_i w_i \cos(2\alpha_i) \sin(2\alpha_i) & \frac{1}{w} \sum_i w_i \sin(2\alpha_i) \sin(2\alpha_i) \end{pmatrix} \times \begin{pmatrix} Q \\ U \end{pmatrix} \quad (6.5)$$

where the first equation on the right-hand side is the QU correlation matrix. This matrix is inverted to solve for the Q and U maps from the weighted accumulation of the pair-differenced timestreams with the sine or cosine of the orientation angle for the sample. In order to invert this matrix each pixel must be measured at two orientation angles. These should be separated by 45° for an optimal estimation in the presence of noise. Since observations are conducted at four different

Table 6.1: Comparison of BICEP2 2010 and BICEP1 2006–2007 data quality. The number of pairs for BICEP2 is shown after the cuts made for the early-2010 and late-2010 observing configurations. The number of pairs shown for BICEP1 are the pairs included in analysis split between the 2006 and 2007 seasons. The BICEP2 map depth and integration time has been calculated from the deepest 400 deg² in the 2010 maps. The minimum integration time for the regions is quoted for both BICEP1 and BICEP2.

	number pairs	map depth $\mu\text{K}/\text{deg}^2$	integration time det – s/deg ²	NET_{det} $\mu\text{K}\sqrt{\text{s}}$
BICEP2 2010	171 / 196	0.22	$> 1.7 \cdot 10^6$	386
BICEP1 2006–2007	–	–	$> 6 \cdot 10^5$	–
100 GHz	19 / 22	0.81	–	530
150 GHz	14 / 15	0.64	–	450

boresight angles, each pixel in the central portion of the maps is measured at roughly four orientation angles, spaced by 45°. This is not necessary to invert the matrix, but is useful for integrating down systematic effects. The Q and U variance maps are determined from the inverted QU correlation matrix, along with weighted sums of the products of the sine and cosine terms with the variance of the pair-difference timestreams.

The maps for the full 2010 observing season are shown in figure 6.2. The noise levels in the maps have been measured from scan-direction jackknife maps that remove signal. The deepest 400 deg² — roughly half the map area — of the Q and U maps have rms noise levels per square degree of 0.22 μK . The same map depth has been confirmed from non-jackknifed, noise-only simulated maps. In the same region, the integration time exceeds $1.7 \cdot 10^6$ detector-seconds per square degree. After two years of observations BICEP1 measured map depths of 0.81 μK and 0.64 μK for the 100 and 150 GHz channels, respectively, [13]. Due to the nearly fivefold increase in detectors the integration time is nearly three times higher than the $6 \cdot 10^5$ detector-seconds per square degree reported by BICEP1 for two years. The map-based sensitivity analysis presented in § 4.1.9.2 has been repeated on the full 2010 maps. The mean sensitivity from the Q and U scan-direction jackknife maps is $NET_{det} \simeq 386 \mu\text{K}\sqrt{\text{s}}$.

6.4 power spectra estimation

Angular power spectra are estimated using heuristically-weighted quadratic estimators, also known as pseudo- C_ℓ estimators. The spectra are estimated assuming a flat sky and follow the calculations used for the QuAD analysis [49]. The maps are inverse-variance weighted, which improves the estimation by down-weighting the noisy regions at the edges of the map. The spacing of the observed

band powers is carried over from the BICEP1 choice of nine bins of $\Delta\ell = 35$ covering a range $20 \leq \ell \leq 335$.² The power spectra estimation generates the three auto-correlation spectra (TT , EE , BB) and three cross-correlation spectra (TE , TB , EB).

The flat-sky estimators were demonstrated to be appropriate for the size of the BICEP1 field given the agreement between the curved-sky and flat-sky power spectra. However, this technique leaks temperature into polarization and will eventually limit the search for r in the BICEP2 maps. Eventually the flat-sky power spectrum analysis will need to be replaced with a curved-sky analysis.

6.4.1 pseudo- C_ℓ estimators

The maps (M_X) are first apodized using the reciprocal of the variance maps, $A_X = 1/\sigma_X^2$, where $X = \{T, Q, U\}$. The cross-spectra are apodized using a mask calculated from $A_{XY} = \sqrt{A_X A_Y}$. The resulting masks show no discontinuities or peaks and smoothly apodize the signal maps to zero at the perimeter, where the coverage is lower. The apodization mask is smoothed just in case, though it's generally not necessary. Discrete point sources in the maps are ignored since they contribute negligibly to the power spectra due to the low-resolution, 0.5° beams [13]. A discrete two-dimensional Fourier transform is then taken of the product of the apodization mask and signal mask:

$$m_X = c FT(M_X A_X) \quad (6.6)$$

where c is a constant that removes the effect of the apodization mask. The result is a two-dimensional map of complex Fourier modes m_X . The modes are a function of the coordinates u and v , which are the Fourier conjugate variables to angular distance from the map center. The corresponding multipoles are given by $\ell = 2\pi r$, where $r = \sqrt{u^2 + v^2}$ is the radius from the origin in the Fourier plane. For the standard pseudo- C_ℓ estimation the E and B spectra are formed from Q and U as:

$$m_E = m_Q \cos(2\phi) + m_U \sin(2\phi) \quad (6.7)$$

$$m_B = m_Q \sin(2\phi) - m_U \cos(2\phi) \quad (6.8)$$

where ϕ is the polar angle of each mode m_X with respect to the origin in the Fourier plane, or $\phi = \tan^{-1}(u/v)$.³

²The bins are centered on $\ell = [37.5, 72.5, 107.5, 142.5, 177.5, 212.5, 247.5, 282.5, 317.5]$.

³The sign of the B-modes matches the Healpix convention and has the opposite sign as the reference [49].

The final binned angular power spectra are found by averaging over the quantity:

$$C_b^{XY} = \langle \frac{\ell(\ell+1)}{2\pi} m_X m_Y^* \rangle \quad (6.9)$$

where the complex modes m_X and m_Y are selected from the set of signals T, E, and B. The averaging is performed over annuli in the two-dimensional Fourier plane encircling the origin with multipole moments that fall within the band power bin b . From here on, the pseudo- C_ℓ power spectra estimators all include the prefactor $\frac{\ell(\ell+1)}{2\pi}$, and the superscript XY will be dropped.

6.4.2 pure pseudo- C_ℓ estimators

The finite sky coverage of the map means there are ambiguous E- and B-modes that cannot be differentiated. The bright EE spectrum is leaked into the BB spectrum and must be removed in the mean, as explained in § 6.5.2.⁴ This leakage also adds to the error bars of the BB band powers. The EE variance is leaked into the measurement of the BB variance. The mean of the ensemble can be removed from the real spectrum, but the fluctuation of the ensemble cannot. For a high-sensitivity, small-coverage experiment like BICEP2, this $E \rightarrow B$ mixing will ultimately limit the constraint on r . The theoretical lower bound for this limit is $r = 0.04\text{--}0.05$ [10][55].

The improved sensitivity of the BICEP2 experiment requires the use of an estimator that reduces the mixing of E- and B-modes. An alternative pseudo- C_ℓ method is now being used to determine EE and BB, which eliminates mixing. It does so by projecting out purely B-modes and throwing out any ambiguous modes, which are prevalent around the perimeter of the map. It follows the treatment referred to as “pure pseudo- C_ℓ estimators” in the reference [55]. The two-dimensional Fourier modes are still calculated for E using equation 6.7, but equation 6.8 has been replaced by the calculation given in the reference. This involves a linear sum of combinations of Fourier-transformed products of Q and U with the spatial derivatives of the apodization mask. There are no other changes to the power spectrum estimation techniques described in § 6.4.1.

The pure estimators have shown a reduction in the BB leakage by a factor of four using the standard apodization mask, as shown in figure 6.6. The size of the BB error bars can be further improved by a choice of a more optimal weight mask in the future, as described in the reference [56]. Additional improvements can be made by replacing the “ra-dec” map projection, and by moving to a curved-sky analysis. As shown in figure 6.3, the noise variance is slightly increased relative to the standard estimators for the 2010 power spectra, due to the loss of the ambiguous modes.

⁴There is of course $B \rightarrow E$ leakage as well, but this is negligible for the Λ CDM model.

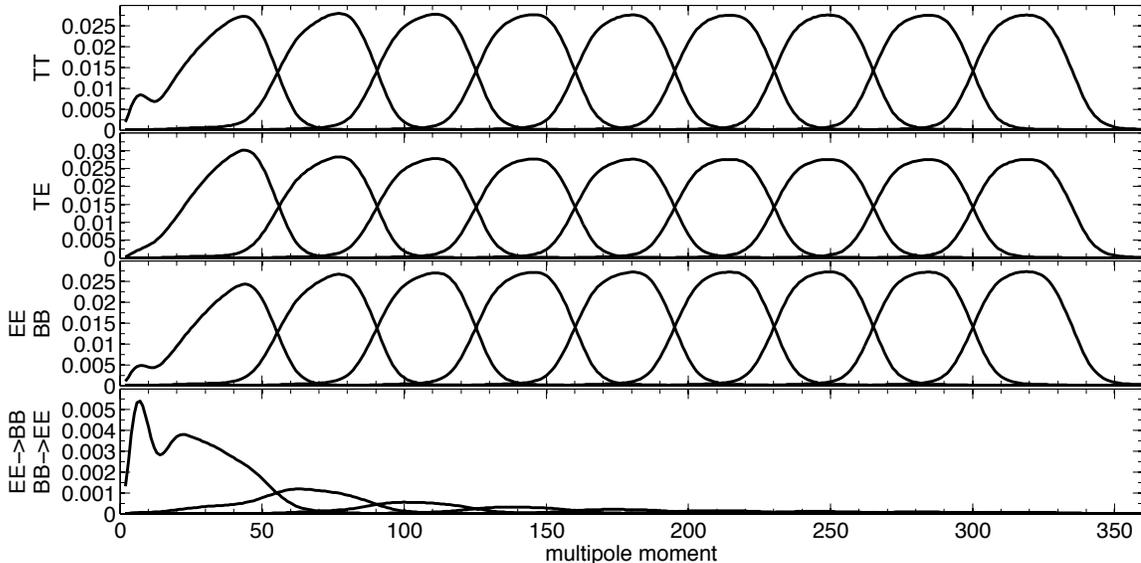


Figure 6.1: The band power window functions calculated from the maps and apodization mask. There is one window function w_ℓ^b for each of the nine band powers b . These window functions are summed with the continuous input model C_ℓ^{mod} to form the expectation values from the model at each band power. The $EE \rightarrow BB$ and $BB \rightarrow EE$ window functions describe the leakage in the maps due to the finite sky coverage and pixelization. These functions are calculated but are not used to debias the leakage from the BB spectra.

This penalty is compensated by the reduction of the leaked signal variance. The noise variance will further decrease with integration time and will ultimately be limited by the signal variance. Reducing the signal variance to the lowest level is the goal in the application of the pure estimators. It has been observed that this pure estimation procedure mixes spatial modes within the map. This does not affect the validity of the estimation of cosmological signals. It does have consequences for the BB power spectrum and constraints on r in the presence of instrument systematics, as described in § 6.8.

