HARVARD
Kenneth C. Griffin

GRADUATE SCHOOL
OF ARTS AND SCIENCES

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Astronomy

have examined a dissertation entitled

Windows to the Universe: Improving the Sensitivity of High Throughput Millimeter Telescopes

presented by Miranda Eiben

candidate for the degree of Doctor of Philosophy and hereby

certify that it is worthy of acceptance.

Signature _____

Typed name: Prof. Josh Grindlay, Chair

Signature _____

Typed name: Prof. John Kovac

Signature _____

Typed name: Prof. Doug Finkbeiner

Signature _____

Typed name: Dr. Scott Paine

Signature _____

Typed name: Prof. Kirit Karkare

Date: May 6, 2025

Page left intentionally blank.

Windows to the Universe: Improving the Sensitivity of High Throughput Millimeter Telescopes

A dissertation presented

by

Miranda Eiben

to

The Department of Astronomy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Astronomy & Astrophysics

Harvard University

Cambridge, Massachusetts

May 2025

© 2025 — Miranda Eiben

All rights reserved.

Dissertation Advisor: Professor John Kovac Miranda Eiben

Windows to the Universe: Improving the Sensitivity of High Throughput Millimeter Telescopes

Abstract

Millimeter-wave refracting telescopes targeting the degree-scale structure of the cosmic microwave background (CMB) have recently grown to diffraction-limited apertures of over 0.5 meters. These instruments are entirely housed in vacuum cryostats to support their sub-kelvin bolometric detectors and to minimize radiative loading from thermal emission due to absorption loss in their transmissive optical elements. The large vacuum window is the only optical element in the system at ambient temperature, and therefore minimizing loss in the window is crucial for maximizing detector sensitivity. This motivates the use of low-loss polymer materials and a window as thin as practicable. However, the window must simultaneously meet the requirement to keep sufficient vacuum, and therefore must limit gas permeation and remain mechanically robust against catastrophic failure under pressure. I report on the development of extremely thin composite polyethylene window technology that meets these goals. Two windows have been deployed for two full observing seasons on the BICEP3 and BA150 CMB telescopes at the South Pole. On BICEP3, the window has demonstrated a 6% improvement in detector sensitivity.

The larger apertures also produce a challenging anti-reflection (AR) design problem for refracting and transmissive optics. AR layers are required to minimize the light lost at each transmissive interface within a receiver, which therefore improves the optical

throughput of the instrument. The plastic optics require consistently thin polymer coats across a wide area, while wide bandwidths require multilayer designs. I present multilayer AR coats for plastic optics of the BICEP Array receivers (30–300 GHz) utilizing an expanded polytetrafluoroethylene (ePTFE) membranes. These ePTFE membranes can be fine-tuned to the ideal quarter wavelength solutions with heated compression, and some types of ePTFE can be layered and compressively heat-bonded to generate thicker stacks. This heated compression process allows for a range of densities (from 0.3 g/cc to 1 g/cc) and thicknesses (>0.05 mm) over a wide radius (33 cm), opening the parameter space of potential AR coats in interesting directions. The ePTFE anti-reflection coats have produced band average reflections close to the ideal for the polyethylene optics on all four receivers in BICEP Array.

The four receivers that make up BICEP Array target different frequencies to aid in measuring and separating the polarized astrophysical components along the line-of-sight to the CMB. We measure the frequency dependent optical throughput of these receivers with a Fourier transform spectrometer (FTS), and our ability to constrain cosmological parameters is dependent on our understanding of the systematics of this optical throughput measurement. I report the biases induced in the eight parameter cosmological model by the bandpass shifts within our maximum likelihood search framework; the maximum shift in the tensor to scalar ratio r was found to be $9.1 \pm 3.8 \times 10^{-4}$ with conservative 2% shifts to bandcenters.

Contents

Title Page **i**

Copyright **ii**

Abstract **iii**

Contents **v**

List of Figures **ix**

List of Tables **xv**

Acknowledgments **xvii**

Dedication **xx**

1 Introduction **1**

1.1 The Universe and Everything In It 1

1.1.1 The Expanding Universe 3

1.1.2 The Cosmic Background 9

1.1.3 The Initial Conditions 19

1.2 The BICEP/*Keck* Experiment 23

1.2.1 The Receivers 26

1.2.2 The Observations 37

1.2.3 The Model 40
1.3 The Thesis 44

2 Thin Window Mechanics 46
2.1 Thin Windows 51
2.2 Strength 53
2.2.1 Lamination Temperature 55
2.2.2 Lamination Pressure 58
2.2.3 Comparison to Bulk Polyethylene 61
2.3 Safety Factor Measurement 63
2.4 Creep 67
2.5 Gas Permeation 69
2.6 Environmental 75

3 Wide Bandwidth Anti-reflection Coats 79
3.1 Ideal AR Coats 82
3.1.1 Chebyshev Polynomials 85
3.1.2 Ideal AR Coats Compared to Achieved AR Coats 89
3.2 Low Frequency AR Coats (25–45 GHz) 92
3.3 Mid Frequency AR Coats (85–170 GHz) 95
3.3.1 Teadit: Compressible PTFE 95
3.3.2 Using Teadit 98
3.4 High Frequency AR Coats (200–300 GHz) 102
3.4.1 DeWAL: Stackable ePTFE 103
3.4.2 Using DeWAL 106
3.5 Problems with Using ePTFEs 114

4 Thin Window Optics 120

4.1 Optical Tests and Modeling at Home . . . 121
4.1.1 Index of Refraction . . . 121
4.1.2 Birefringence . . . 124
4.2 Optical Characterization on BICEP3 . . . 135
4.2.1 Measured Material Optical Load . . . 137
4.2.2 Change in Instrument Optical Load . . . 140

5 Instrumental Noise in BICEP3 151
5.1 Receiver White Noise . . . 152
5.1.1 Optical (Photon) Noise . . . 152
5.1.2 Other Sources of Noise . . . 155
5.1.3 Noise Equivalent Temperature . . . 158
5.2 Measured Noise in BICEP3 . . . 159
5.2.1 Uncertainty on Absolute Calibration . . . 160
5.3 Modeled Noise in BICEP3 . . . 164
5.3.1 Previous Model . . . 165
5.3.2 Previous Model with New Temperatures . . . 171
5.3.3 New Model with New Temperatures . . . 176

6 Receiver Bandpass Characterization 181
6.1 Fourier Transform Spectroscopy . . . 184
6.1.1 Interferometer . . . 184
6.2 BA150 Spectra . . . 190
6.2.1 Analysis . . . 193
6.2.2 Spectral Uncertainties . . . 198
6.3 Component Separation . . . 202
6.3.1 Pseudosimulations . . . 203

7 Conclusions **214**

7.1 Laminate Windows 214

7.2 Anti-reflection Coats 220

7.3 Spectra Analysis 224

A How to Build a Laminate Window **229**

A.1 Designing the Window 231

A.2 Cutting the HMPE (or anything else) on the Rotating Template 232

A.2.1 Using the Template for Window Materials 233

A.2.2 Using the Template for Cutting AR Layer 237

A.3 Making the Laminate Stack 239

B How to Run the Optics Autoclaves **243**

B.1 Small Autoclave 244

B.2 Large Autoclave 249

C How to Layer DeWAL **259**

D Ideal (Chebyshev) AR Indexes for Millimeter Optics **264**

D.1 PTFE (Teflon): $n \sim 1.43$ 266

D.2 HDPE: $n \sim 1.55$ 267

D.3 Nylon: $n \sim 1.73$ 268

D.4 Quartz: $n \sim 2.10$ 269

D.5 Alumina: $n \sim 3.10$ 270

D.6 Silicon: $n \sim 3.44$ 271

References **272**

List of Figures

1.1 *Planck* all-sky maps from 30 GHz to 857 GHz. 11

1.2 Cartoon representation of an electron sitting in a monopole, dipole, and quadrupole temperature field. 15

1.3 Cartoon representation of quadrupole temperature distribution seen by an electron in a density wave. 16

1.4 ΛCDM fit to the *Planck* satellite power spectra. 18

1.5 Cartoon representation of quadrupole temperature distribution seen by an electron in a gravitational wave. 22

1.6 ΛCDM predictions of primordial temperature, E-mode, and B-mode angular power spectra. 23

1.7 Atmospheric transmission at the South Pole and the measured bandpasses of the BICEP/*Keck* receivers. 25

1.8 BA 150 GHz detector tile, and the photon path through a TES. 27

1.9 Cutaway view of the top of BICEP3. 31

1.10 Cutaway view of BA150. 32

1.11 Exterior of the BICEP Array mount in the austral summer of 2025. 36

1.12 Example of the scan strategy of BICEP3 on the sky. 38

1.13 BB auto-spectra decomposed into lensed-ΛCDM, dust, and upper limit synchrotron. 42

1.14 Recent suite of published BB power spectrum measurements. 43

2.1 Modeled optical loading from windows with different thicknesses at band centers from 30 to 270 GHz through a BICEP-style instrument. 49

2.2 Photo of a window material sample in tensile strength machine and tensile strength tests of three samples. 54

2.3 Photo of the highest temperature laminate sample. 57

2.4 Photo of several small scale laminated window samples baked at different pressures. 59

2.5 Photo of the window at the highest pressure achieved during the hydrostatic test. 65

2.6 Example photos of 'stress striping' behaviors in laminate windows. 66

2.7 Measured center deflections of a full-scale window and estimated radial strain. 68

2.8 Measured permeation rates of HMPE samples compared to the predicted permeation rate. 71

2.9 Pictures of the BICEP3 thin window taken with an IR camera in situ. . . . 76

3.1 Cartoon representation of reflected waves off of a single layer AR coated material. 81

3.2 Reflections from ideal single, double, and triple layer anti-reflection coats for HDPE. 84

3.3 The Lorenz-Lorentz and linear relationship vs measured densities and indexes of PTFE membranes. 90

3.4 Band average reflections for ideal AR coats on infinitely thick cores of HDPE, nylon, and alumina. 91

3.5 VNA reflection measurements for the AR coats on polymer optics for BA30/40. 94

3.6 Optical path length, thickness, and index estimates over a sheet of Teadit. 96

3.7 Compression pressures, estimated final optical path of the two formulations of Teadit. 97

3.8 VNA reflection measurements of AR coated optics samples for 95 and 150 GHz optics. 99

3.9 Design corner plots for thin window AR coats on BICEP3 and BA150. . . 100

3.10 Measured VNA reflections of the BICEP3 thin window and modeled reflections of the laminate. 102

3.11 DeWAL density and estimated index at different compression pressures. . . 106

3.12 Design corner plot for multilayer HF HDPE optics. 108

3.13 Measurements of HF (BA220/270) lens AR coats. 110

3.14 Out of band measured and modeled in-band reflections of the science grade BA220/270 window. 115

3.15 Photo of the DeWAL stretching rig and the AR coating of the Short Keck (SK) window. 116

3.16 Thickness distribution across width over the roll of two DeWAL rolls. . . . 118

4.1 VNA measurements of HMPE window cores baked at different pressures. . 122

4.2 Diagram of phase offset measurement and fitted phase slopes. 125

4.3 Example Poinecaré sphere and polarization ellipses of a quarterwave plate. 131

4.4 Poinecaré sphere and polarization ellipses of a maximally birefringent (n_f = 1.536 and n_s = 1.556, h =1.4mm) thin window. 132

4.5 Histograms of excess loading and scattered power per detector from two different windows placed above BICEP3. 139

4.6 Cartoon representation of various named powers on a light TES detector. . 141

4.7 Example per detector fits to the PLCs in the elevation vs loading schedules. 143

4.8 Comparison of the fit per detector P_{const} and P_{sky} between 2017 and 2023. . 144

4.9 Comparison of median yearly per dark detector measured P_{sat} between 2017 and 2022. 147

4.10 Comparison of median yearly per light detector between 2017 and three other representative seasons. 148

5.1 BICEP3 median white noise distribution over all detectors per year between 2017 and 2023. 161

5.2 Estimates of the absolute calibration uncertainties for BICEP3 for multipole bins. 163

5.3 Median per year cold optics temperatures in BICEP3. 172

5.4 Median measured and modeled NETs (white noise) from as a function of the window thickness. 180

6.1 The modeled optical power on the detectors of a BICEP3-like receiver with varied bandcenter from 30 to 270 GHz. 183

6.2 Near field beam slice and modeled FTS aperture coupling vs frequency. . . 187

6.3 Cartoon diagrams of the instrument near field beam coupling with the FTS apertures. 188

6.4 FTS entrance aperture truncation of the near field beam. 190

6.5 The expected percent bandcenter shift caused by the aperture truncation. . 191

6.6 Example BA150 interferograms and spectra from A and B polarization for the three years of data. 192

6.7 Per detector symmetry coefficients for all measured BA150 detectors from the 2024 and 2025 FTS campaigns. 195

6.8 Aperture corrected bandcenters of all detectors that passed cuts in 2024. . 197

6.9 Scatter plot and histogram of per detector bandcenters from 2023 and 2024 FTS campaigns. 199

6.10 Coadded per detector aperture corrected spectra for BA150. 201

6.11 The multi-frequency BICEP/*Keck* maximum likelihood framework for multi-component parameter constraints. 205

6.12 The BK24 BB auto-spectra BICEP/*Keck* bandpowers plotted at the nominal and shifted bandcenter frequencies and the change in model expectation values. 206

6.13 Parameter constraints for pseudosim BK18/BK24/delensed-BK24 and realization to realization parameter biases from bandcenter uncertainties. . . 207

6.14 Maximum likelihood parameters and realization to realization comparison of parameters fit to pseudosims with and without shifted bandcenters. . . . 208

6.15 Corner plot of realization to realization shifts in maximum likelihood parameters. 210

6.16 Projected BICEP/*Keck* $\sigma(r)$ vs. bandcenter shift induced Δr. 213

7.1 Initial combined thicknesses in mils versus the final measured laminate thicknesses of two and three layer full scale windows. 216

7.2 Design corner plot for a hypothetical dual band 95/150 GHz receiver. . . . 221

7.3 Realization to realization comparison of maximum likelihood parameters fit to pseudosims with and without shifted bandcenters. 228

A.1 Examples of 'artifacts' woven into the HMPE weave. 230

A.2 Small scale creep of laminates with three different layers of HMPE. 232

A.3 Photo of the magnets placed below the edge of the template on the table. . 234

A.4 Photo of HMPE sheet over the template. 234

A.5 Photo of placing magnet on the annulus over a sheet of polyester (release material). 236

A.6 Photo of the HMPE roll around the template after a rough cut. 237

A.7 Photo of cutting polyester (release material) with scissors and a small magnet stack. 238

A.8 Photo of a close cut layer of HMPE. 239

A.9 Photo of the razor blade with a length of magnets to make a circular cut of an AR layer. 240

A.10 Photo of full laminate stack in large autoclave, with paper clips around the edge. 241

A.11 Photo of full laminate stack in large autoclave, with top release layer. . . . 242

A.12 Photo of full laminate stack in large autoclave after lamination bake. . . . 242

B.1 Picture of the interior of the small autoclave. 245

B.2 Pipe circuit of small autoclave. 246

B.3 Photo of the pipe circuit for the small autoclave on top of the oven. 247

B.4 Picture of the oven with the small autoclave set up inside. 249

B.5 Cutaway view of the large optics autoclave, with components color-coded and arrows pointed at small features. Large block arrows show direction of forced air flow. 250

B.6 Cartoon of a full stack in a thin window autoclave run. 251

B.7 Picture of the interior of the large autoclave at Step 14. 253

B.8 The patch in the top o-ring groove on the autoclave ring greased with high temperature E grease. 254

B.9 Pipe circuit of large autoclave. 256

List of Tables

2.1 Temperature sequence ultimate tensile strengths. 56

2.2 Pressure sequence of laminates ultimate tensile strengths. 61

2.3 Strength properties of Bulk HDPE, UHMWPE, and HMPE laminates. . . 63

2.4 Measured upper bound permeation rates from thin windows on two receivers and upper bound permeation rates associated with cryogenic issues. 73

2.5 Window surface temperature at different radial points on cold receivers at the South Pole. 77

3.1 Chebyshev single, double, and triple anti-reflection coat layer parameters on HDPE. 88

3.2 Generated AR coats for the current generation instruments in the BICEP/*Keck* collaboration. 93

3.3 Design, measured and modeled AR coat parameters for the BICEP3 and BA150 deployed thin windows. 101

3.4 Peel strengths of various kinds of DeWAL compressed together. 105

3.5 Design, measured and modeled AR coat parameters for the second and third BA220/270 thin windows. 113

4.1 Stokes parameters and polarization efficiency (ϵ) through a worst case laminate. 134

4.2 Relative offset in loading compared to 2017 of per season median per light detector P_{meas}. 150

5.1 Table of parameters in the BICEP3 loading spreadsheet (circa 2018), including calculated loading. 165

5.2 Measured NETs and insertion optical loads compared to the predicted NETs and optical loads from the old model. 166

5.3 Measured NETs and insertion optical loads compared to the predicted NETs and optical loads from the old model with yearly temperature changes. 173

5.4 Table of parameters in the modified BoloCalc model with the updated optical parameters, including calculated loading. 175

5.5 Measured NETs and insertion optical loads compared to the predicted NETs and optical loads from the new model. 176

5.6 Difference in loading compared to 2017 of both per season median per light detector P_{meas} and modeled P_{opt} differences compared to 2017. 177

5.7 Summary full instrument parameters produced by the new optical model for each year. 178

6.1 Bias on r from 2% shifts on individual BK24 bands. 209

D.1 Chebyshev single, double, and triple anti-reflection coat indexes on PTFE. 266

D.2 Chebyshev single, double, and triple anti-reflection coat indexes on HDPE. 267

D.3 Chebyshev single, double, and triple anti-reflection coat indexes on nylon. . 268

D.4 Chebyshev single, double, and triple anti-reflection coat indexes on quartz. 269

D.5 Chebyshev single, double, and triple anti-reflection coat indexes on alumina. 270

D.6 Chebyshev single, double, and triple anti-reflection coat indexes on silicon. 271

Acknowledgments

I have often joked (at least to myself) that it takes a village to teach a PhD, and that was never more true than with my doctorate. I must thank my latest teacher first: many thanks to my advisor John Kovac who has been a stalwart and earnest supporter in research, pushing me to be better. The rest of my committee were fantastically curious, and each had unique expertise that made this work stronger, so thanks to Doug Finkbeiner, Josh Grindlay and Scott Paine, with special thanks to my external examiner Kirit Karkare (also for many a helping hand in my last year).

I am lucky enough to have started and ended my schooling journey within the great state of Massachusetts, and have a plethora of teachers at all levels to thank. From the first few years of my schooling, thanks to Mrs. Entholt; from high school, thanks to Ms. Whitesel nee Fischer, Mr. Van Buren, and (of course) Mr. Dailey. I cannot thank my advisors from UMass enough: Min Yun introduced me to the very concept of instrumentation, and Grant Wilson followed through with a solid foundation in millimeter cryogenics. Kim Ward-Duong taught me python, and how to fill an observing trip with both work and Netflix.

The amazing people populating the labs I have worked must be thanked profusely! Everyone's patience with my constant bothering about their work was very appreciated: I must first thank my current office mates, Annie Polish and Brodi Elwood, who helped me debug many an issue and whom I occasionally helped in turn. Mainly they got to listen to me complain—sometimes about the pipeline or the philosophy of science, but a lot of the time about books. I can't forget my previous office mates either! James Cornelison gave me a lot of great advice, and many chances to pet his dog; Victoria DiTamosso

helped me stumble through most of my first year problemsets, my first steps into the Kovac lab, and was very gracious after I abandoned her to said lab. Many other Kovac Lab persons must be thanked; Clara Vergès for her advice on instrument calibration (and baking), Matthew Petroff for his diligence in making optics (and baking), Marion Dierickx for her steady head and initiative (sometimes about organizing baked goods), and Mike Echter for his exceptional help in saving the science-grade BA220 window (and for helping bake it). Other Kovac alums deserve thanks, such as Denis Barkats (who was the original mad genius behind the window project), David Goldfinger (for the discussions of time-multiplexing), Ben Schmidt (I hope the Japanese shoes work with your other outfits!), and Tyler St. Germaine (your thesis and FTS code was lifesaving).

My previous labmates in the Wilson Lab also need thanks! In particular, Nat Denigris was instrumental (ha) in teaching me many aspects of instrumentation, and in providing an amazing example of a graduate student to emulate! Steve showed me both what a project manager is and what it should be.

The BICEP/*Keck* Collaboration as a whole is first-class: there are too many collaborators to thank among them, but thanks to all. I also need to thank the students I've had the luck to supervise: this thesis would not be the same without the sweat, blood (not literally!), and tears (also not literally, hopefully) of Sage Crystian, Matthew Miller, and Destiny Santalucia. The wonderful receiver lab folks were a huge help in both experience, help and good humor: thanks to Robert Kimberk (for stopping feeding George [a turkey] when I asked), Edward Tuong (for many a receiver lab lunch talk), Paul Grimes (for lots of good optics advice), and Keara Carter (for being DeWAL layup champion).

The entire astronomy graduate student body deserves my undying gratitude for their generosity, kindness and support. Not only would I not have survived grad school without weekly social hours, I don't think I would have come to Harvard at all if not for the graduate community here. Special thanks goes out to my mentees, Erandi Chavez and Jea Redai Adams and my mentor Juliana Garcia-Mejia. And general thanks goes to every graduate student who let me ramble at them at social hour, a noncomprehensive list of which includes Mouza Almualla, Wasundrara Athukoralalage, Angus Beane, Olga Borodina, Vedant Chandra, Anna Fehr, Will Golay, Jesse Han, Alan Hsu, Ralf Konietzka, Claire Lamman, Matthew Leung, Marissa Maney, Tintin Nguyen, Theo O'Neill, Kevin Ortiz Ceballos, Misha Rashkovetskyi, Ana Sofia Uzsoy, Abby White, Rebecca Woody, and Qijia Zhou.

Of my wonderful friends, I must start with the fantastic Amanda MacLean, my best-friend-from-elementary-school, the first extrovert to farm me for engagement, the Foundation of the Group Chat, and the best adult I know my age–who is somehow two years older? Next is the marvelous Locke Patton; my best-friend-from-grad-school, the second extrovert to farm me for engagement, and the champion of my sanity. Others include Katelynn Szybillo-Vicare (whose many names never pass a spell check), Mike-and-Joey (for the watercolors), Ben-and-Rebecca (for the trip to the beach), and Mike Foley (for the jam, games, and puzzles).

My family also needs many thanks, but I find I lack the words. Thanks to Patrick Brugh for the many loads of laundry, and my grandfather for the quiet support through my schooling. This thesis is dedicated to all the Eiben women in my life, but in particular: my grandmother, my Aunt Doe, and my best-friend-from-birth (my mother). Thank you: for the laughter, and the strength, and the support.

For the roaring Eiben women.

Chapter 1

Introduction

1.1 The Universe and Everything In It

The universe started with a bang. The ripples from the very start of existence influence everything, on all scales, from the number distribution of atoms to spatial distribution of galaxy clusters. We live in the aftermath, trying to reconstruct the explosion from the echos. Thankfully, the universe is very big and the speed of light is very fast but finite, which allows us to peer back in time by looking further away.

We are interested in the very beginnings, however, so we look to the furthest and oldest reaches that we can see, the background of the entire universe. Precision measurement of a diffuse background source invites many challenges: everything along the line-of-sight is a foreground and source of noise. "Everything along the line-of-sight" includes the instrument itself, the atmosphere of our planet, the gas and dust in our galaxy, and other galaxies along the line-of-sight. Therefore, we must carefully design

our instruments to produce the smallest possible impact on the measurement, and design the experiment to be able to disentangle sources in the way.

The Cosmic Microwave Background (CMB) has been stretched by the expansion of the universe to be very cold (about three degrees above absolute zero) which prompts us to build very cold instruments to observe it. The design of these instruments is exacting and peculiar. Lenses are made of plastic and optical coatings are macroscopic. Nearly every material property (optical or mechanical) matters, and often must be determined by the instrumentalist.

Our instruments (the only line-of-sight component we have full control over) are designed to maximize our sensitivity to millimeter light with well-defined beam patterns on the sky. Therefore we must increase our optical throughput as much as possible, and avoid beam distortions from internal optical reflections. Individual optical components' impact on throughput is reduced by minimizing absorption losses within the optic, accomplished by reducing component thicknesses or reducing material absorption coefficients. Carefully designed anti-reflection (AR) coats reduce reflective losses at each interface, further reducing transmission loss and, critically, minimizing internal reflections.

To disentangle astrophysical sources we observe at multiple frequencies. The galactic foregrounds have different spectral emission than the CMB, and measuring the sky at multiple frequencies (carefully avoiding atmospheric absorption lines) allows us to separate them from the CMB. However, if we are to accurately subtract one component from another we need to ensure we know what the frequency responses of the instruments. We utilize a Fourier Transform Spectrometer (FTS) to input a signal at a

known frequency to find the spectral response of our detectors, which we then can use to both measure and model the universe.

But before we dive into the instruments' design and calibration, let us explore the universe we are attempting to measure.

1.1.1 The Expanding Universe

Humanity first found evidence of the Big Bang in the motion of galaxies. When Hubble measured the relative speed along the line-of-sight of the galaxies around us, he not only found that galaxies move at a prodigious pace but nearly all of them were moving away from us: the further away the galaxy, the faster the speed (Hubble 1929). This is very simply defined by:

$$v = H_0 d \tag{1.1}$$

where v is measured radial velocity, d is the distance to the galaxy. The slope of this change in speed with distance was named after its discoverer, and therefore called the Hubble constant (at the current time), H_0. But we also trust in the Copernican principle that we are not located in a "special" place. We therefore assume that the universe is isotropic (having no preferred direction) and homogeneous (having no preferred location). The fact that further galaxies move away faster in all directions has a natural—if mind-boggling—conclusion; the universe is isotropically expanding.

It is natural, having discovered this, that one may be concerned about the structural integrity of the Earth on which one stands. If the largest structures in the universe are being pushed apart, what chance does our small rocky marble have? To alleviate these

concerns, let us define the scales and the metrics that we can use to describe and predict this expansion.

Energy and matter are fundamentally entwined (recall $E = mc^2$), and, in turn, both influence the structure and dynamics of spacetime[1]. This is elegantly described by the Einstein field equations:

$$G_{\mu\nu} - \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \tag{1.2}$$

where the terms on the left hand side $G_{\mu\nu}$ (the Einstein tensor) and $g_{\mu\nu}$ (the metric tensor) describe the curvature and geometry of spacetime, respectively, and Λ is the cosmological constant; on the right hand side $T_{\mu\nu}$ (the stress-energy tensor) describes the energy and momentum attributed to energy/matter, and the constants G and c are the gravitational constant and the speed of light in a vacuum, respectively. Essentially, this equation describes how matter and energy distort spacetime which generates gravitational fields, a paradigm commonly referred to as general relativity (Einstein 1916).

We can define a metric (essentially, an equation to describe the path in spacetime between points) to describe the shape of spacetime. Assuming that the universe is isotropic and homogeneous (trusting still in the Copernican principle), we can define the spacetime metric ds^2 with the Friedmann-Robertson-Walker (FRW) metric:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = -c^2 dt^2 + a(t)^2 \left[\frac{dr^2}{1 - Kr^2} + r^2 \left(d\theta^2 + \sin^2(\theta) d\phi^2 \right) \right] \tag{1.3}$$

[1]One of Einstein's many innovations was the realization that the universe is four dimensional: three spatial dimensions and one time dimension. Together this dimensional fabric is referred to as *spacetime*, as opposed to one or the other.

where t is the temporal coordinate and r, θ, and ϕ are the spatial coordinates. We also have a dimensionless scale factor a, and a representation of the curvature of space, K. If the universe is flat, $K = 0$.

On universal scales, the energy-density distribution that shapes spacetime can be modeled as a perfect fluid: $T_{\mu\nu}$ (the stress-energy tensor) therefore becomes a function of pressure, p, and energy density, ρ. The combination of Equations 1.2 and 1.3 become:

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

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

where $H(t)$ is the Hubble parameter, which varies in time with the dimensionless scale factor (at current time the Hubble parameter is H_0). If we continue to assume that the energy-density is a perfect fluid, then

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

The equation of state of some general energy component is $p = w\rho c^2$, which we can use with Eq. 1.6 to find a relationship between an energy density and scale factor:

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

Matter[2] density scales with volume, so $\rho_m \propto a^{-3}$ and therefore $w = 0$; radiation density scales like $\rho_r \propto a^{-4}$ ($w = \frac{1}{3}$), as the wavelength evolves along with the scale factor.

[2]To be clear, matter in this case refers to both regular visible baryonic matter (matter made up of protons, neutrons, electrons, etc.) and the mysterious invisible 'dark' matter, of which the constituent components remain unknown. There is ample evidence for invisible matter that only interacts gravitationally with other energies, notably in the measured rotation curves of galaxies (Rubin & Ford 1970) and (as we shall discuss shortly) the early universe (Planck Collaboration et al. 2020a)

When the scale factor was small in the early universe, the combined energy density of the universe (and thus the density driving the expansion of the universe) was dominated by radiation, as radiation energy density scales with a^{-4}. After a time, the expansion was dominated by matter. Relatively recently (cosmologically), we have found that the acceleration has accelerated, indicating that some energy density remains constant regardless of scale factor, and is now dominating the total energy density ($\rho_\Lambda \propto$ constant, $w = -1$). We call this mysterious energy density 'dark' energy, because cosmologists are unoriginal.

Back to our existential concern about the Earth's ability to withstand the expansion of the universe. Firstly, the assumed homogeneity required for expansion, as defined here, needs a homogeneous universe on the largest scales, above 100 Mpc (the Milky Way Galaxy, by contrast, is only about 30 kpc across). Secondly, the density required to produce expansion, as emphasized from the Friedmann equations, is very low (equivalent to 1 proton in a liter) (Ryden 2016). Not a lot of matter or energy is required to cause this expansion, and therefore it also is easily overridden by local overdensities' gravity. The hyperdense Earth (by universe standards!) easily produces enough gravitational force to hold itself together, thankfully. Though, *how* the Earth managed to become so hyperdense is an interesting question, which we will return to later.

It is also worthwhile, at this point, to clearly outline all the assumptions that have led us here. They are:

1. The universe is isotropic and homogeneous, and (implicitly assumed so far) the physics we find on Earth remain the same throughout the entire universe.

2. General relativity adequately describes the relation between spacetime and energy

densities (recall Equation 1.2).

3. Each of the individual energy densities resembles a perfect fluid on universal scales. The total energy density of the universe consists of matter (both baryonic and dark), radiation (both photons and neutrinos) and some dark energy component that acts like the energy density of the vacuum (what we have been carrying through in our equations so far as Λ).

These assumptions underpin the ΛCDM (Λ cold dark matter) model, which is currently the most widely accepted model of the universe. As we shall see, the ΛCDM model is the simplest model that best describes the universe at a wide variety of scales, which is quite the feat.

However, the Friedmann equations do lead us to our first oddity with the ΛCDM model of the universe: the overall shape of the universe. At different overall densities the curvature term K will flip. We can use Eq. 1.4 to define a critical density at which the overall shape of the universe flips by setting $K = 0$ (and, temporarily, $\Lambda = 0$ as well) and solving to find:

$$\rho_{crit} = \frac{3H^2}{8\pi G}. \tag{1.8}$$

If the universe is closed (spacetime is shaped like a hypersphere) and will eventually reverse expansion and collapse, $\rho > \rho_{crit}$; if it is open (spacetime is shaped like a hyperboloid) and expanding indefinitely, $\rho < \rho_{crit}$. If, by some happenstance, the universe is flat, $\rho = \rho_{crit}$. For any given energy density, we can define a $\Omega_i = \rho_i/\rho_{crit}$. We can then rewrite Eq. 1.4 in the density parameters like

$$\frac{H^2}{H_0^2} = \Omega_{0,r}a^{-4} + \Omega_{0,m}a^{-3} + \Omega_{0,k}a^{-2} + \Omega_{0,\Lambda}, \tag{1.9}$$

where $\Omega_{0,i}$ denotes the current energy density of the given energy component (i).

The actual total energy density of the universe is incredibly, exceptionally, unnervingly close to the critical density. Recent measurements put $\Omega_k = 0.0007 \pm 0.0019$, and the rest of the energy densities subsequently sum to unity (Planck Collaboration et al. 2020a). This is referred to as the *Flatness Problem*: there is no reason within the standard model of cosmology for the curvature of the Universe to balance so perfectly on the critical density knife edge. Let us not forget the scaling of the each of the component densities has also changed with time, and the spatial curvature ($\Omega_k \propto a^{-2}$, $a \propto t^{5/3}$) evolves slower than radiation ($\Omega_r \propto a^{-4}$, $a \propto t^{1/2}$) or matter ($\Omega_m \propto a^{-3}$, $a \propto t^{2/3}$): if the universe is exceptionally flat now, it was near impossibly flat in the distant past.

Having now determined that the universe is expanding and the mechanisms driving that expansion, the next logical leap is to reverse the clock, and imagine what the universe must have looked like over the course of its life. If the universe is expanding and therefore becoming less dense with time, then logic dictates it must have been more dense in the past. At some point, everything in the universe must have been so close and compacted that baryonic matter was ionized into a plasma much like the interior of the Sun. Charged particles (electrons and protons) prevented the free motion of photons through Thompson and Coulomb scattering, creating a fully coupled photon-baryon fluid. As the universe expanded, the matter adiabatically cooled enough to allow the baryons to form neutral atoms[3], freeing the photons, which have been traveling mostly unobstructed ever since. As the universe expanded underneath them, the spectra of the

[3]This is period is known as 'recombination.' I personally dislike the term, as this was the first time chronologically that baryons combined into neutral atoms. I will use the term recombination to describe this period going forward, but consider this my objection to having to do so.

photons (originally sourced from a hot plasma) was redshifted down to a black-body at a much lower temperature[4].