6.4.3 band power window functions

The multiplication of the maps by the apodization mask causes correlations between modes in the Fourier domain. The estimated band powers are not just the mean over the $\Delta\ell = 35$ range of the bins, but are sensitive to power at moments belonging to adjacent bins. To determine the extent of this effect the band power window functions are estimated from the real maps and apodization mask [34]. More importantly, the band power window functions give the contribution to the band power of interest from each multipole moment.

The recipe for calculating the band power window functions follows the method given in the

reference [49]. One small modification has been made to accurately characterize the leakage in the larger BICEP2 maps, using the leakage calculation given in the reference [10]. This technique produces six band power window functions for each of the spectra. Each band power window function (w_ℓ^b) has a window for each of the nine band powers (b) that are continuous functions of ℓ .

The functions have been calculated from the real 2010 observing data and apodization mask, and are shown in figure 6.1. Four of the window functions are the response of TT , TE , EE and BB to true sky power. Two functions, $EE \rightarrow BB$ and $BB \rightarrow EE$, describe the leakage of modes due to the finite extent and pixelization of the maps. There are unique window functions for TT and TE . The window functions for EE and BB are the same, as are the functions $EE \rightarrow BB$ and $BB \rightarrow EE$. The leakage functions are not used in analysis. Instead the $EE \rightarrow BB$ leakage is calculated from the ensemble of signal-only simulations as described in § 6.5.2. For the choice of $\Delta\ell = 35$ band powers, adjacent bins are correlated at a level of 17%.

The band power window functions are used to determine the expected discrete band powers from a continuous input model spectrum. The expectation value for a bin b is calculated from w_ℓ^b and the input model C_ℓ^{mod} using:

$$C_b^{exp} = \sum_\ell w_\ell^b C_\ell^{mod} \quad (6.10)$$

for each of the auto- and cross-spectra. The window function is normalized so that $\sum_\ell w_\ell^b = 1$ for each bin b .

6.5 simulations

Monte Carlo simulations are run to remove the effects of noise, filtering, and leakage from the real maps. Signal-only simulations are run to determine the transfer function the timestream filtering and beam impart on the real maps. Noise-only simulations are used to determine the noise bias in the real spectra due to instrument noise. The error bars for the real spectra are determined from the band power covariance matrix estimated from the combination of signal and noise simulations.

The mapmaking is a linear operation. This allows the signal-only and noise-only maps to be generated independently and then added together to create an ensemble of signal-plus-noise maps. The simulation maps are debiased and processed in exactly the same way the real maps are processed. The independence of the signal and noise maps allow the effects of the noise bias and filtering to be measured and removed separately.

Due to the heavy computation time required in generating the simulated timestreams, an en-

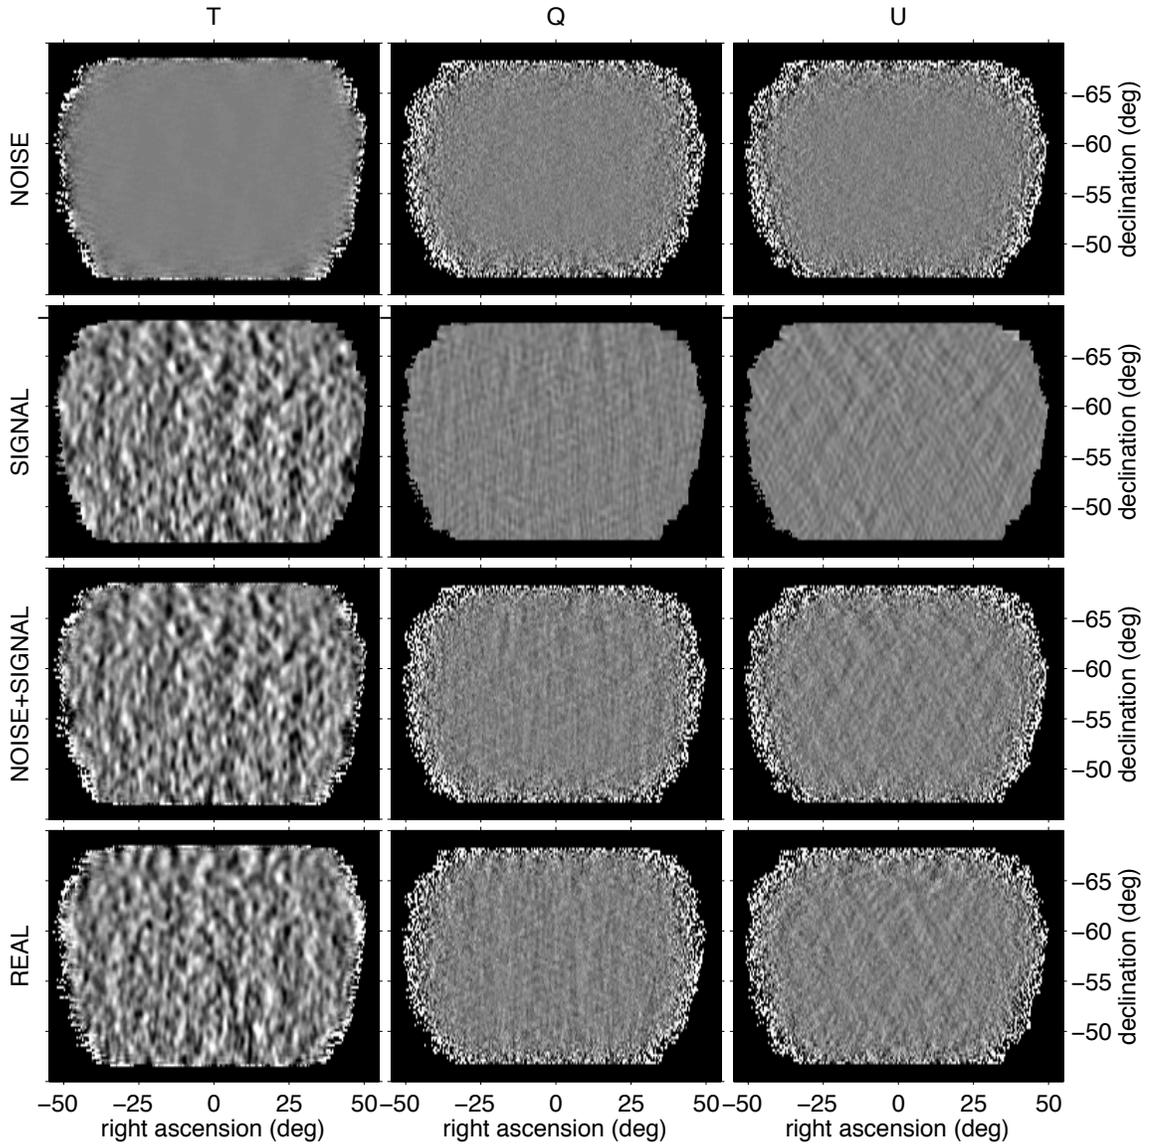


Figure 6.2: Simulated T, Q, and U maps and the real maps for the 2010 observing season. Each of the simulated components of the T, Q, and U maps are shown for a single realization. The noise-only and signal-only maps are created independently from simulated timestreams and then added together to form the signal-plus-noise maps. The signal-plus-noise simulated map faithfully reproduces the structure in the real map.

semble of only 40 realizations has been produced thus far. A recent effort was made to reduce the computation time by making the simulation code “parallel”. Instead of creating a single simulated realization for each real scan set that is read in, the parallel code generates as many as 100 realizations. The parallel simulation code should increase the throughput by a factor of 20.

6.5.1 noise-only simulations

Noise-only simulations are used to remove the noise bias from the real power spectra. The fluctuations in the ensemble of the simulations also contribute to the error bars. The mean of the power spectra of the ensemble of noise-only simulations is removed from the real power spectra. The noise must be accurately modeled in order to properly subtract the noise bias. This is particularly important for the constraint on r from the BB power spectrum, since over-subtraction of the noise bias will bias the result toward a lower 95% confidence limit.

A noise-only simulated timestream is created for each real timestream. A noise model is used to recreate the shape of the PSD of the real timestream as accurately as possible. The noise-only timestreams are pair-summed and pair-differenced, filtered, and binned into maps in the same way as real data. The noise is modeled for the complete scan set, and then the half scans are cut to the portions of linear azimuth motion. Noise-only maps are shown for the 2010 season in figure 6.2 for T, Q, and U. The faint striping in the temperature map is due to the measured $1/f$ noise in the timestreams, and is present in the real T map. It is largely removed by pair-differencing and does not show up in the polarization maps.

The noise-only 2010 maps are processed into the power spectra shown in figure 6.3. The expectation values from the fiducial model with $r = 0.1$ are also shown, in order to illustrate the signal to noise of each of the auto- and cross-spectra. The TT and EE power spectra are signal-dominated, while the BB spectrum is completely noise-dominated for any allowable r . The mean of the ensemble — shown in solid red or blue for the two estimators — is subtracted from the spectra to remove the noise bias. The noise model accounts for correlated noise, which is assumed to be due to the atmosphere. These correlations lead to noise in the cross-spectra, as seen for TB and EB in the figure. These levels are very low, as shown by the mean spectra being consistent with null TB and EB power. As a precaution, the cross-spectra are also debiased by the mean of their corresponding noise-only spectra.