1.1.2 The Cosmic Background

In 1964, Arno Penzias and Robert Wilson were attempting make absolute intensity measurements of radio sources in the galaxy when they measured an excess antenna temperature of approximately 3.5 ± 1.5 K (Penzias & Wilson 1965). This excess antenna temperature was constant over the entire sky, and could not be explained by instrument noise; Penzias and Wilson had measured a uniform cosmic background. In a companion paper, Dicke, Peebles, Roll and Wilkinson came to the conclusion that this background was the relic radiation from the hot plasma of the early universe (Dicke et al. 1965). After the photons were freed, they continued to travel in the direction they were last scattered. When we look in any given direction at the right frequencies, we see the photons scattered in our direction from that location during recombination (at $z \sim 1100$, or 380,000 years after the Big Bang). As recombination occurred relatively rapidly, these photons generate a spherical surface that bounds the observable universe known as the surface of last scattering, which is the source of the cosmic microwave background (CMB).

We have since constrained the temperature of the background to be a bit lower than Penzias and Wilson's initial measurement. Measuring the temperature of the CMB

[4]Recall that redshift is a function of frequency that goes like $z = (f_{emit} - f_{obs})/f_{obs}$. The peak of a blackbody at 3000 K is at 3×10^{14} Hz, while the 3 K peak is at 3×10^{11} Hz; therefore the redshift of these photons is $z \approx 1000$.

requires multi-frequency component separation to disentangle the CMB from foreground sources that produce similar power. Galactic sources of emission, such as thermal emission from dust or synchrotron emission from relativistic photons in the magnetic field, can be seen in Figure 1.1 at high and low frequencies respectively (Fixsen 2009; Planck Collaboration et al. 2020a). We will discuss component separation further in §1.2.3. The current best measure of the CMB temperature is 2.7260 ± 0.0013 K; the CMB is the most ideal black-body measured in nature (Fixsen 2009). Too ideal. This is our next issue: the *Horizon Problem*.

For two points of the universal primordial plasma to be the same temperature they must have been in causal contact; essentially, information must have been able to be shared between the two points, and therefore they must have been within the same light 'horizon' of each other. As light has a speed limit and the universe has an age, this limits which portions of our background should have been within each others' horizons at the time of recombination[5]. The furthest distance that a photon can travel in conformal time $d\tau = dt/a$ is:

$$\eta = \int \frac{c\ dt}{a(t)} = \int \frac{c}{Ha} d\ln(a) \tag{1.10}$$

where the co-moving Hubble distance is $1/Ha$. Vast swaths of the CMB should have been able to diverge much further in temperature then they actually did: at the time of recombination the horizon scale would correspond with 2° angular separation on the sky today. Positions on the CMB further apart have no reason to be the same temperature.

[5]The CMB itself is a light horizon: it has taken the scattered light from recombination this long to reach us, and therefore we cannot directly gain information from locales further (i.e., earlier) than the surface of last scattering. This does not mean that we cannot infer the state of the universe at earlier times than the CMB, but those inferences must originate from times during or after recombination.

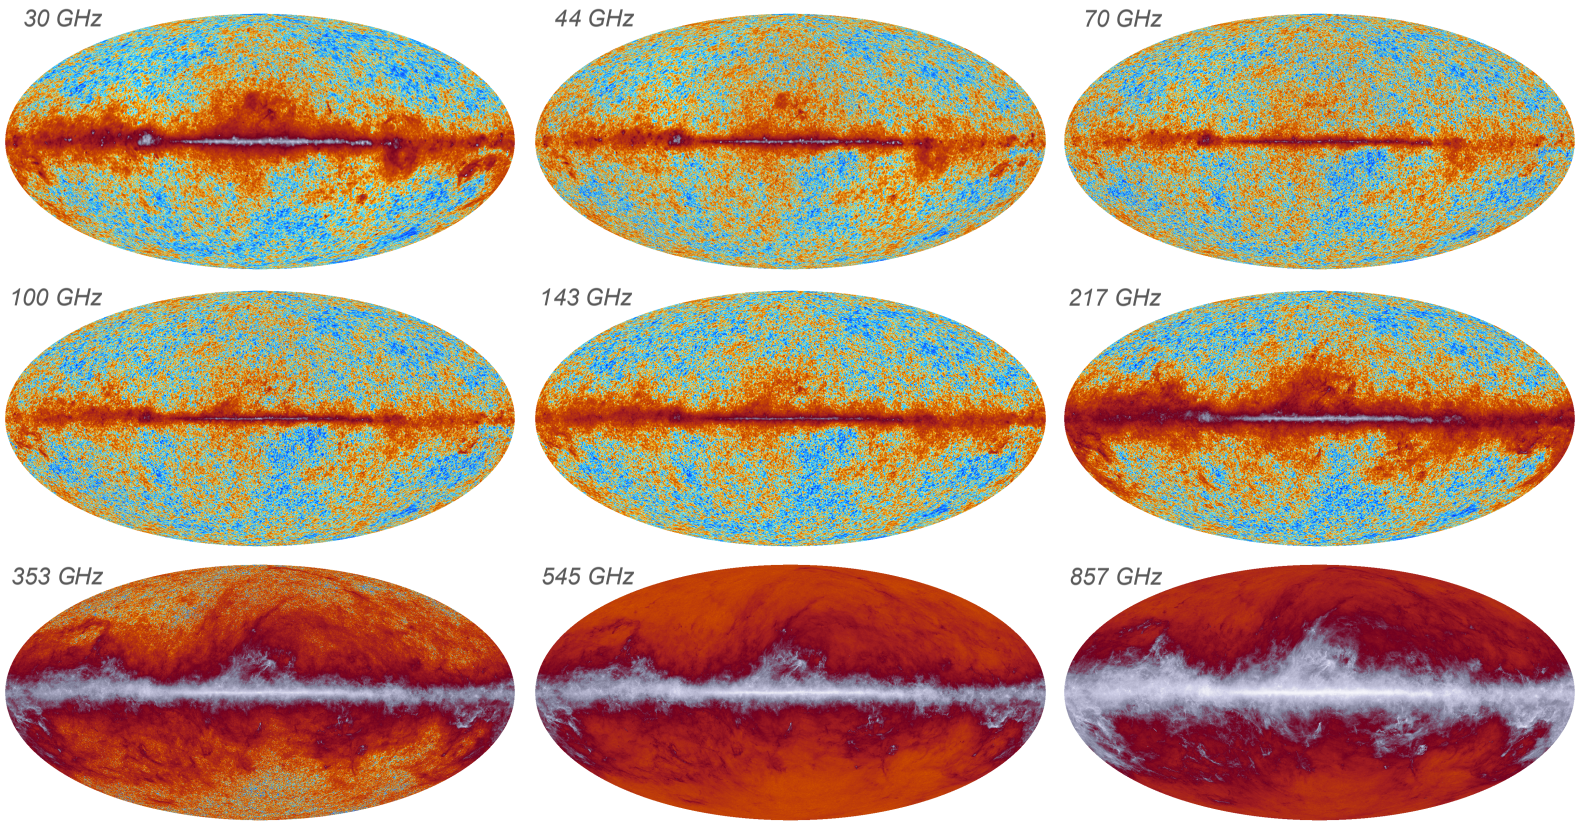


Figure 1.1: *Planck* all-sky maps at different frequencies (from 30 GHz to 857 GHz), after removal of a common dipole component. At the lowest frequencies synchrotron emission from the galaxy dominates, at the highest frequencies galactic dust emission dominates (the units also flip from $\mu\mathrm{K}_{CMB}$ to kJy/sr). The CMB temperature anisotropies are visible in the mid frequencies behind the galactic emissions. Figure from Planck Collaboration et al. (2020a).

There are small deviations from the average temperature within the CMB, which we call temperature anisotropies. The photon-baryon fluid had two competing forces acting on it: gravitational attraction pulling the photons and baryons together (compression), and photon pressure pushing the photons and baryons apart (rarefication). This led to acoustic oscillations in the fluid. When recombination occurred and the universe became neutral, the photons kept moving with the information about the temperature conditions from their last scatter. Some acoustic modes were maximally rarefied for the first time, and therefore produced cold spots, while some acoustic modes were maximally compressed for the first time, and therefore produced hot spots. Other acoustic oscillations were undergoing their second round of maximum rarefication/compression, producing hot and cold spots half the size of the larger modes, and so on. The temperature anisotropies in the CMB are visible in the low frequency all sky maps (less than 353 GHz) in Figure 1.1.

To characterize these temperature anisotropies ($\Delta T(\hat{n})$) as a function of a direction vector pointing at the celestial sphere ($\hat{n}$) we decompose them into spherical harmonics $Y_{\ell m}(\hat{n})$:

$$\begin{aligned} \Delta T(\hat{n}) &= \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} = a_{\ell m} Y_{\ell m}(\hat{n}) \\ a_{\ell m} &\equiv \int_{\Omega} \Delta T(\hat{n}) Y^*_{\ell m}(\hat{n}) d\hat{n}. \end{aligned} \tag{1.11}$$

The star denotes complex conjugation. The statistical distribution of the $a_{\ell m}$'s is assumed to be Gaussian and isotropic[6], so the variance between them can be described

[6]Primordial non-Gaussianity has been constrained to be very small, so this assumption holds (Planck Collaboration et al. 2020b).

with angular power spectrum[7], C_ℓ:

$$\langle a_{\ell m} a'_{\ell m} \rangle = C_\ell \delta_{\ell \ell'} \delta_{m m'}. \tag{1.12}$$

As $\Delta T(\hat{n})$ is the change from the mean temperature of the CMB, the average $\langle a_{\ell m} \rangle$ is zero. Because we only have the one universe and one CMB, we only get one set of modes to observe: this limits the amount of information we can gather, as we get a limited number of draws from the distribution[8] across the entire sky. This is the cosmic variance limit, which is calculated for a given ℓ as

$$\frac{\Delta C_\ell}{C_\ell} = \sqrt{\frac{2}{2\ell + 1}}. \tag{1.13}$$

We evaluate the power spectrum C_ℓ by "crossing" maps' $a_{\ell m}$ (taking a finite sum over the modes instead of the ensemble average defined in Eq. 1.12), like so:

$$C_\ell^{XY} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} a_{\ell m}^{X*} a_{\ell m}^{Y} \tag{1.14}$$

where X and Y are the respective maps. We can cross a map with itself to generate an "auto-spectrum" or with a separate map to make a "cross-spectrum". The maps could be of temperature or polarization anisotropies (discussed shortly), or different frequency maps, or be from different experiments (we commonly generate cross-spectra with our maps and *Planck* or WMAP maps, for example (BICEP/Keck Collaboration et al. 2021)). It is also common to plot or discuss the $\mathcal{D}_\ell \equiv \ell(\ell+1)C_\ell/2\pi$ power spectrum (like in Figure 1.4), as it is approximately flat for low ℓ.

Light does not just hold information in its intensity (i.e. temperature), but also in its polarization. We can gain additional information from the CMB using the polarization

[7]We will tend to refer to the angular power spectrum as just a power spectrum from now on.

[8]At $\ell = 2$ we only get 5 a_{2m}, at $\ell = 3$ we get 7 a_{3m}, at $\ell = 4$ we get 9 a_{4m}, and so on.

states (particularly useful due to the cosmic variance limit), which we can define as the Stokes parameters for a light wave traveling in the $+\hat{z}$ direction like (Stokes 1851):

$$\begin{aligned} I &= \langle E_x^2 \rangle + \langle E_y^2 \rangle \\ Q &= \langle E_x^2 \rangle - \langle E_y^2 \rangle \\ U &= 2\langle E_x E_y \rangle \cos\phi \\ V &= 2\langle E_x E_y \rangle \sin\phi \end{aligned} \tag{1.15}$$

where ϕ is the phase offset between the complex electric field amplitudes aligned along the Cartesian axes (E_x and E_y), and $\langle E \rangle$ is the average of the field over time. The Stokes parameters denote the total intensity I of the electric field and the polarization state of the field in linear polarizations Q and U (which are offset by 45° to each other) and circular polarization V. We do not expect any circular polarization to be generated in the early universe, so we shall ignore it for cosmological polarization studies going forward[9]. We do, however, expect linear polarization to be produced by the primordial plasma.

An electron in the plasma will scatter (via either Thompson or Compton scattering) an incident unpolarized photon into a linearly polarized photon with polarization perpendicular to the line-of-sight. However, if the electron is in a uniform temperature bath (a monopole) the incident unpolarized photons all have the same power and therefore the scattered radiation is unpolarized. If the electron sits in a temperature dipole (a gradient between hot and cold), then the polarization will also average out, as the net power from the hot/cold directions will average to be the same as the power

[9]If our instrument optical components transform incident Q or U radiation into V, *then* we care about V, as that decreases our instrument polarization efficiency. More on this in §4.1.2.

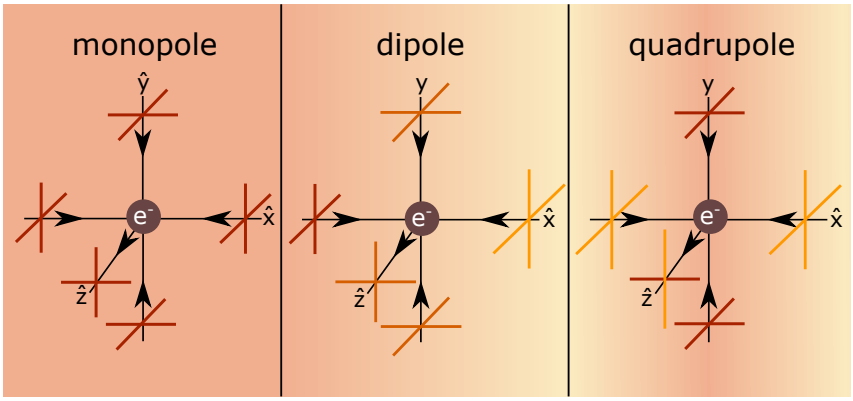


Figure 1.2: A cartoon representation of an electron sitting in a monopole (left), dipole (center), and quadrupole (right) temperature field. Incident unpolarized light approaches the electron from the $\hat{x}$ and $\hat{y}$ directions and get scattered in the $\hat{z}$ (observer) direction. Only the quadrupole produces net linear polarization. Cartoon based on those in Hu & Dodelson (2002).

from the orthogonal uniform temperature direction (i.e., the direction at the average temperature where the electron sits). If the electron sits in a temperature quadrupole, however, then there will be net linear polarization. Cold spots in the arbitrary $\hat{y}$ direction would produce scattered photons aligned with the $\hat{x}$ direction and the orthogonal hot spots $\hat{x}$ direction would produce more energetic (higher amplitude) scattered photons aligned with the $\hat{y}$ direction, producing linear polarization (Q) (Rees 1968). A cartoon representation of each of the temperature fields and subsequent observed polarization (or lack thereof) are shown in Figure 1.2.

However, the linear Stokes polarization states (Q and U) are not rotationally invariant as they are dependent on which arbitrary orientation the Cartesian coordinates are placed. It is generally more useful to work in an invariant basis so that we can easily

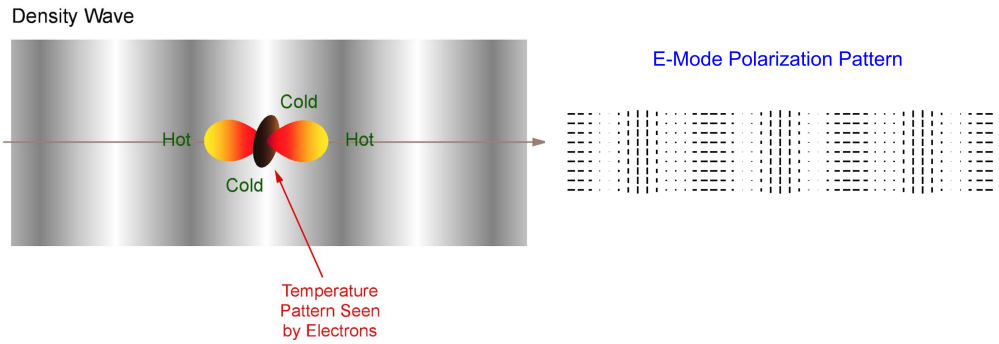


Figure 1.3: Cartoon representation of quadrupole temperature distribution seen by an electron in a density wave. Density waves produce 'E-mode' (curl-free) polarization in the CMB. Cartoon from Kovac (2004).

cross maps. We can express a vector of linear polarization Q and U like

$$\begin{bmatrix} Q(\hat{n}) \\ U(\hat{n}) \end{bmatrix} = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} = a^E_{\ell m} Y^E_{\ell m}(\hat{n}) + a^B_{\ell m} Y^B_{\ell m}(\hat{n}), \tag{1.16}$$

where $Y^E_{\ell m}(\hat{n})$ and $Y^B_{\ell m}(\hat{n})$ are the orthonormal basis functions for what we term as E and B modes. E-modes are "curl-free" (scalar) polarization fields, while B-modes are "gradient-free" (pseudoscalar) polarization fields. E-modes are generated in the CMB from the same acoustic oscillations that source the temperature anisotropies. However, at earlier times the quadrupole temperature anisotropies are suppressed by the scattered photons sharing energy between electrons in the plasma, equilibrating the temperatures. Only when the optical depth rapidly decreased during recombination were the quadrupole temperature anisotropies possible. The density waves produce a temperature quadrupole like that shown in the left side of Figure 1.3, resulting in the polarization field on the right of the figure.