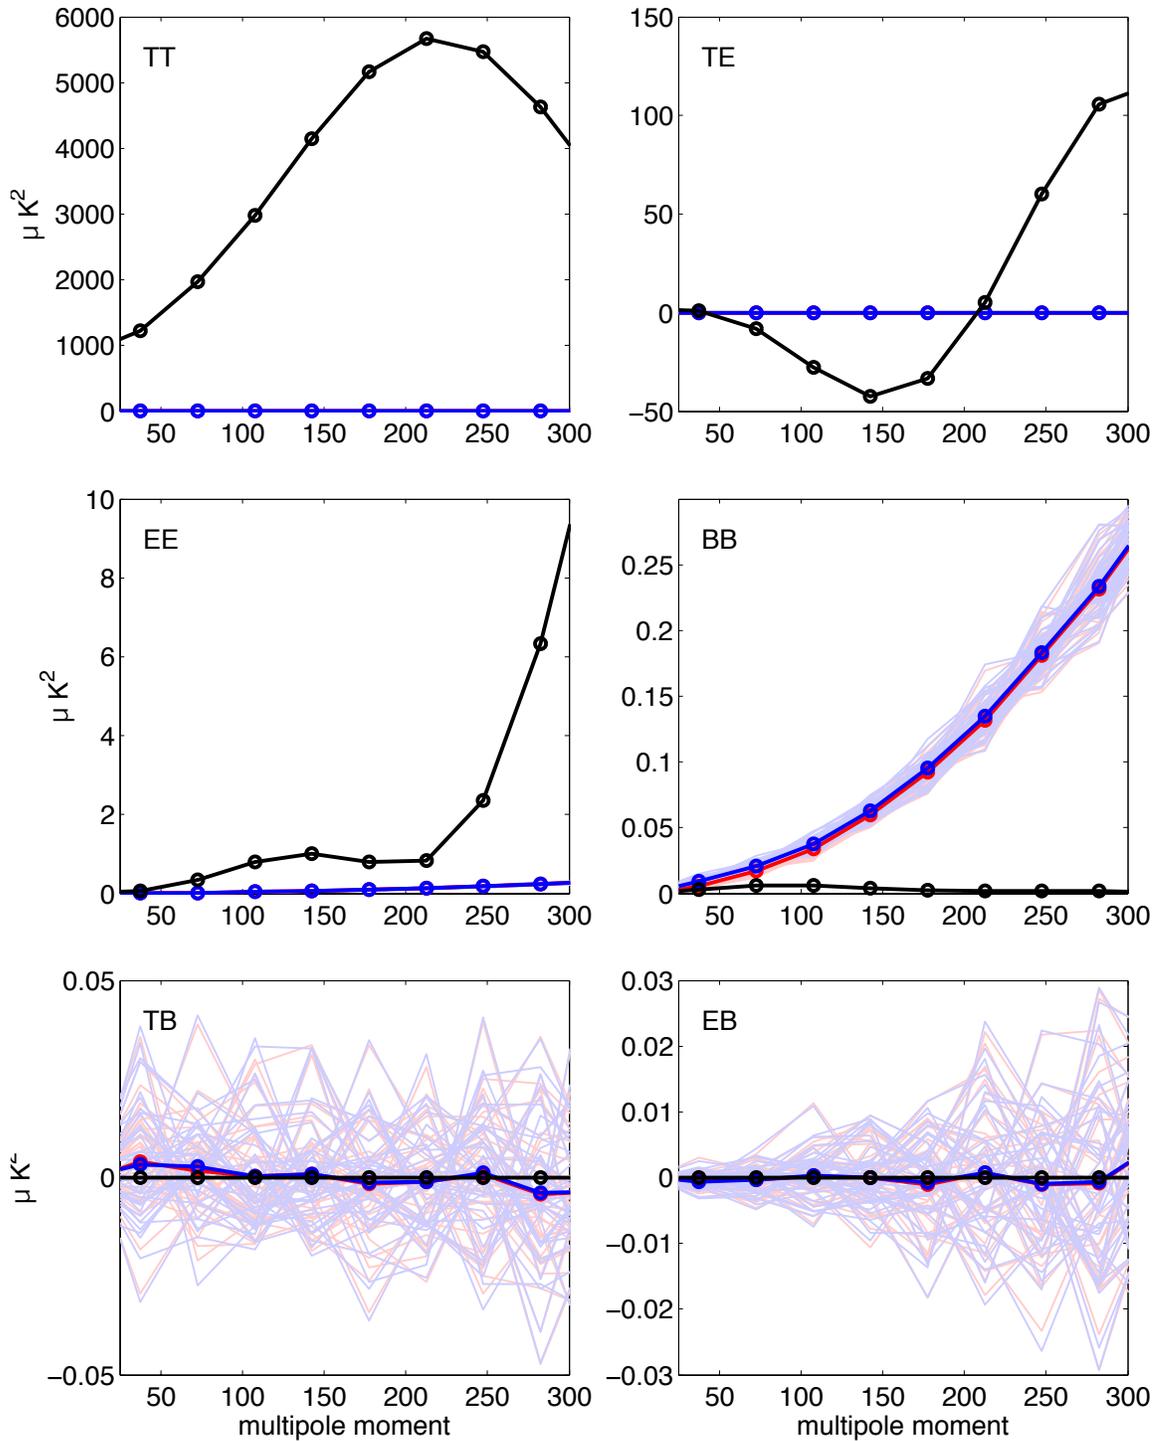


Figure 6.3: Angular power spectra for noise-only simulations. The welter of lightly colored lines represents the ensemble of the simulations, with the mean of the ensemble indicated by the heavy line. The blue lines indicate the results of the pure pseudo- C_ℓ estimator and the red lines indicate the standard estimator. A fiducial model with $r = 0.1$ is shown to compare the signal to noise for each spectrum. While the TT and EE band powers are signal-dominated, the BB band powers are noise-dominated for the 2010 season. The solid line, which indicates the mean of the ensemble, is used to debias the real and signal-plus-noise power spectra.

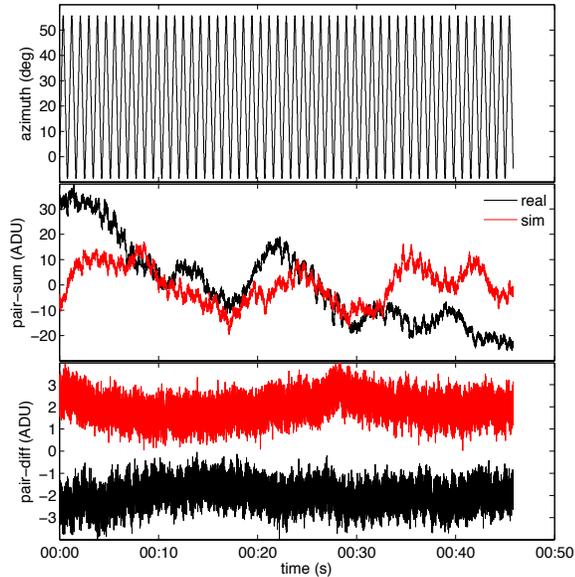


Figure 6.4: A set of real and simulated noise-only timestreams are shown for a single scan set and detector pair. The real and simulated pair-sum and pair-difference timestreams are shown along with the azimuth motion during the scan set. The scan-synchronous response in the real pair-sum timestream is reproduced in the simulated timestream.

6.5.1.1 noise model

The noise model is used to accurately reconstruct the structure of the timestreams in the time and Fourier domains and reproduce any channel-channel correlations. The noise model also inadvertently reconstructs any signal in the timestream when trying to determine the noise. Signal-only and noise-only simulated PSDs were compared to assess the signal to noise expected in the pair-sum and pair-difference timestreams. The pair-sum timestreams show a signal to noise ranging from ~ 3 –10% across the 0.1–1 Hz science band. The pair-difference timestreams are almost completely noise-dominated, with signal to noise of $\sim 0.2\%$ across the science band. The pair-sum timestreams have a large signal component from the CMB anisotropy and atmosphere compared to the noise. The goal of BICEP2 is to measure polarization, so the modeling and removal of the noise bias from the temperature spectra is not critical. There is little concern for this in the polarization maps due to the low levels of signal in the pair-difference timestreams.

Noise is modeled over a complete 50-minute scan set, including the nonlinear portions of the scan that are cut from mapmaking. Modeling over a complete scan set more accurately reproduces the correlations and low-frequency structure in the timestreams. The first step is to read in a real timestream for a scan set. The real timestreams are not filtered, summed or differenced since this step does not occur for the real data until the mapmaking stage. The low-frequency noise for each channel

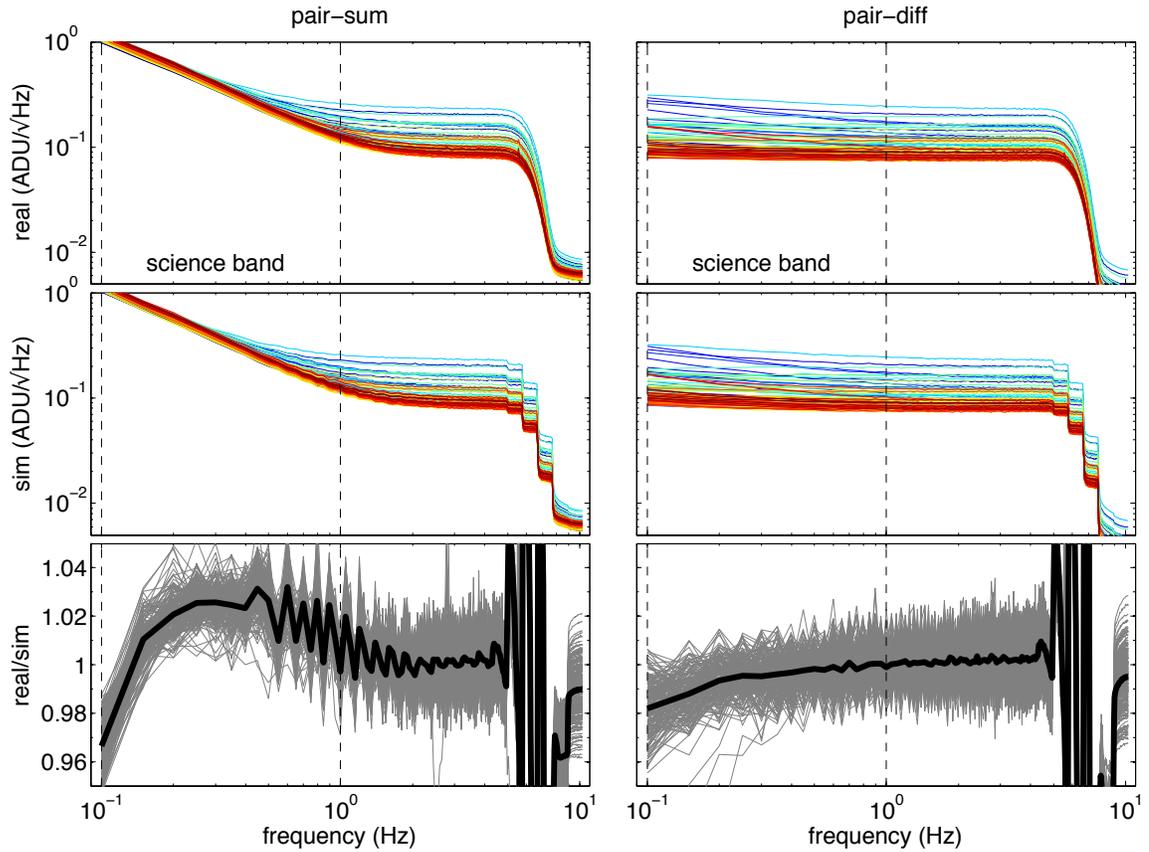


Figure 6.5: Real and simulated PSDs are shown for the pair-sum and pair-difference timestreams. The PSDs have been averaged over a subset of timestreams from 80 scan sets. For clarity, only the PSDs for detector pairs on tiles 3 and 4 have been shown. The 0.1–1 Hz science band is indicated by the vertical dashed lines. The simulated PSDs faithfully reproduce the shapes of the real PSDs within the science band. The blocky features in the simulated PSDs at high frequency are the result of the binning when forming the simulations. To test the noise model the ratio was taken of the real and simulated PSDs for each of the 80 scan sets, indicated by the welter of gray traces. The average over all of the traces is indicated by the solid black line. For the pair-sum PSDs the deviations are $\pm 3\%$ in the science band. The pair-difference PSDs have far less deviation at $< 2\%$ for the 80 scan sets shown.

is modeled by regressing it against a Chebychev polynomial. Only a first-order polynomial is used, so the timestreams are fit to a straight line. A channel-channel covariance matrix is formed from the coefficient for the slope, which is then Cholesky decomposed. The high-frequency noise is modeled from the channel-channel correlations found for the Fourier modes located in 70 logarithmically spaced bins. The real data are Fourier transformed and the channel-channel covariance is found for all modes within a certain bin. The correlations are computed for all bins and then Cholesky decomposed.