There are four non-zero power spectra we can now find for various anisotropies[10]:

$$\begin{aligned} C_\ell^{TT} &= \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} a_{\ell m}^{T*} a_{\ell m}^{T} \\ C_\ell^{EE} &= \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} a_{\ell m}^{E*} a_{\ell m}^{E} \\ C_\ell^{BB} &= \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} a_{\ell m}^{B*} a_{\ell m}^{B} \\ C_\ell^{TE} &= \frac{1}{2\ell+1}\sum_{m=-\ell}^{\ell} a_{\ell m}^{T*} a_{\ell m}^{E} \end{aligned} \tag{1.17}$$

The TT power spectrum alone can tell us a lot about the conditions in the early universe: for example, the angular scale of the first peak puts constraints on the curvature of the universe (and the energy density Ω_Λ of dark energy), the ratio between the first and the second peak puts constraints on the ratio of regular matter to dark matter, etc. The TE and EE power spectra each put similar constraints on the early universe conditions (with an expected phase offset), but are temporally restricted to the relatively short period of recombination, as the quadrupole temperature anisotropies required to source E-modes are only possible when the optical depth became sufficiently low (Hu & Dodelson 2002).

The ΛCDM model fits the TT, TE, and EE power spectra very well, as seen in Figure 1.4. However, it does not describe the source of the perturbations nor how matter managed to get close enough to condense (recall that the ΛCDM model assumes

[10] TB and EB are expected to be zero, as the sky should be invariant under parity transformation in ΛCDM, so any cross between a scalar ($a_{\ell m}^T$ or $a_{\ell m}^E$) and a pseudoscalar ($a_{\ell m}^B$) is zero. Cosmic birefringence (or some other parity violating effect) can generate TB or EB correlations: see Cornelison (2023) or BICEP/Keck Collaboration et al. (2024b) for a more in-depth overview of these effects and recent constraints.

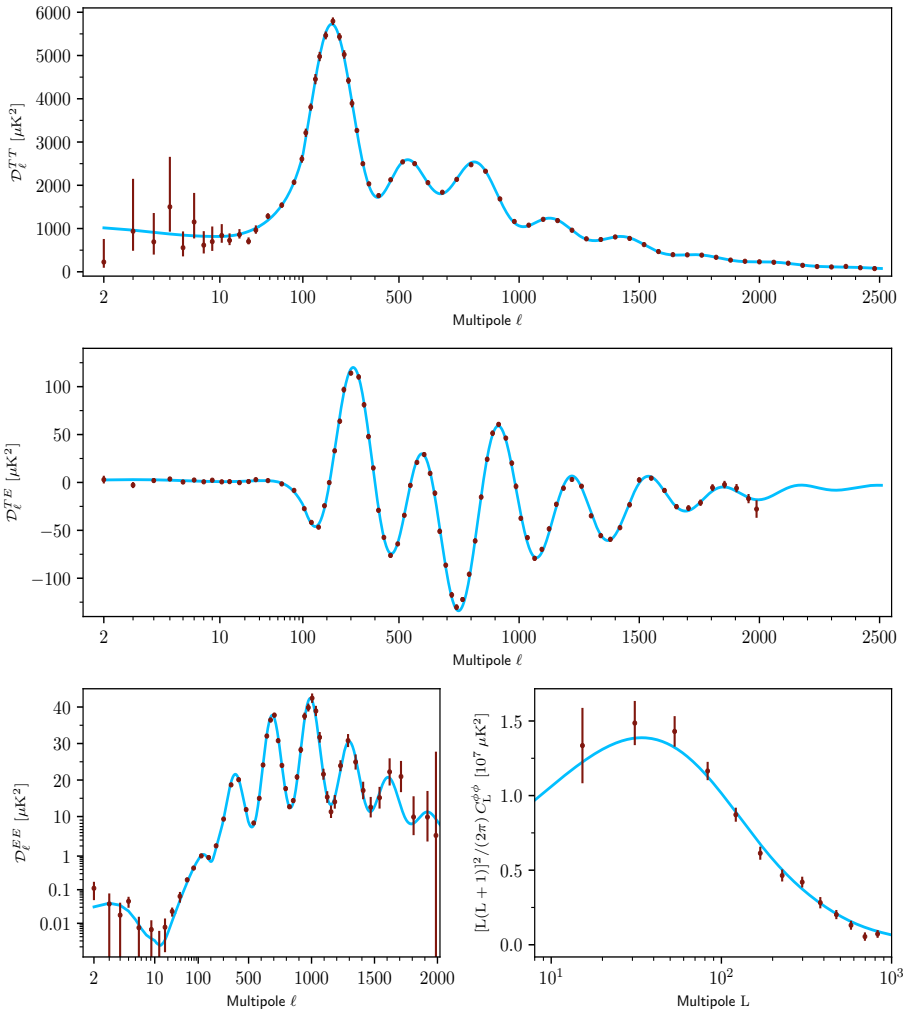


Figure 1.4: ΛCDM fit to the *Planck* satellite power spectra. Top panel is the TT auto-spectra, middle are the TE cross-spectra, bottom left are the EE auto-spectra and bottom right are the foreground gravitational lensing spectrum. Figure from Planck Collaboration et al. (2020a).

a homogeneous and isotropic universe with very low overall density, and our hyperdense Earth; every planet, star, and galaxy had to somehow get matter close enough to gravitationally condense), only their eventual effects on the observable universe. This is our final problem with ΛCDM: the *Perturbation Problem*.

1.1.3 The Initial Conditions

Let us quickly review the three main problems we have identified with the ΛCDM model of the universe:

1. *The Flatness Problem:* the energy density of the universe is very close to the critical density, resulting in a flat, Euclidean universal spacetime (Planck Collaboration et al. 2020a). There should have been infinite chances for the Universe to be open or closed. How is the universe so close to flat?

2. *The Horizon Problem:* the cosmic microwave background (relic radiation from the early universe) is 2.7260 ± 0.0013 K, but most points in the CMB should not have been within each other's light horizons, and therefore should not have been able to equilibrate (Fixsen 2009). Why is the CMB so uniform?

3. *The Perturbation Problem:* the source of perturbations in CMB is not explained by the ΛCDM model. Where did the perturbations come from?

A possible solution to these issues is an extension to ΛCDM called Inflation (first proposed by Starobinsky (1980), Guth (1981) and Linde (1982)). If there was a period of rapid, superluminal (faster than the speed of light), exponential expansion prior to the standard period of expansion predicted by ΛCDM, each of the issues with ΛCDM model could be explained: any initial curvature would be "smoothed out" by the expansion, as a curved surface *appears* flat on small scales; during inflation the smallest scales were expanded past the light horizon, freezing the current observable universe at the same temperature; random quantum fluctuations are scaled up to seed the density perturbations, and neatly explain the perturbations' Gaussian distribution and near

scale invariance.

The scale of the inflation required to solve these problems requires some explanation. Superhorizon modes—the portions of space that should never have been in causal contact in ΛCDM—must be brought within the light horizon, somehow, to equilibrate and solve the *Horizon Problem*. Recall the comoving Hubble radius (the maximum distance that a photon can travel in conformal time), Eq. 1.10. If we take the derivative of Eq. 1.10, we find:

$$\frac{d}{dt}(Ha)^{-1} = \frac{d}{dt}(\dot{a})^{-1} = -\frac{\ddot{a}}{(\dot{a})^2} \tag{1.18}$$

where $\dot{a}$ and $\ddot{a}$ are the derivative and second derivative of the scale factor a.

If $\frac{d}{dt}(Ha)^{-1} < 0$ (the Hubble radius shrinks, bringing superhorizon modes within the light horizon), then $\ddot{a} > 0$. Eq. 1.5 therefore tells us that this requires a negative pressure (not unlike dark energy) like $p < \rho c^2/3$. If the equation of state of the inflationary expansion is $p = -\rho c^2$, then the scale factor of inflation has to be $a \propto e^{Ht}$. Depending on the energy scale of inflation, between 40 to 60 e-fold (increases by a factor of e) expansions are required during inflation in order to bring the correct amount of superhorizon modes within the comoving Hubble radius, and additionally explain the spectral tilt n_s (described further below) (Kamionkowski & Kovetz 2016).

For the *Flatness Problem*, we rearrange the Friedmann equation (Eq. 1.9) by setting the energy density $\Omega_k = \Omega - 1$ to get:

$$\Omega_k = \frac{Kc^2}{H^2 a^2} \tag{1.19}$$

The same dark energy-like scale factor of $a \propto e^{Ht}$ reduces the above relationship to $\Omega_k \propto e^{-2Ht}$; with an expansion of 60 e-folds this becomes e^{-120}, which is very flat.

Quantum mechanical fluctuations (which could be thought of as virtual particles being separated by the rapid expansion of the universe before they can annihilate) generate both scalar and tensor perturbations in the spacetime metric. The scalar perturbations, as before, correspond to density fluctuations in the matter that are expected to be described by a power spectrum:

$$P_s(k) = \Delta_s \left(\frac{k}{k_{s*}} \right)^{n_s - 1} \tag{1.20}$$

where k is the wavenumber in Fourier space, and the amplitude Δ_s is defined at the pivot scale k_{s*}. Because inflation ended, n_s (the power law index, or the spectral tilt) is expected to be close to but not exactly one. Recent constraints put $n_s = 0.9626 \pm 0.0057$ to the 68% confidence level (Planck Collaboration et al. 2020c).

Tensor perturbations (gravitational waves) are expected to also be described by a power spectrum:

$$P_t(k) = \Delta_t \left(\frac{k}{k_{t*}} \right)^{n_t} \tag{1.21}$$

where the amplitude Δ_t is defined at the pivot scale k_{t*}. The tensor power law is expected to be zero[11], but as the tensor perturbations have not yet been found this has not been measured. We can define the tensor to scalar ratio r by the power spectrum amplitudes[12]:

$$r \equiv \frac{\Delta_t}{\Delta_s}. \tag{1.22}$$

[11]The discrepant use of $n_s \approx 1$ and $n_t \approx 0$ to set scale invariance is standard in literature.

[12]It is common to reference a pivot scale for the tensor modes in the subscript of r; so $r_{0.05}$ would have a pivot scale of 0.05 Mpc^{-1}.

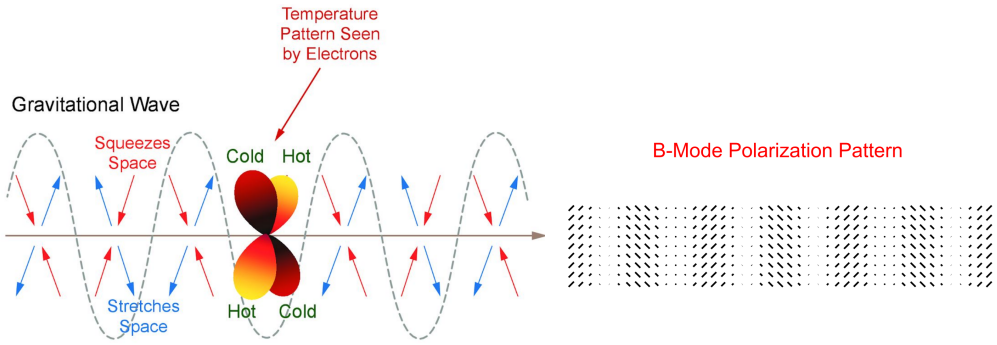


Figure 1.5: Cartoon representation of quadrupole temperature distribution seen by an electron in a gravitational wave. Gravitational waves produce 'B-mode' (curl) polarization in the CMB. Figure source: Kovac (2004).

The relative scale of the tensor perturbations is directly related to the scale of the primordial gravitational waves generated by inflation. These gravitational waves should produce another form of quadrupole temperature anisotropy in the primordial plasma, as shown in Figure 1.5. Critically, the tensor perturbations source a B-mode polarization pattern within the CMB: a potentially observable direct tracer of cosmic inflation.

The primordial B-mode power spectrum is expected to peak at around ℓ of 100, as seen in Figure 1.6. Another important thing to note from Figure 1.6 is the expected difference in power between the different anisotropies. If the temperature anisotropies fluctuations are at the mK level (approximately 10^3 μK^2 in Figure 1.6), the E-mode anisotropies are roughly three (at some ℓs more like five) orders of magnitude below the power of the temperature anisotropies (roughly the μK^2 level). The primordial B-mode anisotropies are at least an order of magnitude below even that. Sensitivity then becomes the name of the game; precision measurement down to the *nano-kelvin* level is required to detect the primordial B-modes.

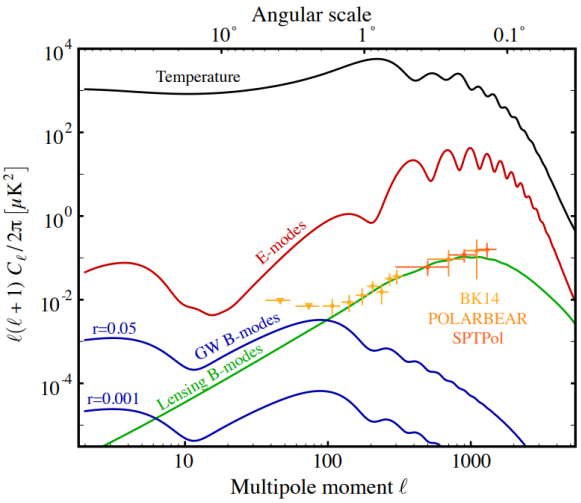


Figure 1.6: ΛCDM predictions of primordial temperature (black), E-mode (red), and B-mode (blue) angular power spectra. The predicted B-modes for two different representative values of r are shown at $r = 0.05$ and $r = 0.001$. The predicted foreground gravitation lensed B-modes (sourced from the primordial E-modes) are in green. Measured B-mode angular power from BICEP2/*Keck* Array (yellow), POLARBEAR (orange), and SPTPol (dark orange) are shown. Note the expected differences in angular power between the different anisotropies. Figure from Abazajian et al. (2019).

1.2 The BICEP/*Keck* Experiment

BICEP is a long running CMB experiment located at the geographic South Pole, with the goal of detecting B-modes in the CMB sourced from primordial gravitational waves. Increased sensitivity for detecting primordial B-modes is accomplished in two ways: either by increasing the amount of power captured by the instrument (by increasing number of detectors, which occasionally requires increasing the aperture size), or by decreasing noise on detectors. The experiment has had three "stages" so far. Each of the stages saw an order of magnitude increase in detector counts, and occasionally

larger apertures: BICEP1 had a 250 mm aperture for 49 detector pairs at 95, 150, and 220 GHz (Takahashi et al. 2010); BICEP2 (150 GHz) and the *Keck Array* of five receivers (150 GHz, later adding 95, 220 and 270 GHz) had 256 detector pairs each for their 264 mm apertures (BICEP2 Collaboration et al. 2015b; Kernasovskiy et al. 2012; Staniszewski et al. 2012); the currently operating BICEP3 (95 GHz) has 1200 detector pairs in a 520 mm aperture (BICEP/Keck Collaboration et al. 2022). A single BICEP Array receiver—the set of receivers currently being deployed which range in frequency from 30 to 300 GHz—has between 100 and 4000 detector pairs (depending on the designed frequency) for a 560 mm aperture (Moncelsi et al. 2020; Schillaci et al. 2023; BICEP/Keck Collaboration et al. 2024c).

The combined 15 years of experience mapping the CMB at various frequencies has led to a robust set of calibration and analysis practices that I leveraged for this thesis. While I will primarily focus on results from the current generation of BICEP/*Keck* receivers (BICEP3 and BICEP Array) for the rest of this thesis, many of the tests I have adapted and utilized were originally developed for previous generations of BICEP/*Keck*. The relatively long history of BICEP/*Keck* has also resulted in a genuine engineering prototyping and development cycle. This is relatively unique for astronomical instruments, which tend towards highly technical and bespoke one-off devices.

The environment of the polar plateau is uniquely suited for a high sensitivity millimeter experiment. The geographic South Pole is at an altitude of 2.8 km (9,200 ft), but feels like 3.4 km (11,000 ft) due to centripetal force moving the atmosphere away from the poles. The shape of the Antarctic continent also works to make the interior among the driest places in the world. The lack of daily diurnal cycles provides additional atmospheric stability, particularly in the winter, when the Sun is below the horizon for

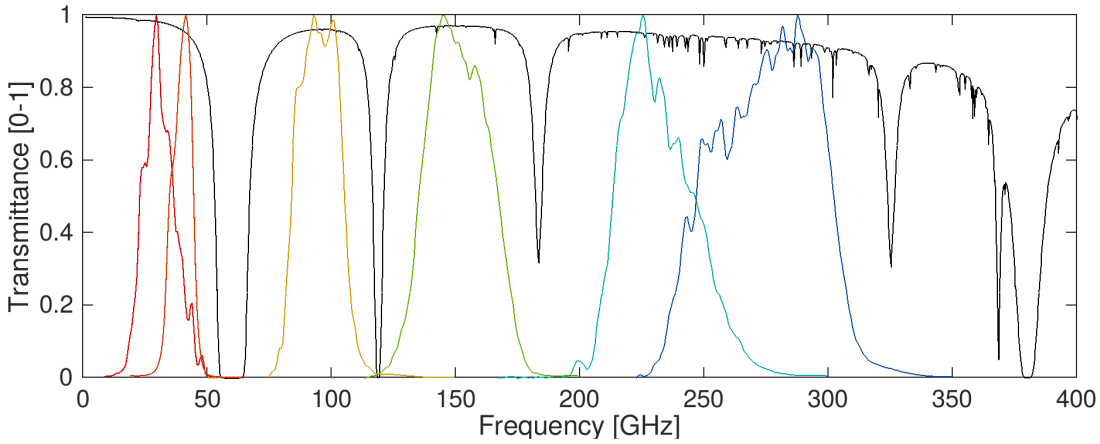


Figure 1.7: Comparison of atmospheric transmission at the South Pole and the measured bandpasses of the BICEP/*Keck* receivers. The median atmospheric transmission (black line) is derived from the *am* model of atmospheric absorption at an elevation of 60 degrees (Paine 2024). The colored lines are real peak normalized coadded receiver spectra, from left to right: BICEP Array 30 GHz, BICEP Array 40 GHz, BICEP3 (95 GHz), *Keck* Array 150 GHz, *Keck* Array 220 GHz, *Keck* Array 270 GHz.

six months. Both the high altitude and the low water vapor content in the air at the South Pole make for ideal millimeter observational conditions (Kovac & Barkats 2007).

There are four important atmospheric absorption features in the millimeter bands; two water lines (at approximately 183 and 325 GHz), a band of oxygen lines at approximately 57 GHz, and a satellite line associated with this band at 118 GHz. These absorption features contribute directly to atmospheric loading, and thus we carefully design our receivers' frequency response (aka bandpass) to avoid the absorption lines (shown in Figure 1.7). We typically refer to our receivers by the (rounded) center frequency of their bandpasses (ie, Keck220 is the Keck receiver with rounded bandcenter at 220 GHz, BA150 is the BICEP Array receiver with rounded bandcenter at 150 GHz,

and so on). The exceptions to this convention are the stand-alone receivers BICEP1 (95, 150, and 220 GHz), BICEP2 (150GHz) and BICEP3 (95GHz).

1.2.1 The Receivers

Each of the receivers made for the BICEP/*Keck* Experiment have been designed to maximize sensitivity at degree angular scales, which results in a set of experimental design principles for each receiver. These are:

- Modular detector design to allow for quick adaptation of newer, higher sensitivity detectors.
- The telescopic front ends are 'small' aperture (about 0.25 to 0.5 meter diameter) on-axis refractors which prioritize reducing instrumental noise and polarization systematics over increasing resolution[13].
- Continuous observation of a low foreground (small galactic dust column density) patch of the sky, which generates deep and high sensitivity maps.

The Detectors

Modern telescopes require a way to convert incident radiative power to a digital signal that can be processed and analyzed. Many CMB experiments use bolometers[14], which

[13]As we are targeting degree angular scales, we do not need high resolution (or larger apertures, recall $\theta \propto \lambda/D$). We instead desire high sensitivity, so the simpler the optical system, the better.

[14]Bolometers have a long history going back to the 1880s as highly sensitive sensors (Langley 1880). There is a recent trend in millimeter detectors towards kinetic inductive devices (KIDs), which operate on

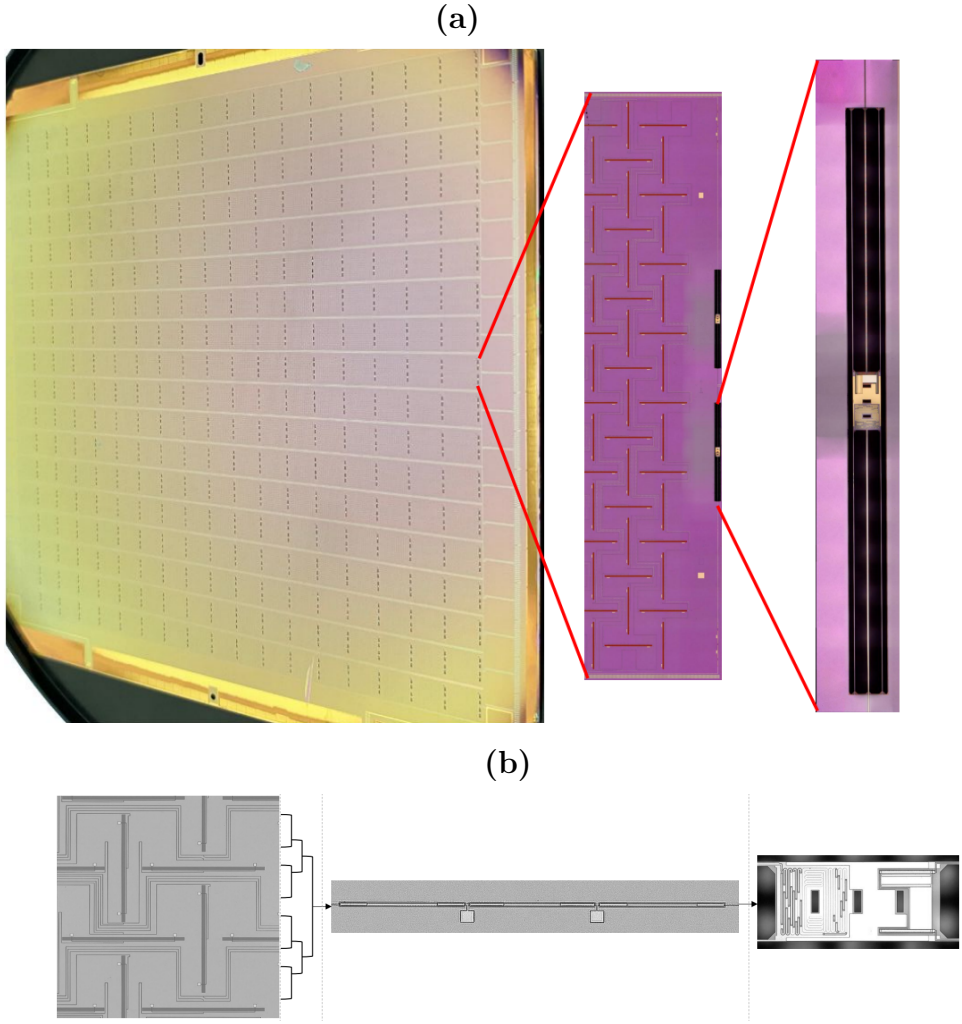


Figure 1.8: (a) A 6" BA 150 GHz detector tile (left), a zoomed in view of the bottom of one pixel (middle), and a further zoomed in view of a single polarization TES island (right). Right photo credit: Alessandro Schillaci; middle and left photo credit: Cheng Zhang; aggregation credit Cheshire (2024). (b) Photon path through a TES (left to right): optical power is collected by two co-located slot antenna arrays, which is coherently summed by a microstrip summing tree, passed through a band-pass filter, and deposited onto the thermally isolated TES island with a gold meander resistor. Figure credit: Zhang (2023).

convert thermal energy into a change in electrical resistance. Recent BICEP/*Keck* instruments have used transition-edge-sensors (TES), which operate by holding a superconductor on the transition edge (the temperature at which the superconductor quickly transitions between normal resistance and no resistance) with an applied bias voltage. The applied bias results in a negative electrothermal feedback loop: when radiative power is coupled to the TES, the resistance increases and the TES current reduces. We measure the amount of current supplied by the feedback loop required to bring the TES back to normal resistance, which corresponds to the amount of radiative load on the TES (BICEP2 Collaboration et al. 2015a).

BICEP3's focal plane consists of twenty 4" square detector tiles in a 5-by-5 grid (corners and an additional spot are left empty); each tile has a grid of 60 photo-lithographically deposited pixels which in turn consist of two orthogonal linearly polarized detectors. The optical power is coupled from free space with a phased array of co-located slot-dipole antennas. Each set of linear slot arrays feed to a separate TES via a microstrip summing network and a band-pass filter (an example is shown in Figure 1.8) (BICEP/Keck Collaboration et al. 2022). Each detector in a pair form a co-located beam on the sky with either vertically or horizontally aligned polarization response, so we take the difference between each detector pair (termed a 'pair-diff') to measure a linear polarization (Q or U) signal (BICEP/Keck Collaboration et al. 2021).

BICEP Array receiver focal planes utilize the same photolithographic detector technology, but with twelve 6" square detector tiles in a 4-by-4 grid (again with corners

a different physical principle (photons break Cooper pairs resulting in a measurable change in inductance) and have the potential to be highly sensitive, but have not been widely adopted among CMB receivers.

left empty). The number and physical scale of the pixels in a tile change depending on the designed frequency—at a lower desired frequency, fewer pixels can fit on a 6" tile. As such the BA30/40 tiles have 16/25 pixels per tile (respectively), while BA150 and BA220 have 324 pixels per tile (Zhang 2023; Hui et al. 2018; BICEP/Keck Collaboration et al. 2024c).

The TES islands are thermally isolated from the rest of the tile with thin silicon-nitride legs, which also serve as bridges for the microstrip and TES bias lines to the island. There are two superconducting metals (titanium and aluminum) in series on the TES island, which have different transition temperatures at 500 mK and 1.2 K, respectively. For regular CMB observations, we use the titanium transition, while for higher temperature calibration measurements (such as the beam characterization or the FTS measurements discussed in Chapter 6) we use the aluminum transition. We can characterize the detectors with different superconductors as both metals share an antenna network, bandpass filter, and TES island (BICEP2 and Keck Array Collaborations et al. 2019).

Each detector is inductively coupled to a superconducting quantum interference device (SQUID), which acts as an amplifier. The detector SQUIDs, called SQ1s, are flux-locked with a feedback coil: as the current in the TES changes in response to power, we compensate by applying current to the feedback coil to cancel the change (Irwin et al. 2012). Another set of SQUID series array (SSA) at 4 K impedance matches between the cold focal plane and the warm readout electronics. The SQUIDs are sensitive to changes in magnetic flux—including external fields, such as the Earth's—so the entire focal plane is surrounded by a pure niobium cylinder. At operating temperatures the niobium is superconducting, preventing external magnetic fields from interacting with the SQUIDs

(BICEP/Keck Collaboration et al. 2022).

We switch between sampling different detectors with 'time domain multiplexing' (TDM), which involves an additional four-Josephson-junction (SQUID-like device) with high critical current acting as a flux switch (FS) in series with each detector's SQ1. When no flux is applied to the FS it is super conducting and the SQ1s are shorted. When a detector is selected by the multi-channel electronics (MCE), the detector's entire row is biased and the row select line running along the column applies a half quantum flux to the selected detector's FS. The rest of the detectors in the row's switches remain superconducting and shorted, resulting in the row bias current flowing through only the selected detector. The MCEs cycle between rows at 12.6 kHz for BICEP Array and 25.3 kHz for BICEP3 (BICEP/Keck Collaboration et al. 2022; Cheshire 2024).

The Cryostat

We make our transmissive optics as cold as possible to reduce the confounding optical power produced by our instrument. To make our optics cold we must both remove and prevent heat from entering the receiver. Therefore, we carefully control the three forms of heat transfer: radiation, conduction and convection.

We do not want to fully eliminate radiative transfer through our instrument: after all we want to measure radiation from the CMB. We do wish to limit radiation generated by non-astrophysical sources (such as the building or the ground), so our cryostats are open ended nested shells. We nest shells of decreasing temperatures to reduce radiative thermal power from reaching the focal plane (as the detectors are surrounded by a 4 K aluminum shell, they only couple to a thermal source at 4 K, not the external 300

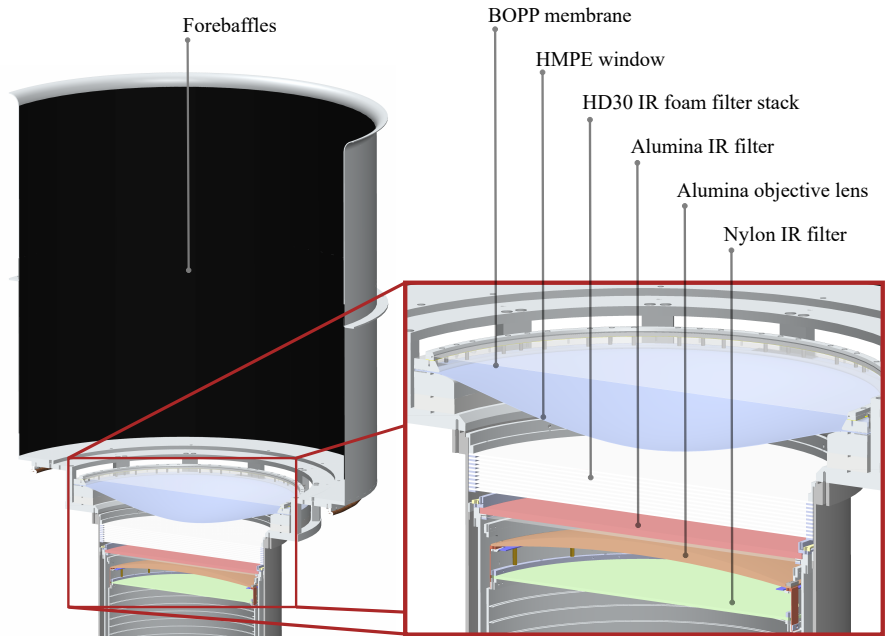


Figure 1.9: Cutaway view of the top of the SolidWorks model of BICEP3, with forebaffle attached. A zoom-in on the optical components at the top of the receiver is included to the right, with transmissive optical components listed. The full cutaway of the BICEP3 receiver (without the new thin HMPE window) is shown in BICEP/Keck Collaboration et al. (2022). The primary optical differences between BICEP3 and the full BICEP Array receiver cutaway shown in Figure 1.10 are the alumina lenses, second nylon filter below the field lens (not shown), and the flat focal plane (not shown) in BICEP3.

K), and to reduce the radiative thermal load on cold stages. Between each shell are multilayer stacks of aluminized mylar (commonly referred to as 'multi-layer insulation' or MLI) to further reduce radiative heat transfer.

Conduction is carefully utilized to move thermal power to our refrigerators by adding high conductivity thermal supports. We require two refrigerators to remove heat from the system: one the 'workhorse' cryocooler (a helium pulse tube) to remove the

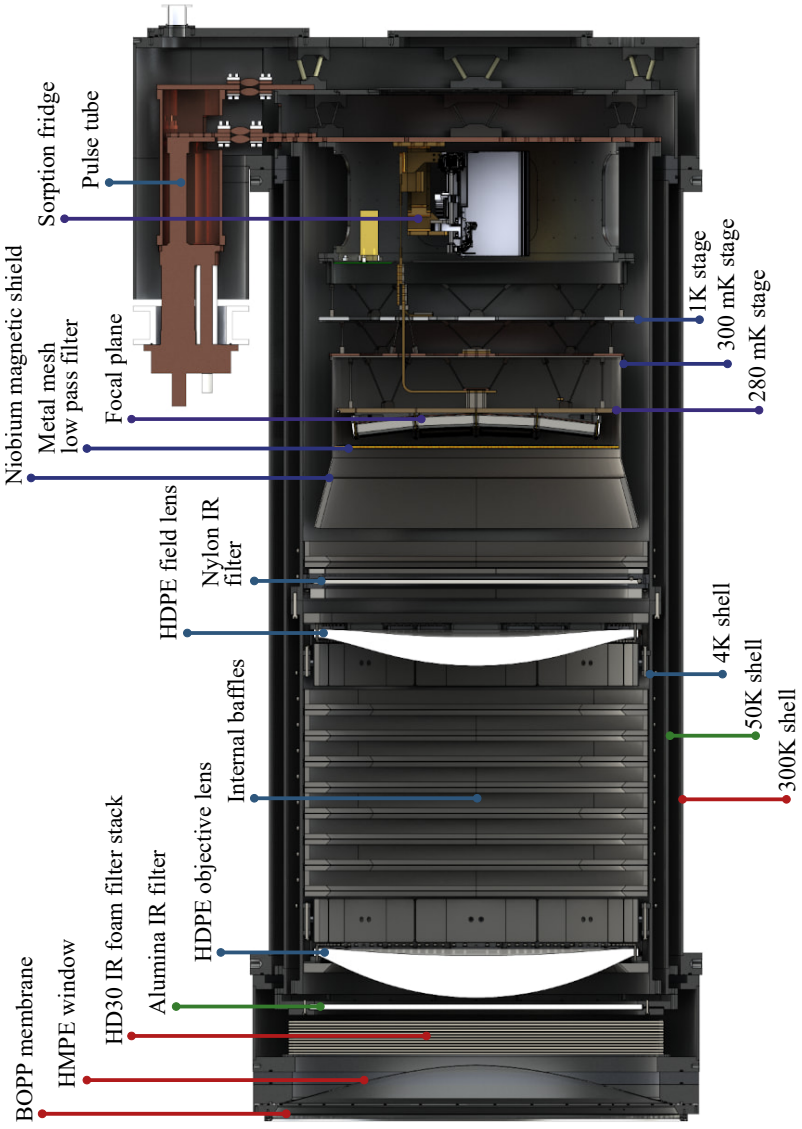


Figure 1.10: Cutaway view of the SolidWorks model of BA150. The temperature of a component is roughly color coded to the temperature at which they sit: components to the left of the HD30 IR foam filter stack are generally at or close to room temperature (300 K), the alumina IR filter is the only transmissive optical component at 50 K, components between the HDPE objective lens and the nylon IR filter are at 4 K, and components to the right of the niobium magnetic shield are mostly at or below 300 mK.

majority of the heat, and the other a finer and colder fridge for cooling the magnetic shield to 300 mK and our detectors to 280 mK (a sorption fridge).