These two Cholesky decomposition matrices are then used to create the simulated timestreams from a Gaussian distribution of pseudorandom numbers. An array of random first-order Chebychev polynomial coefficients are generated from the low-frequency Cholesky matrix. The low-frequency component of the timestream is simulated, amounting to an ensemble of straight lines with random but correlated slopes. The high-frequency Fourier modes are mixed between adjacent bins to reproduce the mode-mode correlations in the PSDs. This is then multiplied by a matrix of random numbers for each Fourier-domain sample and channel, and inverse-Fourier transformed to get the high-frequency component of the timestreams. The low- and high-frequency components are then added together.

A real timestream as well as a single realization of a noise-only timestream are shown for a single detector pair and scan set in figure 6.4. These timestreams have been gain-corrected and processed for mapmaking. The scan-synchronous structure in the real timestream is faithfully reproduced by the noise model. To further gauge the accuracy of the noise model, PSDs were created from both the real and noise-only simulated pair-sum and pair-difference timestreams. An example of the real and simulated PSDs is shown in figure 6.4. For clarity, only channels from tiles 3 and 4 with well-matched detector parameters are shown. PSDs were formed for each half-scan within an 8-phase subset and then averaged together. The ratio of the real and simulated PSDs was also determined for each half-scan and averaged. This procedure was performed for the entire 2010 observing season. The ratio for the pair-difference PSDs had less than 1% error across the science band. The error was higher at $\sim \pm 3\%$ for the pair-sum PSDs. Again the adequacy of the pair-sum noise model is less of a concern since polarization power spectra are the primary science goal.

The noise model is sensitive to steps and the stationary state of the noise in the real timestreams. A glitch or step introduces power across a wide range in the Fourier domain. Since the noise-only timestreams reproduce the structure of the real timestream in the Fourier domain and the modeled bins are broad, the glitches will produce increased noise that is broad in the time domain. Thus, steps will create an overestimate of the noise in the simulated timestreams. Steps and glitches must

be properly removed from the timestreams in the low-level processing steps. The steps are generally caused by flux jumping of the output DACs, which induces pickup on other channels.

6.5.2 signal-only simulations

Signal-only simulations serve a number of functions in the analysis of the real power spectra, as well as in the characterization of instrument systematics. They are used to empirically determine the suppression of signal power that occurs from the beam, and filtering of the pipeline. They give the amount of $EE \rightarrow BB$ leakage that must be removed from the BB spectra. The signal-only simulations give the contribution to the band power error bars resulting from the loss of modes due to finite sky coverage. This, along with the leakage of EE variance, determines the contribution of the signal to the total error bars. Signal-only simulations are also used to simulate the measured systematic errors in the instrument (§ 5).

Two types of model power spectra were created in CAMB from the WMAP five-year best-fit parameters. In one model the tensor-to-scalar ratio was set to $r = 0$ to create the standard Λ CDM model. In the other model $r = 0.1$ was used, but the EE power spectra have been nulled. The output of CAMB is a set of C_ℓ fits files for the six auto- and cross-spectra. These models were then used to create different unconstrained realizations of the CMB temperature and polarization anisotropies, and their spatial derivatives. The C_ℓ fits files were then input to the Healpix synfast code to create a set of 500 maps for each input model. The output synfast maps were specified with 31.2' arcminute smoothing to match the average BICEP2 beam FWHM.

The signal-only maps are used to determine the so-called filter and beam suppression factor. This is the transfer function imparted by polynomial filtering and scan-synchronous template subtraction that reduces the power recovered in the real spectra. The suppression factors are calculated from the mean of the ensemble of signal-only simulations. The real power spectra are then corrected by the mean to recover the true input power spectra. At low ℓ , scan-synchronous template subtraction and polynomial filtering cause signal suppression. At high ℓ , the suppression factor is dominated by the effect of filtering by the 31.2' beam window function, which scales as $e^{-\ell(\ell-1)\sigma^2}$. The empirically-derived suppression factors also include other effects such as the suppression of power by the finite map pixels or pixel window functions. In some formalisms the effects of filtering, the beam and pixel window functions are estimated analytically. In the BICEP1 analysis the filter transfer function was separated from the beam filtering by the use of non-smoothed input maps. In this analysis, due to the use of pre-smoothed input maps, the suppression factor includes the combined effect of the

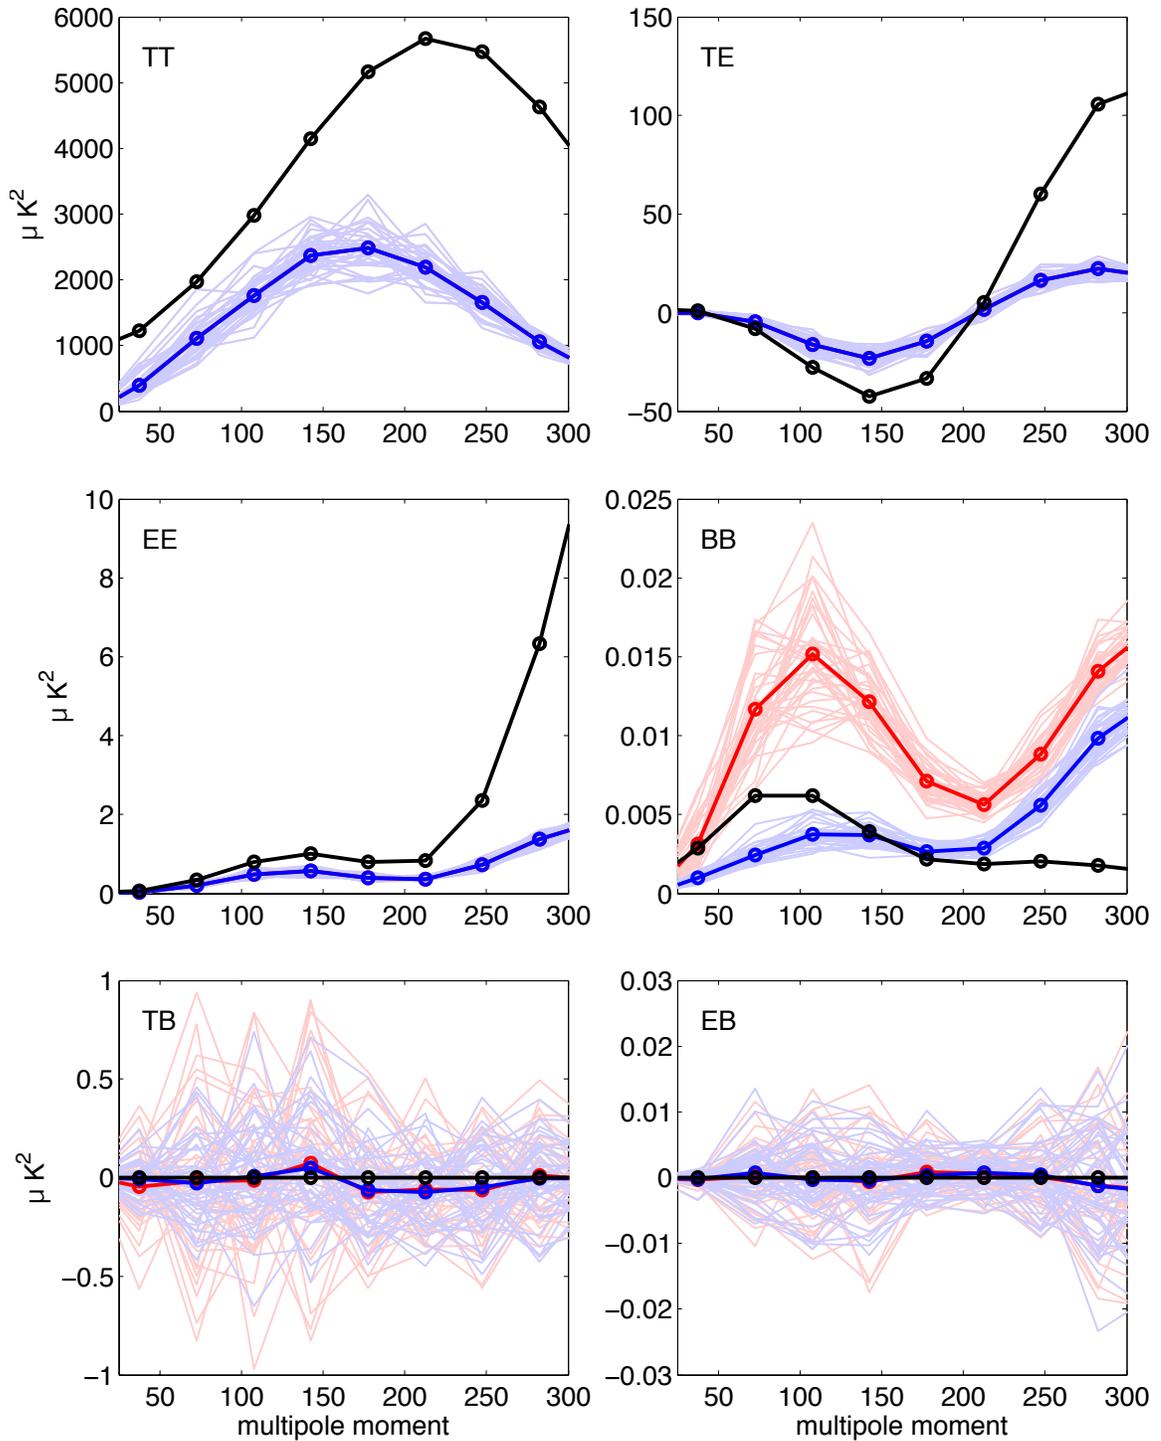


Figure 6.6: Angular power spectra for signal-only simulations shown with the $r = 0.1$ fiducial model. The difference between the expectation values from the model and the result of the signal-only simulations is used to calculate the suppression factor. The leakage in the BB spectrum due to $EE \rightarrow BB$ mixing is shown for the standard (red) and pure (blue) estimators. The mean and fluctuation of that leakage is reduced by a factor of ~ 4 by the pure estimators. In both cases the mean of the ensemble, shown by the heavy red or blue lines, is used to debias the leakage in the BB spectrum.

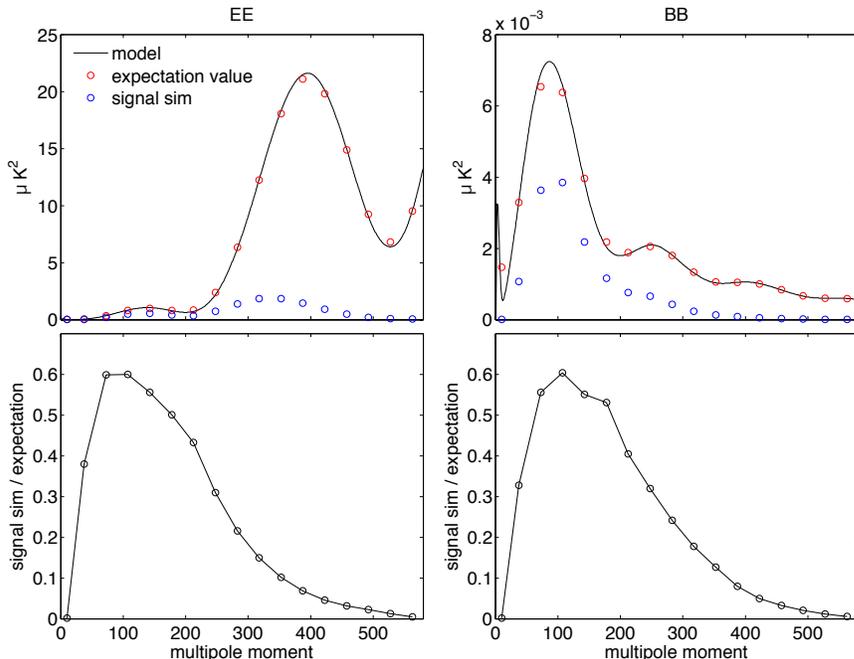


Figure 6.7: The calculation of the EE and BB filter and beam suppression factors. The EE suppression factor is given by the ratio of the mean of the ensemble of signal-only power spectra in which $r = 0$ to the expectation values from the input model. The BB suppression factor is calculated from signal-only simulations that use the $r = 0.1$ input model in which the EE spectrum has been nulled.

beam and timestream filtering.

The real trajectories of the common centroid of each detector pair is used to sample the smoothed input map to create a simulated signal-only timestream. The sampling is performed using a Taylor interpolation of the input map along with its spatial derivatives. The polarization maps sample the smoothed Q and U input maps given the trajectory and projection of the polarization orientation angle of the pair. The simulated timestreams are then processed into maps using the same technique applied to the real maps. Since the signal-only timestreams include no noise they are weighted during the coadd by the variance calculated for the real data. An ensemble of signal-only angular power spectra are then created from the maps.

6.5.3 filter and beam suppression factor

The filter and beam suppression factors are calculated from the ratio of the recovered signal-only power to that expected from the input model. The band power window function is used to calculate the expectation values for each of the spectra from the two input models. The mean of the ensemble of signal-only power spectra is divided by the expectation value to determine the suppression factor.

The procedure and outcome is illustrated in figure 6.7 for the EE and BB suppression factors. The real, noise-only, signal-only and signal-plus-noise power spectra are then divided by the suppression factor to recover the sky power.

The EE suppression factor is calculated from the $r = 0$ input model, while the BB suppression factor is calculated from the model with $r = 0.1$ and nulled EE spectrum. By nulling the signal in EE to calculate BB, and vice versa, this ensures that EE and BB mixing from the finite map does not interfere with the suppression factor extraction. The TT suppression factor is found from the $r = 0$ input model. The mean of the TT and EE or BB suppression factors gives the factor for TE or TB, and EB is formed from the mean of the EE and BB suppression factors. For the standard C_ℓ estimators the EE and BB suppression factors are equivalent, as shown in figure 6.7. When the pure-B estimators are used the suppression factors can differ and must be calculated individually as outlined.

6.6 signal and noise simulations

The error bars for the band powers are set by the fluctuations in the signal-only and noise-only simulated power spectra. The error bars have a component of sample variance as well as instrument noise. The error bars for the band powers are determined from the fluctuations in the power spectra of the ensemble of signal-only and noise-only simulations.

The first step in determining the error bars is to add the signal-only and noise-only simulated maps. A single realization of a signal-only and noise-only map is shown for comparison with the real data in figure 6.2. The ensemble of signal-plus-noise maps are then processed into angular power spectra. The signal-plus-noise spectra are corrected for noise bias, leakage, and suppression factor, just as with the real angular power spectra. The ensemble of signal and noise simulated power spectra are shown in figure 6.8, with the noise and filter bias removed.

The band power covariance matrix $\mathbf{C}_{bb'}$ is then found from the covariance of the ensemble of each of the signal-plus-noise spectra. The square root of the diagonal terms of the 9×9 matrix are the error bars assigned to the nine bins of the real spectra:

$$\delta C_b = \sqrt{\mathbf{C}_{bb}} \quad (6.11)$$

The error bars derived from the 2010 signal-plus-noise simulations have been plotted with the simulated power spectra in figure 6.8. A fiducial model is included in the figure as well to illustrate the

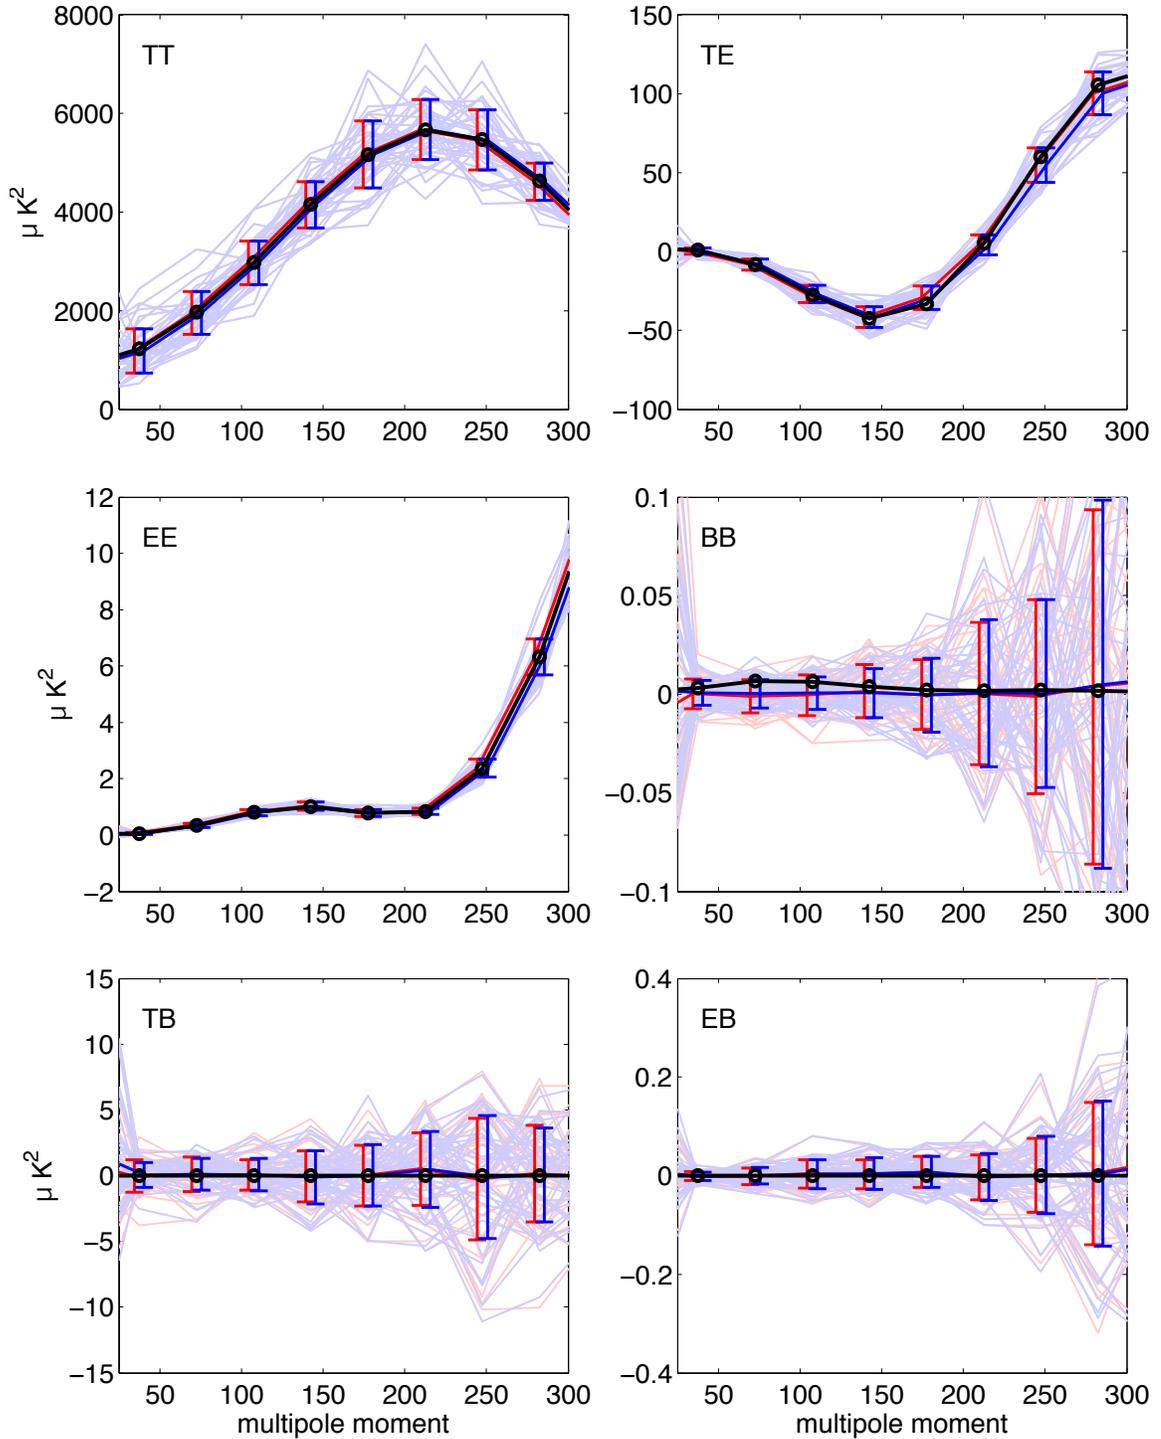


Figure 6.8: Angular power spectra for signal-plus-noise simulations. The ensemble of simulated power spectra are shown along with the $r = 0.1$ input model. These spectra have been corrected for the noise bias, leakage, and suppression factors in the same way the real data are corrected. The resulting means of the distributions are consistent with the $r = 0$ input model they were derived from. The error bars are formed from the signal-plus-noise simulations and are plotted with centers at the mean of the signal-plus-noise spectra. The error bars in the BB spectrum from the pure estimators (blue) are smaller than the standard estimators (red) due to the reduced leakage of the EE variance. The error bars do not include the offset lognormal transformation.

size of the BB error bars in relation to $r = 0.1$.