The pulse tubes are continuous cryocoolers provided by Cryomech[15], which operate by pulsing a gas (in our case, helium) through a porous medium with high specific heat. The 'cold head' of the pulse tube is the primary heat sink of the cryostat, and both the 50 K and 4 K shells have high conductivity heat straps connecting to the appropriate stages of the pulse tube, as seen in Figure 1.10.

The sorption fridge is a three-stage set of closed loop chambers filled with either ^{3}He or ^{4}He and active charcoal pumps, provided by CEA/SBT[16]. Flow through the chambers is controlled via a set of resistive switches and heaters. Helium is adsorbed by the charcoal and released by increasing the pressure and boiling point in the charcoal via external heaters. The gas then reaches a cold condensation point, and is condensed into a reservoir through gravity. After a sufficient amount of liquid is collected, the external heat is cut to the charcoal pump, which begins adsorbing and 'pumping' helium again, lowering the pressure in the loop and forcing helium to evaporate to maintain equilibrium vapor pressure, therefore reducing the temperature of the liquid and the surrounding reservoir. The ^{4}He loop cools to 1.4 K, and precools the inter-cooler (IC) ^{3}He stage, which, in turn, cools to 350 mK and precools the final stage—the ultra-cold (UC) ^{3}He stage—which cools to 230 mK. The UC stage is connected to the focal plane

[15]Cryomech, Inc.: `www.cryomech.com`. The *Keck* Array used PT410s, while BICEP3 and BICEP Array use PT415s.

[16]French Alternative Energies and Atomic Energy Commission, Low Temperature Systems Department; `https://www.d-sbt.fr/en`

through a stainless steel block to act as a thermal filter. The fridge must be 'cycled' (charcoal pumps heated to release helium) every few days for a few hours, which imposes a modest cost on the observational efficiency.

For mechanical supports, we utilize strong materials[17] with poor conduction, which are typically composite struts to support the bottom of the shells (either G10, a fiberglass composite, or carbon-fiber), or titanium 'boomerangs' separating the shells at the top.

Convection is the easiest heat transfer method to eliminate, conceptually: we just remove the gas. In practice, pulling a vacuum in our receivers had a domino effect that led to most of this thesis. To observe external signals, we need a transmissive vacuum window on the exterior of the receiver, which has to remain in contact with the ambient environment. These windows need to be both highly transmissive and strong enough to withstand atmospheric pressure. Wide aperture millimeter vacuum windows are commonly made of plastic, which is not very strong, requiring thicker (i.e. less transmissive) plastic slabs. We can make the plastic windows thinner with new technology, which is discussed further in Chapter 2.

The Optics

Our lenses collect and focus incident radiation onto the focal plane by bending incident rays. The lenses are meant to concentrate the light without adding too much excess thermal optical emission, which generally results in lenses made of highly transmissive, high index of refraction materials to make the lenses as thin as possible. This is why

[17]Using strong materials allows us to make supports with smaller cross sections, making them even worse thermal conductors.

the BICEP3 lenses were made of alumina (index of refraction about 3.1); however, alumina is difficult to source, machine and coat with anti-reflective layers (more on this in Chapter 3), particularly at higher frequencies (Dierickx et al. 2023; Jeong et al. 2023). For these reasons, we switched to using high density polyethylene (HDPE) lenses (index of refraction about 1.53) for BICEP Array[18].

An optical redesign conducted between deployment of BA30/40 and BA150 changed the lens and focal plane design. By allowing the focal plane to become 'curved' (really, flat tiles oriented tangent to a curved focal surface) and by removing the requirement that the lenses be convex-plano (top lens surface convex, bottom lens surface planar[19]), the lenses could be made thinner and the entire focal plane could be more evenly illuminated. The previous optical design (used for BA30/40) is shown in Moncelsi et al. (2020), while the current design (used for BA150 and BA220/270) is shown in Figure 1.10.

The majority of the transmissive optics in our cryostats are IR filters: IR would be a significant portion of the radiative thermal load into our receivers if we allowed it to pass through to the sub-kelvin stages—potentially preventing the detectors from fully cooling—and our antennas would readily couple with IR radiation—potentially significantly increasing optical noise). We need to limit the amount of IR radiation

[18]This was not the first set of BICEP/*Keck* receivers to use HDPE lenses, as both BICEP2 and the *Keck Array* had utilized HDPE (BICEP2 and Keck Array Collaborations et al. 2015).

[19]This requirement was intended to make the anti-reflection (AR) coat application possible, as layers could be easy placed below the lens on a flat bottom surface. We apply AR coats to the concave bottom surface of the new lenses by placing the layers on a blank negative made out of bulk PTFE; the blank acts as reference plane like the flat surface did previously during the lamination process.



Figure 1.11: Exterior of the BICEP Array mount in the austral summer of 2025. Around the edge of the mount is the ground shield made of aluminum plated plywood. Humans in foreground (for scale) left to right: Steve Jungst, Markus Ayasse, Mike Echter, Yuka Nakato, Dave Pernic, Sofia Fatigoni, Matthew Petroff, Brianna Cantrall, Nolan Maher, Baibhav Singari, and Cheng Zhang. Above the people are three BA forebaffles (from top left counter-clockwise BA150, BA30/40, and the recently installed BA220/270) and one remaining Keck forebaffle (Keck270). Photo credit: Matthew Petroff.

reaching the detectors as much as possible, which we do primarily with IR absorbing dielectric filters (the HD30 foam filter stack, alumina filter, and nylon filters) and with a metal mesh low pass filter just above the focal plane[20] (Halpern et al. 1986; Ade et al. 2006; Choi et al. 2013; BICEP2 and Keck Array Collaborations et al. 2015).

[20]'Low pass' means that it allows lower frequencies to pass through, reflecting higher frequencies like IR.

The main beam of our receiver should exit the forebaffle completely unobstructed, excluding a small amount of truncation at the 4 K aperture stop (typically situated above the objective lens). Scattered power out to high angles should be terminated at the warm absorptive forebaffle. See Figure 1.9 for a model of the top of the BICEP3 cryostat, including forebaffle. Further side-lobe response is terminated at the reflective ground shield, which redirects the side-lobes to cold sky instead of coupling to the (relatively) hot ground. Both the BA forebaffles and ground shield are shown in Figure 1.11, with humans for scale.

1.2.2 The Observations

The geographic South Pole also provides the same sky year round: this allows us to continuously observe the same field. Compact, on-axis refracting optics are rotationally symmetric, allowing us to rotate the receiver along the 'boresight' (center axis of the mount) to sample the linear polarization of the sky at different angles. This breaks any degeneracy between an individual pair-diff polarization angle and the sky signal, as multiple detectors sample the same portion of the sky at different angles. We observe at seven different 'deck' angles for BICEP Array (23°, 68°, 113°, 158°, 248°, 293°, 338°), but only four deck angles in BICEP3 (23°, 68°, 203°, 248°) due to interference between the cryostat and the mount.

For most observations (approximately 90% of on-sky time), we track the same patch of the CMB (known as the "Southern Hole", a patch far from the galactic plane with low foregrounds) as it moves in sidereal time, as shown in Figure 1.12. The telescope raster-scans in azimuth over the patch for approximately 50 minutes (a 'scanset' or 'tag')

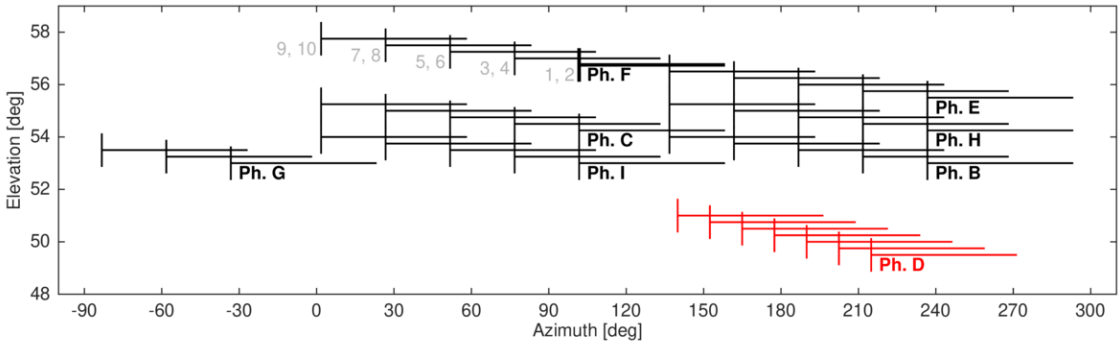


Figure 1.12: An example of the scan strategy of BICEP3 on the antarctic sky. A schedule starts with the Phase B (Phase A is the fridge cycle), and runs through the other phases alphabetically. Phase D (highlighted in red) scans over the galaxy, not the CMB patch. The horizontal lines denote the azimuthal extent of an individual scanset; the vertical lines denote the elevation extent of the bracketing elevation nods (elnods) that start and end a scanset. Ten scansets make up most phases (except for Phase G, which only has six, and Phase D, which only has seven), and are labeled in gray for Phase F. BICEP Array uses the same strategy, but only covers Phases B through F. Figure credit: BICEP/Keck Collaboration et al. (2022).

to produce time-ordered data (timestreams), with bracketing elevation nods to sample atmospheric stability[21]. Tags are generally cut based on weather-related elnod statistics. After each pair of scansets, the telescope moves to track the CMB patch, and repeats. A set of scansets is referred to as a 'Phase': BICEP3 has nine phases (A–G, though

[21]Just before and after the elnods we typically take partial load curves (PLCs), which vary the bias current in the TES to measure the margin to the saturation power of the detectors (discussed further in §4.2). These PLCs used to be used for relative calibration of the scanset. Due to excess low ℓ noise in BA150, PLCs were fully turned off in BICEP Array part of the way through 2024, and only reinstated per phase (not per scanset) in 2025.

Phase A is reserved to cycle the sorption fridge) lasting approximately three days, while BICEP Array only has six (A–F, with A again reserved for the fridge cycle) lasting approximately two days due to lower fridge hold time[22]. Phase D is the only regular phase that does not observe the CMB patch; instead it covers a portion of the galactic plane.

A CMB schedule refers to the full set of phases for a given mount; other schedules could be designated for other types of observations (such as calibration observations discussed elsewhere in this thesis). The full observing season typically lasts from early March to late October; the weather is very stable during the winter months, though clouds and blowing snow are significant contributors to observing efficiency reduction.

After a season's worth of tags are taken and appropriately trimmed for weather or other systematics, we bin the timestreams into pixels on the sky by utilizing the knowledge of where the detectors were pointed at a given moment in time, accounting for the relative angle of a detector pair based on the boresight rotation. This binning procedure produces Q and U maps of the CMB patch.

The maps (which are generated per frequency band) are then 'matrix purified' to separate out B-mode structures (BICEP2 Collaboration et al. 2014, 2016). We then apodize the maps, take a 2D Fourier transform, and then average within annular rings in the resultant Fourier space to bin the power spectrum into what we term 'bandpowers'. Additionally, the Q/U maps are rotated into E/B (BICEP/Keck Collaboration et al. 2021).

[22]The hold time is the period where liquid helium is condensed in the stages within the sorption fridge. Higher thermal load will increase the evaporation rate of the helium, reducing hold time.

1.2.3 The Model

We compare an eight parameter lensed-ΛCDM+dust+synchrotron+noise model to the bandpowers by exploring the parameter space with the Markov chain Monte Carlo (MCMC) method, and maximizing the joint likelihood (which is approximated with the Hamimeche-Lewis function) of all BB auto- and cross-spectra[23] (Hamimeche & Lewis 2008; BICEP/Keck Collaboration et al. 2021). This model incorporates foregrounds (galactic dust/synchrotron) and gravitationally lensed ΛCDM. There are other potential sources of power along the line-of-sight (such as the Sunyaev-Zel'dovich or integrated Sachs-Wolfe effects); galactic foregrounds and extra-galactic gravitational lensing are the primary non-primordial signals at the relevant angular scales. Let us discuss each component of that model.

The Galaxy

Galactic synchrotron radiation is sourced from relativistic electrons being accelerated in the galactic magnetic field. Synchrotron emission is polarized, as it should be radiated perpendicularly to the magnetic field lines. We expect that synchrotron emission follows a simple power law at our frequencies, $A_s \propto \nu^{\beta_s}$, with $\beta_s \approx -3$ (note that the frequency spectral index is negative, indicating that the intensity decreases at higher frequencies). External low frequency ($\nu < 40$ GHz) maps from WMAP and *Planck* have so far

[23]Recall that we typically cross maps with *Planck* and WMAP; there are a set of 66 BB auto- and cross- spectra in the latest data release, for the eleven frequency maps. The BICEP/*Keck* maps are Keck95, BICEP3 (95 GHz), BICEP2/Keck150, Keck220; the WMAP maps are at 23 and 33 GHz; and the *Planck* are at 30, 44, 143, 217, and 353 GHz (BICEP/Keck Collaboration et al. 2021)

provided the primary constraints on the intensity of polarized synchrotron emission, though the BA30/40 receiver deployed in 2019 is likely approaching those maps' level of sensitivity. We evaluate A_s at an ℓ of 80 and at 23 GHz, the lowest WMAP frequency band.

Galactic dust in the interstellar medium thermally emits a (modified) blackbody spectrum $A_d \propto \nu^{\beta_d} B(\nu, T_d)$; the blackbody is modified because the size of the dust grain affects the efficiency of the emission. Ellipsoidal dust grains will also align with galactic magnetic fields, and preferentially emit along their short axis, producing polarized emission. We assume that the temperature of the emitted blackbody is $T_d \approx 19.6$ K. Recent BICEP/*Keck* analyses have removed the prior on β_d, and we find the frequency spectral index of dust in the CMB patch is around 1.6 (BICEP/Keck Collaboration et al. 2021). We evaluate A_d at an ℓ of 80 and at 353 GHz.

We also fit three additional terms to both the galactic synchrotron and dust emission to model the spatial distribution across the patch: α_s, α_d and ϵ. The spatial power laws go like $\mathcal{D}_\ell \propto \ell^{\alpha_s}$ and $\mathcal{D}_\ell \propto \ell^{\alpha_d}$ for both synchrotron and dust, and ϵ models the spatial correlation between the dust and synchrotron.

A spectral decomposition of the measured 95 GHz and 150 GHz BB spectra are shown in Figure 1.13. Note the similar amplitudes of the BB bandpowers from the CMB and galactic dust and how the each scales with multipole (ℓ) and observing frequency; this emphasizes the strong need to observe to high sensitivity at multiple frequencies.

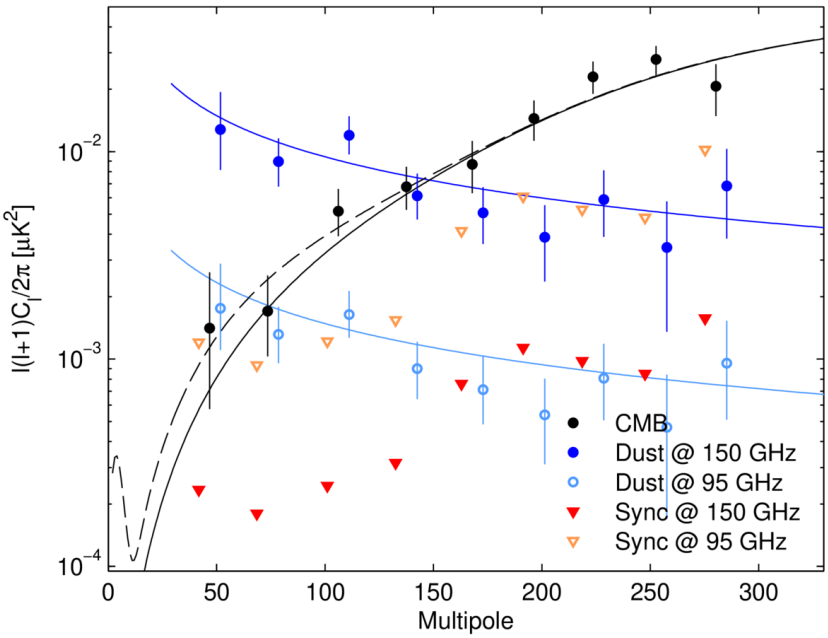


Figure 1.13: 95 GHz (hollow points) and 150 GHz (filled points) BB auto-spectra decomposed into lensed-ΛCDM (black), dust (blue), and upper limit synchrotron (red). The synchrotron powers are offset slightly horizontally for clarity, and have not yet been detected in the BICEP/*Keck* CMB patch (the triangles show the 95% confidence upper limits). The lines denote the model, where the solid black line is nominal ΛCDM, black dashed line lensed-ΛCDM with added $r_{0.05} = 0.011$, and the blue lines representing the maximum likelihood dust model from the baseline BICEP/*Keck* analysis ($A_{d,353} = 4.4\mu K^2$, $\beta_d = 1.5$, and $\alpha_d = -0.66$) at 95GHz (light blue) and 150GHz (dark blue), respectively. Figure source: BICEP/Keck Collaboration et al. (2021).