6.7 consistency checks

6.7.1 jackknives

Real and simulated jackknife maps are formed in order to check the self-consistency of the power spectra. The data are divided into two halves and then coadded into two distinct maps. The maps are then differenced and divided by a factor of two to preserve the noise level in the maps. The jackknife variance maps are created by summing the two halves and dividing by four. The differenced maps are then processed into power spectra. Any residual noise bias is removed in the mean from the ensemble of jackknifed noise-only simulations. The suppression factors are calculated from the non-jackknifed, signal-only maps and applied to the jackknifed spectra so they can be compared with the real spectra.

The resulting maps and power spectra should have no signal and have only statistical fluctuations consistent with the map noise. However, due to interaction between the filtering the signal cancellation between the two maps is not perfect. In general the polarization maps and power spectra have very little residual signal. The temperature maps and spectra can have considerable residual signal.

The BICEP2 pipeline currently has eight defined jackknife data splits. A scan-direction jackknife splits the data between the left- and right-going half scans. This jackknife is sensitive to improper deconvolution of the frequency-dependent gains of the detectors and thermal instability in the focal plane. A tag-split jackknife splits the scan sets that are coadded into the maps chronologically into the first half of the set and second half. This jackknife is sensitive to any transitory systematics and probes systematics caused by changes to the instrument. For example, when coadding maps over the 2010 observing season the first half of the data will include the initial bias and readout rate settings, while the latter data half will include the optimized settings. A phase jackknife divides the data between observing phases (B, C, E) and (F, H, I). These phases cycle between covering different portions of the elevation range of the CMB field. This jackknife is sensitive to ground-fixed or scan-synchronous contamination. A tile jackknife replaces the QU focal plane jackknife used by BICEP1, since all detectors in the BICEP2 focal plane observe only instrument Q. For the tile jackknife the data are split between tiles (1,2) and (3,4) to test for systematics originating in the different detector arrays. Two additional jackknives — multiplexer column and multiplexer row — test for contamination or crosstalk due to the multiplexed readout. For the column jackknife the

data are split between the odd and even columns. For the row jackknife the data are split between rows (0–16) and (17–33), which divides up the focal plane pixels in a checker board pattern.

In BICEP1 a very sensitive boresight angle jackknife was formed by splitting the data between two sets of the four observing angles, where angles within a set are spaced by 45° . This jackknife is sensitive to a number of potential issues in the instrument including differential pointing, thermal stability, gain mismatch, ground pickup and atmospheric opacity. The same jackknife has been attempted in BICEP2 for angle sets $(68^\circ, 113^\circ)$ and $(248^\circ, 293^\circ)$. However, due to the $\sim 9\%$ measured pointing offset in BICEP2 this test is impossible to pass without first regressing out the polarization leakage in the maps. Instead, an alternative boresight jackknife has been established that is less sensitive to differential pointing but still maintains sensitivity to the other effects. The two angle sets are $(68^\circ, 248^\circ)$ and $(113^\circ, 293^\circ)$, so that rotation by 180° will cancel some of the leakage. The tile jackknife shows the same contamination due to differential pointing as in the standard boresight angle. This jackknife should also only be included after leakage in the polarization maps has been cleaned.

6.7.2 χ^2

The consistency of the real and simulated band powers is checked versus the input model using χ^2 tests. χ^2 is computed versus the expectation values from the relevant model across the nine band powers as:

$$\chi^2 = (C_b - C_b^{exp})^\dagger (\mathbf{C}_{bb'})^{-1} (C_b - C_b^{exp}) \quad (6.12)$$

where the band power window function has been used to derive expectation values C_b^{exp} from the model spectra. χ^2 is calculated versus Λ CDM for the non-jackknifed spectra and versus the null model for the jackknives.

This calculation is run for both the real data and for the signal-plus-noise simulations. Due to the low simulation numbers the off-diagonal terms in the band power correlation matrix are noise-dominated. The terms are very small for nonadjacent band powers and average to zero with the inclusion of more simulations. Adjacent band powers are correlated, so they must be retained in the correlation matrix. The main and first two off-diagonal terms are used in the covariance matrix and all others are nulled. When calculating χ^2 for the signal-plus-noise simulations each realization is individually calculated. To prevent large fluctuations in a given realization from erroneously down-weighting the results, the questionable realization is removed from the simulation set when calculating the covariance matrix.

Rather than assuming a Gaussian likelihood function, an offset lognormal distribution is included in the covariance [7]. This involves including the noise bias with the band powers as:

$$Z_b = \ln(C_b + x_b) \quad (6.13)$$

where the noise bias offset term (x_b) is just the mean of the ensemble of noise-only simulations. The quantity in equation 6.13 is calculated in the same manner as the expectation value calculated from the input model. The modified covariance matrix is then given by:

$$(\mathbf{D}_{bb'})^{-1} = (\mathbf{C}_{bb'})^{-1} (C_b + x_b)(C_{b'} + x_{b'}) \quad (6.14)$$

χ^2 is then calculated from the offset lognormal transformation as:

$$\chi^2 = (Z_b - Z_b^{exp})^\dagger (\mathbf{D}_{bb'})^{-1} (Z_b - Z_b^{exp}) \quad (6.15)$$

Absolute gain and beam systematic errors have not been included in the band power covariance matrix for these preliminary results. The error bars for the measured or simulated spectrum are then calculated from the modified covariance matrix $\mathbf{D}_{bb'}$. The effect of including the noise bias in this calculation is an increase in the size of the error bars shown in figure 6.8.

The consistency tests result in 54 χ^2 values for each of the real and signal-plus-noise simulations. Each of the six power spectra from the set of nine standard and jackknifed data types are tested. If the real data are checked for consistency without first cleaning the polarization leakage, then the tile and standard boresight angle jackknives should be removed. This leaves a suite of 42 χ^2 tests. In addition, the TT, TE, and TB spectra have been removed from the consistency tests for both the BICEP1 and BICEP2 analyses. This is mainly due to failure of the temperature-related consistency tests, attributed in part to improper modeling of noise in the pair-sum timestreams. Since the analysis focuses on the measurement of the polarization anisotropy, the inclusion of these consistency tests is not important. This reduces the BICEP2 χ^2 tests to 21 if the leakage is not cleaned. The BICEP1 analysis included 60 tests due to the ability to form cross-frequency jackknives.

The BICEP2 pipeline currently tests the jackknives of the real and signal-plus-noise spectra versus the null model. For the BICEP1 analysis they were evaluated versus the mean of the jackknifed signal-plus-noise spectra [13]. This approach was adopted to account for the lack of perfect cancellation of the signal in the signal-plus-noise jackknives mentioned above. The assumption is made that the

real jackknifed spectra will have the same imperfect cancellation, which will lead to higher χ^2 if compared to the null model. The final BICEP2 analysis may elect to pursue this procedure, though the deviations are generally worst for the temperature-derived spectra.

The suite of χ^2 tests are used to evaluate the probability to exceed (PTE) for the real χ^2 versus the distribution of the signal-plus-noise values. The PTE is a measurement of how well the real data fall within that distribution; it is the measure of the number of simulations with higher χ^2 divided by the number of simulations. A real jackknife χ^2 should fall within the distribution derived from the signal-plus-noise sims. The real data pass consistency checks when none of the real jackknife PTEs are abnormally high or low and when they are distributed equally between zero and one. Currently the real data do not pass the suite of consistency tests. The real jackknives are being used to evaluate several potential sources of failure.

6.8 constraint on r

A constraint on r will be made once the real data pass the consistency tests. Until then, the signal-plus-noise simulations can be used to evaluate the prospects for an upper limit in the absence of systematics. Given the well-measured differential pointing, the choice for evaluating the real r limit should be guided by differential pointing simulations. If the r limit is derived from maps in which the leakage has not been cleaned, then the band powers used for the constraint should be determined from leakage simulations.

The method for placing a constraint on r follows the frequentist approach used for the BICEP1 result [13]. The model power spectra with $r = 0.1$ are read in. The BB spectrum is scaled to a range of trial r values under the assumption the amplitude of the spectrum scales linearly with r . At each trial r value the band power window function is used to calculate the BB expectation value. The noise bias offsets are added to the BB spectrum and the expectation value and χ^2 are calculated for each trial r :

$$\chi^2(r) = (Z_b^{BB} - Z_b^{BB,exp}(r))^\dagger (\mathbf{D}_{bb'}^{BB}(r))^{-1} (Z_b^{BB} - Z_b^{BB,exp}(r)) \quad (6.16)$$

The likelihood function versus r is then given by:

$$\mathcal{L}(r) \propto \frac{1}{\sqrt{\det(\mathbf{D}_{bb'}^{BB}(r))}} e^{-\chi^2(r)/2} \quad (6.17)$$

which has been normalized by the determinant of the modified band powers. The likelihood function

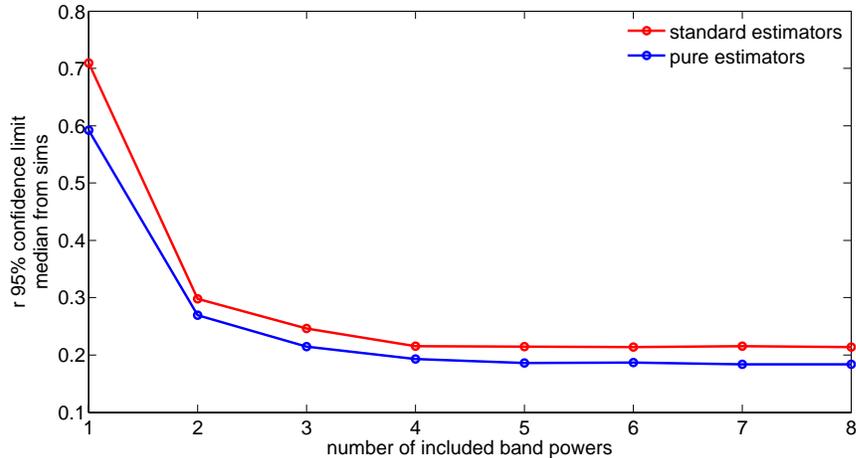


Figure 6.9: The expected 95% confidence limit on r from simulations versus the number of band powers included in the calculation of χ^2 , in the absence of systematic error. The quoted r limit is taken from the median of the 95% confidence limit given by the ensemble of signal-plus-noise simulations including offset lognormal noise bias. The pure estimators result in a lower overall r limit. Both estimates saturate as the number of bins increases due to the increased error bars at higher ℓ . The standard estimators saturate at four bins and the pure estimators saturate at five bins. When differential pointing is included the limits increase and have a minima around three bins.

is integrated to 95% to determine the upper limit on r . The maximum of the likelihood function gives the best-fit value, with σ given by the contour that encloses 68% of the likelihood integral. This procedure is repeated for each of the signal-plus-noise simulations and for the real data when appropriate.

6.8.1 band power selection

The bins that are used for the constraint on r must be carefully chosen to avoid the polarization leakage caused by differential pointing. As shown in figure 5.3, the contamination to the BB spectrum from signal-only simulations including differential pointing increases with ℓ and peaks around the beam scale at $\ell \sim 300$. The constraint on r is sensitive to both the amplitude of the band power and size of the error bar on the BB spectrum.

In the absence of systematic errors, the constraint on r improves with the number of bins (starting from the first) included in equation 6.16. Figure 6.9 shows the median 95% confidence limit from the signal-plus-noise simulations versus the number of included band powers. The constraint saturates with the inclusion of more than four or five bins due to the rapid increase in the error bars at higher ℓ . The r limit from the standard estimators saturates at four bins, while the pure estimators saturate at five bins. In the absence of systematics, the best constraint that should be expected for the 2010

season is $r < 0.215$ for the standard estimators, based on the median of the simulations. The limit improves to $r < 0.186$ when the pure estimators are used.

However, the measured differential pointing calls for using fewer bins in the constraint, due to the increased band powers and errors in the higher ℓ bins. Signal-only simulations have been run using the measured best estimates of the differential pointing without additional scatter for the measurement error. These signal-only simulations were combined with the 2010 noise simulations to form signal-plus-noise simulations and power spectra with the proper 2010 noise levels. The noise bias was removed from the signal-plus-noise simulations based on the noise-only spectra. The $EE \rightarrow BB$ leakage and suppression factor was removed based on the mean leakage from the signal-only simulations that do not include the pointing offsets. The noise bias was included in the offset lognormal transformation and the likelihood functions were calculated using equation 6.17, while letting the number of included bins vary in equation 6.16.

The 95% confidence limit versus number of band powers showed that the minimum constraint is given by using only the first three bins. Above that the limit rapidly degrades due to the increased band powers at higher ℓ . Furthermore, the results from the pure estimators give a much worse constraint for all bin numbers versus the standard estimators. This is believed to be due to the mixing of power between multipole moments and band powers that occurs due to the pure BB band power estimation. The pure estimators mix the large contamination at $\ell > 100$ with the less-contaminated low- ℓ bins. In the presence of differential pointing the pure estimator gives a best constraint of $r < 0.47$. The best constraint from the standard estimators in the presence of differential pointing is $r < 0.32$.

6.8.2 conclusions

Based on the results of simulations, which include the effect of the measured differential pointing, the best constraint on r that can be placed with the real BICEP2 2010 data is $r < 0.32$. This is only a factor of two better than the $r < 0.72$ constraint placed by BICEP1, despite the factor of 10 improvement in mapping speed. In order to do better than this predicted limit the leakage in the polarization maps must be removed using the regression analysis described in § 5.7.2. This procedure has been demonstrated to very effectively clean systematic-induced leakage from the maps being used for BICEP1 three-year result. The regression analysis is currently being integrated into the BICEP2 pipeline, and should begin producing cleaned maps soon. Once the polarization maps have been cleaned the upper limit from the 2010 data should allow a constrain on $r < 0.19$, using the

pure estimators.⁵ This constraint should improve by reducing the residual EE variance leakage into the BB spectrum. This will require the optimization of the map projection, apodization mask and eventually moving to a curved-sky analysis. The detector and readout optimization work has lead to a factor of ~ 2 improvement in the mapping speed of the 2011–2012 seasons over the 2010 season. This increase in raw instrument sensitivity should allow BICEP2 to eventually probe $r \sim 0.01$ in the three-year maps.

⁵Assuming the error bars don't substantially increase due to the regression analysis.

Bibliography

- [1] P A R Ade, G. Pisano, C Tucker, and S. Weaver. A review of metal mesh filters. *Proc. SPIE*, 6275:62750U, 2006.
- [2] Randol W Aikin, P A Ade, S Benton, J J Bock, J A Bonetti, J A Brevik, C D Dowell, L Duband, J P Filippini, S R Golwala, M Halpern, V V Hristov, K Irwin, J P Kaufman, B G Keating, J M Kovac, C L Kuo, A E Lange, C B Netterfield, H T Nguyen, R W Ogburn IV, A Orlando, C Pryke, S Richter, J E Ruhl, M C Runyan, C Sheehy, S A Stokes, R Sudiwala, G P Teply, J E Tolan, A D Turner, P Wilson, and C L Wong. Proceedings of SPIE. In *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V*, pages 77410V–77410V–9. SPIE, 2010.
- [3] ES Battistelli, M Amiri, B Burger, MJ Devlin, SR Dicker, WB Doriese, R Dunner, RP Fisher, JW Fowler, and M Halpern. Automated SQUID tuning procedure for kilo-pixel arrays of TES bolometers on the Atacama Cosmology Telescope. *Proceedings of SPIE*, 7020:702028, 2008.
- [4] D J Benford, M.C. Gaidis, and J.W. Kooi. Optical properties of Zitex in the infrared to submillimeter. *Applied Optics*, 42(25):5118–5122, 2003.
- [5] RS Bhatia, ST Chase, SF Edgington, J Glenn, WC Jones, AE Lange, B. Maffei, AK Mainzer, PD Mauskopf, and BJ Philhour. A three-stage helium sorption refrigerator for cooling of infrared detectors to 280 mK. *Cryogenics*, 40(11):685–691, 2000.
- [6] E M Bierman, T Matsumura, C D Dowell, B G Keating, P Ade, D Barkats, D Barron, J O Battle, J J Bock, H C Chiang, T L Culverhouse, L Duband, E F Hivon, W L Holzzapfel, V V Hristov, J P Kaufman, J M Kovac, C L Kuo, A E Lange, E M Leitch, P V Mason, N J Miller, H T Nguyen, C Pryke, S Richter, G M Rocha, C Sheehy, Y D Takahashi, and K W Yoon. A Millimeter-Wave Galactic Plane Survey With The BICEP Polarimeter. *arXiv.org*, astro-ph.GA, March 2011.

- [7] JR Bond, AH Jaffe, and L. Knox. Radical compression of cosmic microwave background data. *The Astrophysical Journal*, 533:19, 2000.
- [8] J A Bonetti, P K Day, M Kenyon, Chao-Lin Kuo, A Turner, H G LeDuc, and J J Bock. Characterization of Antenna-Coupled TES Bolometers for the Spider Experiment. *IEEE Transactions on Applied Superconductivity*, 19(3):520–523, 2009.
- [9] L.A. Boyle, P.J. Steinhardt, and N. Turok. Inflationary predictions for scalar and tensor fluctuations reconsidered. *Physical review letters*, 96(11):111301, 2006.
- [10] A. Challinor and G. Chon. Error analysis of quadratic power spectrum estimates for cosmic microwave background polarization: sampling covariance. *Monthly Notices of the Royal Astronomical Society*, 360(2):509–532, 2005.
- [11] Anthony Challinor. Cosmic microwave background polarization analysis. *arXiv.org*, astro-ph, February 2005.
- [12] H.C. Chiang. *Observation of Cosmic Microwave Background Polarization with BICEP*. PhD thesis, California Institute of Technology, 2009.
- [13] HC Chiang, PAR Ade, D Barkats, JO Battle, EM Bierman, JJ Bock, CD Dowell, L Duband, EF Hivon, and WL Holzapfel. Measurement of cosmic microwave background polarization power spectra from two years of BICEP data. *The Astrophysical Journal*, 711:1123, 2010.
- [14] J.C. CLARKE and A. I. Braginski. The SQUID handbook, vol. 1: fundamentals and technology of SQUIDS and SQUID systems. *Recherche*, 67:02, 2003.
- [15] Tom Crawford. Power spectrum sensitivity of raster-scanned CMB experiments in the presence of 1/f noise. *arXiv.org*, astro-ph, February 2007.
- [16] P.A.J. de Korte, J Beyer, S Deiker, G C Hilton, K D Irwin, M. MacIntosh, S W Nam, C D Reintsema, L R Vale, and M.E. Huber. Time-division superconducting quantum interference device multiplexer for transition-edge sensors. *Review of scientific instruments*, 74:3807, 2003.
- [17] D Drung, C Assmann, J Beyer, A Kirste, M Peters, F Ruede, and T Schurig. Highly sensitive and easy-to-use SQUID sensors. *Applied Superconductivity, IEEE Transactions on*, 17(2):699–704, 2007.
- [18] J Dunkley, R Hlozek, J Sievers, V Acquaviva, P A R Ade, P Aguirre, M Amiri, J W Appel, L F Barrientos, E S Battistelli, J R Bond, B Brown, B Burger, J Chervenak, S Das, M J Devlin,

- S R Dicker, W Bertrand Doriese, R Dunner, T Essinger-Hileman, R P Fisher, J W Fowler, A Hajian, M Halpern, M Hasselfield, C Hernandez-Monteagudo, G C Hilton, M Hilton, A D Hincks, K M Huffenberger, D H Hughes, J P Hughes, L Infante, K D Irwin, J B Juin, M Kaul, J Klein, A Kosowsky, J M Lau, M Limon, Y-T Lin, R H Lupton, T A Marriage, D Marsden, P Mauskopf, F Menanteau, K Moodley, H Moseley, C B Netterfield, M D Niemack, M R Nolta, L A Page, L Parker, B Partridge, B Reid, N Sehgal, B Sherwin, D N Spergel, S T Staggs, D S Swetz, E R Switzer, R Thornton, H Trac, C Tucker, R Warne, E Wollack, and Y. Zhao. The Atacama Cosmology Telescope: Cosmological Parameters from the 2008 Power Spectra. *arXiv.org*, astro-ph.CO, September 2010.
- [19] J Dunkley, E Komatsu, MR Nolta, DN Spergel, D Larson, G Hinshaw, L Page, CL Bennett, B Gold, and N Jarosik. Five-year wilkinson microwave anisotropy probe observations: likelihoods and parameters from the WMAP data. *The Astrophysical Journal Supplement Series*, 180:306, 2009.
- [20] E. Figueroa Feliciano. *Theory and development of position-sensitive quantum calorimeters*. PhD thesis, Stanford, 2001.
- [21] D.P. Finkbeiner, M. Davis, and D.J. Schlegel. Extrapolation of galactic dust emission at 100 microns to cosmic microwave background radiation frequencies using FIRAS. *The Astrophysical Journal*, 524:867, 1999.
- [22] A.H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Physical Review D*, 23(2):347, 1981.
- [23] M. Hasselfield. Threshold crosstalk in act. private communication, 2010.
- [24] E Hivon, K M Gorski, C B Netterfield, B P Crill, S Prunet, and F Hansen. MASTER of the cosmic microwave background anisotropy power spectrum: A fast method for statistical analysis of large and complex cosmic microwave background data sets. *The Astrophysical Journal*, 567:2, 2002.
- [25] W. Hu and M. White. A CMB polarization primer. *New Astronomy*, 2(4):323–344, 1997.
- [26] Wayne Hu, Matthew M Hedman, and Matias Zaldarriaga. Benchmark Parameters for CMB Polarization Experiments. *arXiv.org*, astro-ph, October 2002.
- [27] M.E. Huber, A.H. Steinbach, and R.H. Ono. Resonance damping in tightly coupled dc SQUIDS via intra-coil resistors. *Physica C: Superconductivity*, 351(2):85–90, 2001.