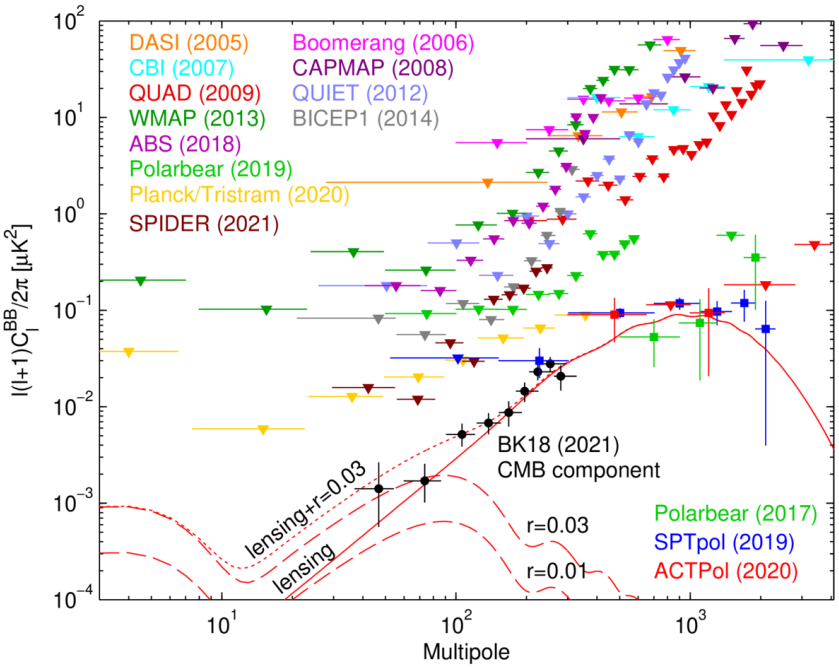


Figure 1.14: Recent suite of published BB power spectrum measurements. Triangles are upper limits, squares are measured values from external collaborations, circles are measured values from BICEP/Keck Collaboration et al. (2021). The predicted spectrum for the lensed BB spectrum is shown with a solid red line, along with the theoretical spectrum of two values of r with dashed red lines, and the summed lensed and r spectrum with dotted red lines. Figure source: BICEP/Keck Collaboration et al. (2021).

The Lensing

The gravitational warping of spacetime from galaxy clusters along the line-of-sight slightly distort the position the CMB anisotropies we observe. The underlying, unlensed fields of $\bar{T}$, $\bar{Q}$ and $\bar{U}$ are related to the line-of-sign gravitational potential ϕ (recall the $\phi\phi$ power spectrum in Figure 1.4) like

$$T(\hat{r}) = \bar{T}(\hat{r} + \nabla_{\perp}\phi) \tag{1.23}$$

$$(Q \pm iU)(\hat{r}) = (\bar{Q} \pm i\bar{U})(\hat{r} + \nabla_{\perp}\phi). \tag{1.24}$$

The lensing does not alter the intensity nor polarization direction, but rather shears (transversely translates) the apparent position. This produces a small but detectable 'softening' of the acoustic oscillations in the T and E anisotropies. Unfortunately, the lensing breaks the symmetry of the E-mode anisotropies, transforming a small fraction into a B-mode pattern: this is source of the lensing B-mode foreground. The primordial tensor B-modes are sub-dominant to the lensing foreground for $r \lesssim 0.01$, which we are quickly approaching in BB measurements, as shown in Figure 1.14. Fortunately, we can estimate ϕ from a combination of other high resolution maps and ΛCDM E-modes, which allows us to 'delens' the B-mode measurements and recover the primordial signal.

1.3 The Thesis

This thesis will describe the three projects I have developed for the BICEP/*Keck* experiments, summarized as follows:

1. Thin vacuum windows development and testing of thin laminate polyethylene

vacuum windows for use on wide aperture BICEP Array cryostats. Thin windows are the primary focus of Chapters 2 and 4 (where their respective mechanical and optical attributes are described), and play a significant role in the instrument noise modeling in Chapter 5. Additionally, procedures for thin window generation are discussed in Appendixes A and B.

2. Wide aperture and wide band anti-reflection coats for polymer optics in BICEP Array. Everything to do with all AR coats (including both idealized cases and designs for both thin windows and other plastic optics) are the primary focus of Chapter 3. Appendixes C and D are related to AR coats as well.

3. Receiver bandpass characterization (i.e., measuring the frequency response of the detectors) in BICEP Array is discussed in Chapter 6, including bias estimates on r from bandcenter uncertainty for the BK24 dataset.

Chapter 2

Thin Window Mechanics

Portions of this thesis chapter originally appeared in the literature as

BICEP/*Keck* Collaboration et al., *arXiv e-prints, arXiv:2410.12089*

There are broadly two ways to increase instrument sensitivity: increase detector counts while keeping the per-detector sensitivity the same or decrease the optical loading thus increasing the per-detector sensitivity. These two goals lead to conflicting physical designs. Increasing the detector count requires a larger aperture size, which leads to thicker optical components and (because the optical emission from a component increases with its thickness) therefore larger apertures tend to degrade the per-detector sensitivity.

To accommodate the order of magnitude increase in detector count between the recent stages of BICEP receivers the aperture size of the receivers has doubled from 264 mm to between 520–560 mm (BICEP2 and Keck Array Collaborations et al. 2015; Moncelsi et al. 2020). The readout electronics and transition edge detectors are carefully designed so that the receivers are photon-noise limited (more on this in Chapter 5).

Therefore reducing in-band optical load can have a profound impact on a given BICEP receiver's per-detector sensitivity.

The in-band radiative load emitted by the optical elements is the product of their absorptive loss and physical temperature. The absorption loss of a material is directly related to the complex permittivity ($\varepsilon = \varepsilon' - i\varepsilon''$) of a material. Absorption is discussed numerically throughout this thesis as a material's loss tangent, which is the ratio of the imaginary part over the real part of the permittivity, denoted as $\tan\delta$.

Recall that the cryostat vacuum windows must have low transmission loss in band and be strong enough to withstand the force of atmospheric pressure. Previously, millimeter-wave vacuum windows have been made of ceramics [such as fused-silica (Datta et al. 2021)], plastic foams [including various formulations of Zotefoam (BICEP2 and Keck Array Collaborations et al. 2015)], or bulk plastic materials [primarily high density polyethylene (HDPE) or ultra high molecular weight polyethylene (UHMWPE) (BICEP/Keck Collaboration et al. 2022; D'Alessandro et al. 2018)]. A more detailed history of the development of millimeter vacuum windows can be found in Barkats et al. (2018).

The larger aperture sizes of modern CMB receivers, however, limit the materials available for use. BICEP3 and BICEP Array have nominal clear apertures at the window of 730 mm, while the full outer diameter of the window is 900 mm. Current plans for the CMB-Stage 4 Small Aperture Telescopes (SATs) and Large Aperture Telescope (LATs) include similarly sized windows (Abazajian et al. 2019). Ceramics become prohibitively expensive at those scales. Plastic foams would become impracticably thick and lossy: the primary advantage of the foam windows are their low index of refraction and low

transmission loss, but they have the disadvantage of being very weak. A suitably thick foam window would include many laminated layers, which would make their loss properties significantly worse. For example, the 150 GHz *Keck* receiver windows were 120 mm thick and estimated to cause $\lesssim$2% transmission loss (BICEP2 and Keck Array Collaborations et al. 2015); the doubled aperture size for new BICEP receivers would require a foam window four times thicker[1]. A window with such high transmission loss ($\lesssim$8%) would add approximately 3 pW optical load to a 150 GHz detector, swamping out all other sources of optical power at that frequency.

If we replace plastic foams with bulk plastics (which are strong compared to plastic foams, but not very strong compared to other materials), the windows must still become thicker to maintain a suitable safety factor for a vacuum vessel (discussed further in §2.2), but do not absorb nearly as much power as a thick foam window. For BICEP3 (the current generation 95 GHz receiver), we previously estimated that over half of the instrument loading originated from the 31.8 mm thick HDPE slab window, though recent measurements have provided evidence to contradict this estimate (discussed further in §4.2 and Chapter 5) (BICEP/Keck Collaboration et al. 2022). That window has since been replaced by the thin window design discussed in this thesis.

In Figure 2.1 I model the optical loading on a large aperture BICEP-style receiver for six observed bands and different window thicknesses. I also show the relative change in per-detector white noise and survey time for different window thicknesses[2]. For the

[1]Imagine with me, for a moment, a window nearly *half a meter* thick (480 mm). It would be ungainly. It would probably have to get even wider and thicker so that the support structure does not clip the beam. It is not practical.

[2]There remains some uncertainty about the exact magnitude of the loss properties of the windows, as

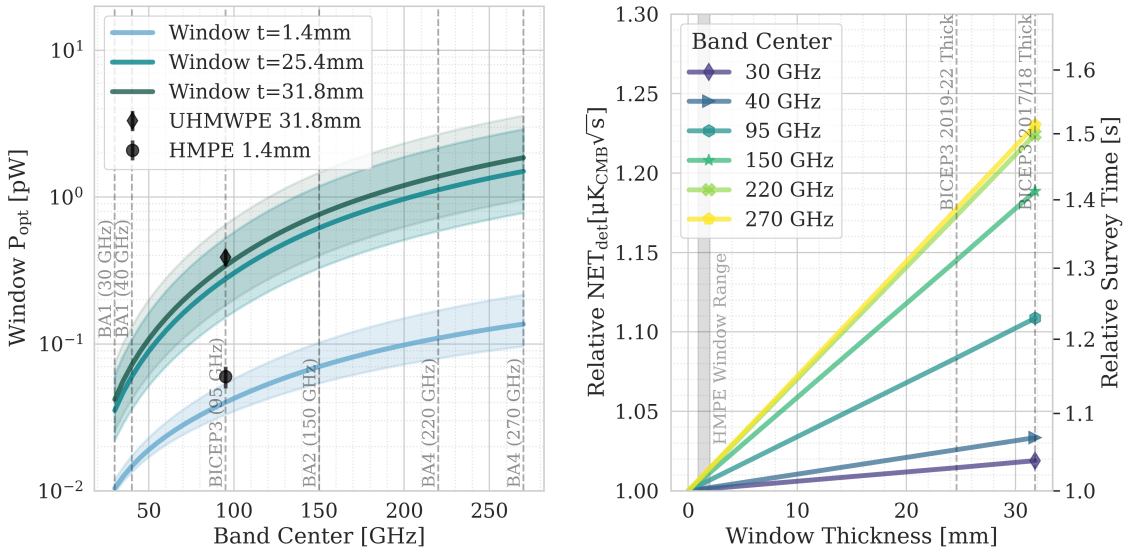


Figure 2.1: [Left] Modeled optical loading from a window with thicknesses of 1.4 mm, 25.4 mm, and 31.8 mm, at band centers from 30 to 270 GHz (fractional bandwidth of 0.25) through a BICEP-style instrument. Black points are measured optical loading on BICEP3 from a 31.8 mm (1.25") UHMWPE spare window and the 1.4 mm laminate window deployed on BICEP3 for the 2023 season. Shaded bands are uncertainty on the loss properties of the window, with $\tan\delta$ between 0.6×10^{-4} to 2.4×10^{-4}. [Right] Modeled relative noise equivalent temperature per detector (NET_{det}) and relative survey time increase compared to an instrument without a window, for each of the BICEP Array observing band center frequencies at an assumed $\tan\delta$ of 1.2×10^{-4}.

lowest frequency band, at 30–40 GHz, decreasing the window thickness only yields marginal improvements in white noise and mapping speed. For bands above 40 GHz, however, a substantial reduction in window thickness from the original could potentially decrease noise by tens of percent. At the high frequencies (such as the 270 GHz band), such a window thickness reduction could potentially decrease the survey time by as much as 50%. The window is the only element in the optical chain where improvement would produce such a significant return, primarily due to the window's relatively high temperature compared to the rest of the cold optics.

We have achieved an order of magnitude reduction in window thickness using a special form of polyethylene we call high modulus polyethylene (HMPE, commercial name Dyneema or Spectra) laminated with low density polyethylene (LDPE) to generated a polyethylene composite material. The high strength of the HMPE allows for thicknesses of around a millimeter at our apertures, in the same range as the wavelength of light we are observing, which will become relevant for designing the anti-reflection coats of the window (discussed further in Chapter 3). The LDPE melts at lower temperature than the HMPE, allowing for the laminating material to fully surround the HMPE fibers without compromising the strength carrier's integrity, which is discussed further in §2.2 (Vasile 2005; Eiben et al. 2022). Because the composite is entirely polyethylene, we expect the window to appear homogeneous to millimeter light, and therefore have similar optical properties—such as index of refraction, loss and scattering—as standard bulk polyethylene (discussed further in Chapter 4).

---

shown on the left hand side of Figure 2.1; for clarity, only one loss value is assumed for the window on the right hand side. A discussion and breakdown of the NET modeling to fit to the optical powers shown in Figure 2.1 and related measured NETs is the primary focus of Chapter 5.

2.1 Thin Windows

There are three parameters that influence the in-band optical emission of a material: the physical temperature, the absorption coefficient, and its thickness. The temperature is not a parameter than can be changed as the window is in thermal equilibrium between the ambient temperature outside and the cold optics inside the cryostat[3]. Similarly, the absorption coefficient cannot be changed; it is a property of the material itself (see Barkats et al. (2018) for further discussion). Therefore the thickness is the only parameter available to reduce the optical loading from the window. Unfortunately, bulk plastic windows tend to be very thick due to bulk polyethylene's low strength. Therefore, to decrease the thickness of the window, we must increase the strength of the materials we use.

HMPE fibers have an ultimate tensile strength over 150 times stronger than bulk HDPE or bulk UHMWPE (Barkats et al. 2018; Eiben et al. 2022). However, the HMPE fibers are woven into a fabric, which is not naturally vacuum tight. We must add an additional processing step, laminating the HMPE with another material that should be optically matched to the strength carrier. As LDPE is also a polyethylene, the millimeter optical properties are expected to be very similar. We add LDPE layers between HMPE layers and then the LDPE, having a lower melting point than HMPE, can be melted to laminate the fabric layers and fill the voids between fibers (Vasile 2005). The ideal

[3]We attempt to warm the window apparatus to prevent snow accumulation by forcing room temperature nitrogen into the volume between the BOPP membrane and the window. This works well enough to prevent snow accumulation and to keep the window above the external ambient temperatures, but the center of the window is still well below room temperature (see §2.6).

lamination temperatures are discussed in §2.2.1. This makes the composite of HMPE and LDPE vacuum tight (Barkats et al. 2018; Eiben et al. 2022). The ideal lamination pressures are discussed in §2.2.2. The detailed window generation procedure is covered in Appendix A; in general, the large thin windows discussed throughout the rest of this thesis have three layers of HMPE oriented at three different angles (0/30/60 degrees) to supplement the stretchy bias (45 deg from weave orientation) direction, with layers of LDPE sandwiching the HMPE to fully surround the fibers during lamination. The final thickness of a window is determined by the number and initial thickness of the LDPE layers.

As shown in Figure 1.9, the window is the first optical element facing the sky. Above the window are "forebaffles" which terminate sidelobe and scattered pickup outside the nominal 27.4° field of view of the telescope at a stable, warm temperature (BICEP/Keck Collaboration et al. 2022). Below the window are a stack of physically separated thin polyethylene foam sheets (Zotefoam HD30) which act as an IR filter (Choi et al. 2013). These thin laminate windows must fulfill a variety of criteria to avoid negatively impacting our sensitive instruments. Mechanically, the window should not allow excessive amounts of gas into the cryostat (§2.5), must withstand the force of atmospheric pressure with a reasonable safety factor (§2.2 and 2.3), and must not deflect too far inwards to avoid contacting the IR foam filter below them (§2.4). The optical design of thin windows are discussed in Chapter 3 and the optical characterization of thin windows are discussed in Chapter 4.