- [28] KD Irwin. An application of electrothermal feedback for high resolution cryogenic particle detection. *Applied Physics Letters*, 66:1998, 1995.
- [29] KD Irwin. SQUID multiplexers for transition-edge sensors. *Physica C: Superconductivity*, 368(1):203–210, 2002.
- [30] KD Irwin and GC Hilton. Transition-edge sensors. *Cryogenic Particle Detection*, pages 81–97, 2005.
- [31] J. Jaycox and M. Ketchen. Planar coupling scheme for ultra low noise dc SQUIDS. *Magnetics, IEEE Transactions on*, 17(1):400–403, 1981.
- [32] R Keisler, C L Reichardt, K A Aird, B A Benson, L E Bleem, J E Carlstrom, C L Chang, H M Cho, T M Crawford, A T Crites, T de Haan, M A Dobbs, J Dudley, E M George, N W Halverson, G P Holder, W L Holzzapfel, S Hoover, Z Hou, J D Hrubes, M Joy, L. Knox, A T Lee, E M Leitch, M Lueker, D Luong-Van, J J McMahon, J Mehl, S S Meyer, M Millea, J J Mohr, T E Montroy, T Natoli, S Padin, T Plagge, C Pryke, J E Ruhl, K K Schaffer, L Shaw, E Shirokoff, H G Spieler, Z Staniszewski, A A Stark, K Story, A van Engelen, K Vanderlinde, J D Vieira, R Williamson, and O Zahn. A Measurement of the Damping Tail of the Cosmic Microwave Background Power Spectrum with the South Pole Telescope. *arXiv.org*, astro-ph.CO, May 2011.
- [33] L. Knox and Y.S. Song. A limit on the detectability of the energy scale of inflation. *Arxiv preprint astro-ph/0202286*, 2002.
- [34] Lloyd Knox. Cosmic Microwave Background Anisotropy Window Functions Revisited. *arXiv.org*, astro-ph, February 1999.
- [35] Lloyd Knox and Michael S Turner. Detectability of Tensor Perturbations Through CBR Anisotropy (final published version). *arXiv.org*, astro-ph, July 1994.
- [36] E Komatsu, K M Smith, J Dunkley, C L Bennett, B Gold, G Hinshaw, N Jarosik, D Larson, M R Nolta, L Page, D N Spergel, M Halpern, R S Hill, A Kogut, M Limon, S S Meyer, N Odegard, G S Tucker, J L Weiland, E Wollack, and E L Wright. Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *arXiv.org*, astro-ph.CO, January 2010.
- [37] J Kovac, EM Leitch, C Pryke, JE Carlstrom, NW Halverson, and WL Holzzapfel. Detection of polarization in the cosmic microwave background using DASI, 2002. *Nature*, 420:772, 2002.

- [38] J M Kovac and D Barkats. CMB from the South Pole: Past, Present, and Future. *arXiv.org*, astro-ph, July 2007.
- [39] C L Kuo, J J Bock, J A Bonetti, J Brevik, G Chattopadhyay, P K Day, S Golwala, M Kenyon, A E Lange, H G LeDuc, H Nguyen, R W Ogburn, A Orlando, A Trangsrud, A Turner, G Wang, and J Zmuidzinas. Antenna-coupled TES bolometer arrays for CMB polarimetry. *arXiv.org*, astro-ph.IM, August 2009.
- [40] D Larson, J Dunkley, G Hinshaw, E Komatsu, MR Nolta, CL Bennett, B Gold, M Halpern, RS Hill, and N Jarosik. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: power spectra and WMAP-derived parameters. *The Astrophysical Journal Supplement Series*, 192:16, 2011.
- [41] A. Lewis, A. Challinor, and A. Lasenby. Efficient computation of cosmic microwave background anisotropies in closed Friedmann-Robertson-Walker models. *The Astrophysical Journal*, 538:473, 2000.
- [42] Tomotake Matsumura, Peter Ade, Denis Barkats, Darcy Barron, John O Battle, Evan M Bierman, James J Bock, H Cynthia Chiang, Brendan P Crill, C Darren Dowell, Lionel Duband, Eric F Hivon, William L Holzappel, Viktor V Hristov, William C Jones, Brian G Keating, John M Kovac, Chao-lin Kuo, Andrew E Lange, Erik M Leitch, Peter V Mason, Hien T Nguyen, Nicolas Ponthieu, Clem Pryke, Steffen Richter, Graca M Rocha, Yuki D Takahashi, and Ki Won Yoon. Absolute polarization angle calibration using polarized diffuse Galactic emission observed by BICEP. *arXiv.org*, astro-ph.IM, July 2010.
- [43] J V Migacz and M.E. Huber. Thermal annealing of Nb/Al-AlO_x/Nb Josephson junctions. *Applied Superconductivity, IEEE Transactions on*, 13(2):123–126, 2003.
- [44] M.D. Niemack. *Towards Dark Energy: Design, development and preliminary data from ACT*. PhD thesis, Princeton University, 2008.
- [45] R W Ogburn IV, Randol W Aikin, P A Ade, S Benton, J J Bock, J A Bonetti, J A Brevik, C D Dowell, L Duband, J P Filippini, S R Golwala, M Halpern, V V Hristov, K Irwin, J P Kaufman, B G Keating, J M Kovac, C L Kuo, A E Lange, C B Netterfield, H T Nguyen, A Orlando, C Pryke, S Richter, J E Ruhl, M C Runyan, C Sheehy, S A Stokes, R Sudiwala, G P Teply, J E Tolan, A D Turner, P Wilson, and C L Wong. Proceedings of SPIE. In

Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy V. SPIE, 2010.

- [46] A Orlando, RW Aikin, M Amiri, JJ Bock, JA Bonetti, JA Brevik, B Burger, PK Day, JP Filippini, and SR Golwala. Antenna-coupled TES arrays for the BICEP2/Keck and SPIDER polarimeters. *American Institute of Physics Conference Series*, 31(1):471, 2009.
- [47] L Page, G Hinshaw, E Komatsu, MR Nolta, DN Spergel, CL Bennett, C. Barnes, R. Bean, O. Doré, and J Dunkley. Three-year Wilkinson microwave anisotropy probe (WMAP) observations: polarization analysis. *The Astrophysical Journal Supplement Series*, 170:335, 2007.
- [48] E. Pascale, PAR Ade, JJ Bock, EL Chapin, J. Chung, MJ Devlin, S. Dicker, M. Griffin, JO Gundersen, and M Halpern. The Balloon-borne large aperture submillimeter telescope: BLAST. *The Astrophysical Journal*, 681:400, 2008.
- [49] C Pryke, P Ade, J Bock, M Bowden, ML Brown, G Cahill, PG Castro, S Church, T Culverhouse, and R Friedman. Second and Third Season QUaD Cosmic Microwave Background Temperature and Polarization Power Spectra. *The Astrophysical Journal*, 692:1247, 2009.
- [50] S.J.E. Radford and R.A. Chamberlin. ALMA Memo 334.1 Atmospheric Transparency at 225 GHz over Chajnantor, Mauna Kea, and the South Pole. 2000.
- [51] C D Reintsema, J Beyer, S W Nam, S Deiker, G C Hilton, K Irwin, J. Martinis, J. Ullom, L R Vale, and M. MacIntosh. Prototype system for superconducting quantum interference device multiplexing of large-format transition-edge sensor arrays. *Review of scientific instruments*, 74:4500, 2003.
- [52] U. Seljak. Measuring polarization in cosmic microwave background. *Arxiv preprint astro-ph/9608131*, 1996.
- [53] C D Sheehy, P A R Ade, R W Aikin, M Amiri, S Benton, C Bischoff, J J Bock, J A Bonetti, J A Brevik, B Burger, C D Dowell, L Duband, J P Filippini, S R Golwala, M Halpern, M Hasselfield, G Hilton, V V Hristov, K Irwin, J P Kaufman, B G Keating, J M Kovac, C L Kuo, A E Lange, E M Leitch, M Lueker, C B Netterfield, H T Nguyen, R W Ogburn IV, A Orlando, C L Pryke, C Reintsema, S Richter, J E Ruhl, M C Runyan, Z Staniszewski, S Stokes, R Sudiwala, G Teply, K L Thompson, J E Tolan, A D Turner, P Wilson, and C L Wong. The Keck Array: a pulse tube cooled CMB polarimeter. *arXiv.org*, astro-ph.IM, April 2011.

- [54] Meir Shimon, Brian Keating, Nicolas Ponthieu, and Eric Hivon. CMB polarization systematics due to beam asymmetry: Impact on inflationary science. *Physical Review D*, 77(8), April 2008.
- [55] K M Smith. Pseudo-C_l estimators which do not mix E and B modes. *Physical Review D*, 74(8):083002, 2006.
- [56] Kendrick M Smith and Matias Zaldarriaga. A general solution to the E-B mixing problem. *arXiv.org*, astro-ph, October 2006.
- [57] D N Spergel, R. Bean, O. Doré, MR Nolta, CL Bennett, J Dunkley, G Hinshaw, N Jarosik, E Komatsu, and L Page. Three-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: implications for cosmology. *The Astrophysical Journal Supplement Series*, 170:377, 2007.
- [58] G M Stiehl, H M Cho, G C Hilton, K D Irwin, J A B Mates, C D Reintsema, and B L Zink. Time-division SQUID multiplexers with reduced sensitivity to external magnetic fields. *arXiv.org*, astro-ph.IM, August 2010.
- [59] Y D Takahashi, P A R Ade, D Barkats, J O Battle, E M Bierman, J J Bock, H C Chiang, C D Dowell, L Duband, E F Hivon, W L Holzapfel, V V Hristov, W C Jones, B G Keating, J M Kovac, C L Kuo, A E Lange, E M Leitch, P V Mason, T Matsumura, H T Nguyen, N Ponthieu, C Pryke, S Richter, G Rocha, and K W Yoon. CHARACTERIZATION OF THE BICEP TELESCOPE FOR HIGH-PRECISION COSMIC MICROWAVE BACKGROUND POLARIMETRY. *The Astrophysical Journal*, 711(2):1141–1156, February 2010.
- [60] JN Ullom, WB Doriese, GC Hilton, JA Beall, S Deiker, WD Duncan, L Ferreira, KD Irwin, CD Reintsema, and LR Vale. Characterization and reduction of unexplained noise in superconducting transition-edge sensors. *Applied Physics Letters*, 84:4206, 2004.
- [61] K. W. Yoon. *Design and Deployment of BICEP*. PhD thesis, California Institute of Technology, 2007.
- [62] K W Yoon, P A R Ade, D Barkats, J O Battle, E M Bierman, J J Bock, J A Brevik, H C Chiang, A Crites, C D Dowell, L Duband, G S Griffin, E F Hivon, W L Holzapfel, V V Hristov, B G Keating, J M Kovac, C L Kuo, A E Lange, E M Leitch, P V Mason, H T Nguyen, N Ponthieu, Y D Takahashi, T Renbarger, L C Weintraub, and D Woolsey. The Robinson Gravitational Wave Background Telescope (BICEP): a bolometric large angular scale CMB polarimeter. *arXiv.org*, astro-ph, June 2006.