2.2 Strength

In order to dimension these new thin windows correctly, we first start with the measured mechanical properties: the elastic modulus, the ultimate tensile strength, and the strain at break. We measure these mechanical properties using a tensile strength machine (shown on the left in Figure 2.2), that collects force versus elongation data over the range of extension of the material up to failure. Those data are then converted into engineering stress (force over the initial cross sectional area) and strain (percent elongation). Example tensile strength measurements for a 10 atm 135°C HMPE laminate are shown on the right in Figure 2.2.

The ultimate tensile strength is the maximum stress a material can withstand before failure, which may vary slightly between samples due to measurement uncertainty, particularly uncertainty in the cross sectional area. The failure strain is the percent elongation a material can withstand before failure. The slope of the stress/strain curve is the elastic modulus, which is particularly useful to predict how a material will deform (strain) in response to a load (stress).

We were initially concerned that the extra thermal processing associated with lamination would reduce the strength properties of the HMPE. Therefore, we ran a series of temperature (holding lamination pressure steady) and pressure (holding lamination temperature steady) sequences, and then attempted to estimate the laminated HMPE fiber strength using tensile testing. These results are discussed in the next two sections. The raw HMPE ultimate tensile strength is around 3-4 GPa, but most of the measured cross section in a laminate is the LDPE (Dyneema 2008; Barkats et al. 2018). To assess that processing does not reduce the fiber strength, we need to estimate the cross section

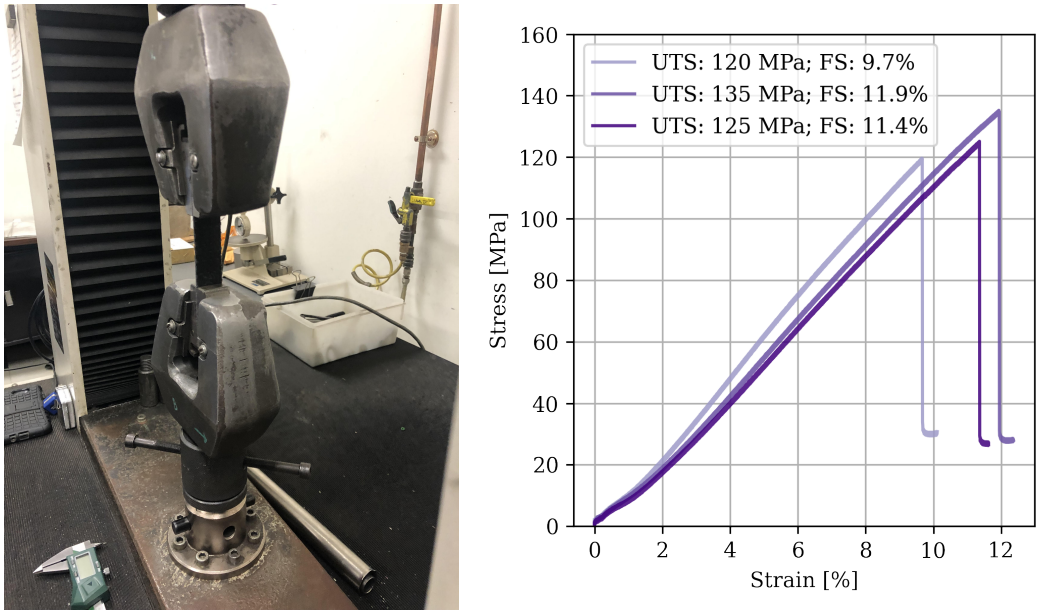


Figure 2.2: [Left] Photo of a window material sample (black strip) in tensile strength machine. The machine measures the force and change in distance between the jaws (big metal claws in the center). [Right] Tensile strength tests of three samples from a high pressure (10 atm, 135°C) HMPE laminate. The stress is calculated by dividing the measured force by the initial cross sectional area of the sample, while the strain is the percent change in elongation. The ultimate tensile strength (UTS) and the failure strain (FS) for each sample are reported in the legend. Laminate HMPE temperature sequence and pressure sequence UTS/estimated fill percentages summarized in Table 2.1 and 2.2, respectively, and compared to other polyethylenes in Table 2.3.

of the HMPE fibers under tension, as they are what is being stressed in the test. If we assume that the fibers are cylinders with radius 0.11 mm, and that there are 32 fibers per inch of laminate, we can estimate the fiber cross-section from the initial width measurement of the tensile sample[4].

[4]The somewhat speculative nature of estimating the cross-section of the HMPE fibers make these

We also estimated the lamination quality by calculating the 'fill factor' from the density of the samples. By measuring the initial volume of the each of the layers (the LDPE layers and the HMPE fibers[5]), then measuring the final laminated volume and mass we can approximate the fill factor within the HMPE fibers by calculating the ratio of the initial and the final densities.

The full cross sectional area measurements are summarized in Table 2.3 for the HMPE laminate material, along with literature values for the bulk HDPE and UHMWPE found in Barkats et al. (2018); D'Alessandro et al. (2018).

2.2.1 Lamination Temperature

The appropriate temperature during lamination must melt the laminating LDPE while the strength carrying HMPE is unaffected. Thankfully, the melting point of LDPE is typically cited to be between 110 to 120°C (Vasile 2005), while the manufacturer claims

estimates particularly susceptible to measurement error (this is one reason they were not reported in BICEP/Keck Collaboration et al. (2024a)). As such, these cross-sections really only confirm that the processing did not significantly reduce the strength, which is useful; however, for the safety factor estimates in §2.2.3 I use the full cross section.

[5]The volume measurement for the HMPE was particularly difficult. We used water displacement to estimate the HMPE volume, but the fabric is somewhat hydrophobic and air bubbles would get trapped between the fibers, biasing the measured volume up. In the end, we put a small HMPE square in a capped off graduated syringe with a pre-measured volume of water, and compressed the syringe by hand to minimize the air bubbles as much as possible. I suspect that the resultant volume estimate for HMPE fabric was still slightly too high, which is why none of the fill factors based on density reach 100% as the initial estimate of the volume to-be-filled is larger than it actually is, so the LDPE cannot fully 'fill' it.

Table 2.1: Temperature sequence ultimate tensile strengths using the estimated HMPE fiber cross section and estimated fill percentages from density measurements.

Baking Temp [°C]	HMPE Fiber Ultimate Tensile Strength [GPa]	Estimated Fill Percentages [%]	Notes
120	3.07–3.10	87	Visibly poor lamination
130	3.24–3.36	94	Fine lamination
135	3.47–3.87	96	Good lamination
140	N/A	96	Good lamination
145	N/A	94	HMPE fibers visibly shrunk after bake
150	2.04–2.68	94	HMPE fibers visibly shrunk after bake

the melting point of HMPE is between 140 to 150°C (Dyneema 2008). However, there was concern about non-obvious impacts of the temperature cycle on the lamination quality and the HMPE. We ran a sequence of samples varying the temperature of the temperature cycle from 120 to 155°C, keeping the laminating pressure the same, and then tested samples' strength with tensile tests.

In Table 2.1, the ultimate tensile strength of a subset of the temperature sequence HMPE laminate samples are reported, as well as the estimated fill factor from density measurements. I also report a qualitative description of the lamination quality[6]. The

[6]In higher quality laminates the HMPE fibers become less visible to the naked eye, seemingly due to less scattering off inadequately infilled HMPE (see Figure 2.4 for an example of lamination quality

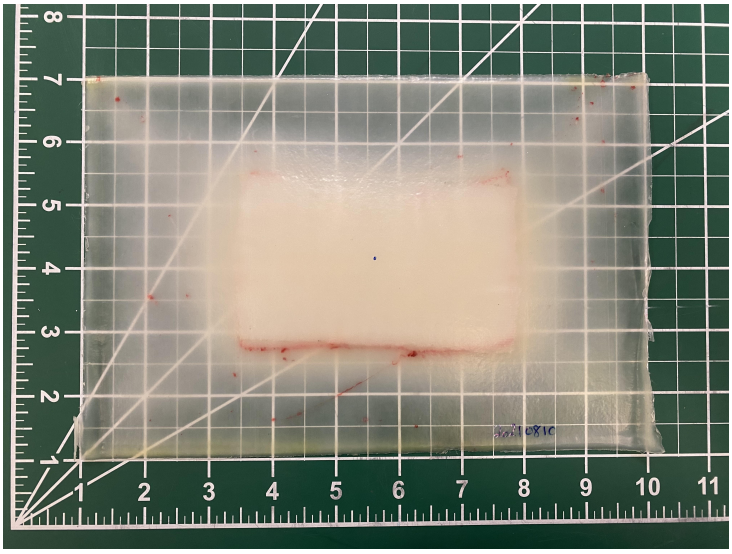


Figure 2.3: Photo of the highest temperature laminate sample generated for the temperature sequence reported in Table 2.1. It was baked at 155°C, above the reported melting point of HMPE (between 140–150°C (Dyneema 2008)). The original dimensions of all layers were 6"x9" (all sequence samples were meant to be cut into a 6"x6" optical sample and three 6"x1" tensile test samples); the central white block is the shrunk HMPE fibers, which appears to have reverted to bulk UHMWPE; the red on the edge of the block is from sharpie marks on the HMPE from cutting prior to the lamination; the see-through portions are the remnants of the laminating LDPE layers that did not get sucked into the central block.

lowest temperature laminate measured was baked at 120°C, and was particularly difficult to strength test; the lamination quality was so poor that the HMPE essentially ripped away from the LDPE when tensile load was applied. The highest temperature laminates (above 140°C) had the HMPE fibers visibly shrink during lamination. I suspect the HMPE fibers have a certain amount of stress ingrained due to the long polyethylene

differences in the pressure sequence).

polymer chains being forced straight, like a spring stretched out. When the polymers gain enough energy to move at or above the melting point, they coil up, releasing the stress of being stretched straight, and making the fibers appear to shrink. This is particularly apparent in the one temperature sequence laminate missing from Table 2.1; the 155°C laminate, shown in Figure 2.3, which shrunk so severely during lamination that it was deemed untestable.

These tensile tests revealed that ideal temperature range is between 130 and 140°C, so we typically aim for 135°C to allow for small temperature variations during processing.

2.2.2 Lamination Pressure

The initial set of prototype laminate windows—generated prior to my arrival—were baked by pulling a vacuum underneath a vacuum barrier with the laminate stack; essentially, they were generated with 1 atmospheres (atm) of pressure applied. These laminates had two primary issues:

1. The index of refraction of the low pressure laminates was estimated to be significantly lower than the bulk index of HDPE—at $n \approx 1.35$ instead of $n \approx 1.53$—which indicated a low fill factor, and made them difficult to anti-reflection coat[7] (Barkats et al. 2018).

2. The low pressure laminates were incredibly leaky, which also indicated low fill

[7]Everything to do with AR coats is discussed extensively in Chapter 3, but to summarize: the ideal single layer AR coat index of refraction is $n_{AR} = \sqrt{n_{core}}$, so if the laminate index is $n_{core} \approx 1.35$, the ideal AR coat would be 1.16, which is difficult (though not impossible) to source.

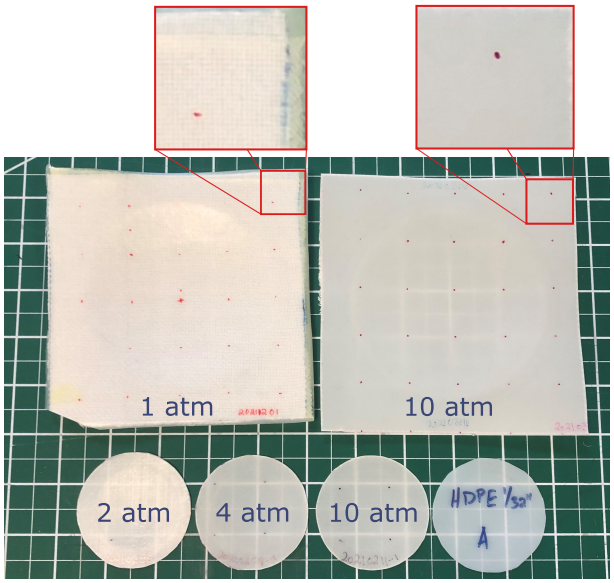


Figure 2.4: Photo of several small scale laminated window samples, labeled with the pressure at which they were laminated in atmospheres (atm). The top two samples were used in the permeation tests in Figure 2.8; the recipes are identical apart from the difference in lamination pressure. Note that the weave of the HMPE fibers remains visible in the 1 atm laminate, whereas the 10 atm laminate appears more homogeneous. Below these is a row of three smaller HMPE laminate samples and a solid extruded HDPE sheet having the same thickness. The highest pressure (10 atm) laminate (top right and bottom second to the right) is visually most similar to the solid HDPE sample. Note that the grid pattern on the mat below the samples laminated at lower pressure is blurred by light scattering associated with the incomplete filling of the HMPE weave with LDPE.

factor[8] (discussed further in §2.5).

In Barkats et al. (2018), the low fitted value of the index of refraction for laminate windows ($n \approx 1.35$) was attributed to a low fill factor of 30–40%: this is partially true, but those index measurements were later found to be affected by systematic error, which is discussed further in §4.1.1. With proper systematics correction, we now estimate the fill factor from fit indexes of 1 atm laminates to be approximately 85%, which is corroborated by the fill factor from density estimates (see Table 2.2).

A single atmosphere of pressure is insufficient to force the LDPE between the HMPE fibers. Thus we conducted a pressure sequence of laminates (holding the temperature steady) from 1 atm to 10 atm applied pressure, which are shown in Figure 2.4[9]. The fit indexes and estimated fill percentages from density for the pressure sequence are also reported in Table 2.2.

The lamination quality plateaus at pressures above 4 atm. For safety we typically laminate at 10 atm.

[8]I suspect that the low fill factor was only partially to blame for the high leak rate through these early laminates. The recipes for the early laminates all used very thin (4 mil or less) LDPE layers around the HMPE. If there is an inadequate volume of LDPE available to be wicked into the HMPE, voids are left in the laminating material, which further reduces leak tightness.

[9]My very first thin window project was the first pass at a pressure sequence of laminates; they were measured on the VNA (discussed in §4.1.1), and then leak tested by cutting them to fit on a KF-40 flange (hence the smaller circles in Figure 2.4) on our helium leak checker. It was those helium leak checks that convinced me that helium permeation rates through windows were not particularly useful; more on that in §2.5.

Table 2.2: Pressure sequence of laminates ultimate tensile strengths using the estimated HMPE fiber cross section, fit index of refraction (measurements discussed in §4.1.1), and estimated fill percentages from density measurements.

Baking Pressure [atm]	HMPE Fiber Ultimate Tensile Strength [GPa]	Fit Index	Estimated Fill Percentages [%]	Notes
1	3.5-3.9	1.437–1.484	87	Visibly poor lamination
2	N/A	1.557–1.565	94	Fine lamination
4	N/A	1.546–1.551	96	Good lamination
6	N/A	1.536–1.557	96	Good lamination
10	3.3-3.5	1.53–1.56	96	Good lamination

2.2.3 Comparison to Bulk Polyethylene

Now that we have determined that processing the laminates does not adversely affect the fiber strength, we must determine how thin a laminate can be safely deployed on a receiver. This is typically estimated with a safety factor (more on that in a shortly), which is dependent on the expected stress put on the window.

With the mechanical properties found from tensile tests we can compute the expected plastic deformation and stress of a window under atmospheric load. Assuming we are in the large angle deflection regime (δ) where the deflection is large compared to the thickness of the window, the initial elastic deformation is described by the flat plate equation:

$$\delta = K\Big(\frac{\Delta P\, R^4}{E\, h}\Big)^{(1/3)}. \tag{2.1}$$

While the maximum stress experienced under such deflection is given by:

$$\sigma_{max} = Z\left(\frac{E\Delta P^2 R^2}{h^2}\right)^{1/3} \tag{2.2}$$

where K or Z are a constant depending on the location of interest of the plate (K = 0.662, Z = 0.423 for the center), R is the radius of the plate, E is the flexural modulus (assumed equal to the elastic modulus in most materials), ΔP is the pressure difference (equal to one atmosphere), and h is the thickness of the material, all in SI units (Timoshenko 1959).

The ratio of the window's ultimate tensile strength to the predicted stress is the safety factor. We can iterate on the window thickness to achieve a desired safety factor, which we set to be at least three. §2.3 presents a detailed motivation for this choice. One way to compare the polyethylene window materials is to calculate the thickness required for each polyethylene window material to reach the required safety factor, using the reported material properties in Table 2.3 and Equation 2.2. These minimum safe thicknesses are reported in Table 2.3; the thickness for a safety factor of three in Table 2.3 is calculated using the lower value of ultimate tensile strength. Both bulk polyethylenes have similar lower limits on the ultimate tensile strength and therefore require an order of magnitude thicker material to achieve the same safety factor as the HMPE window.

It is important to note that these deflection, stress, and strain calculations use the full thickness of the laminate even though a significant fraction of the thickness is taken up by the much weaker LDPE. As such, the safety factor thicknesses reported in Table 1 are a lower limit. To confirm that the thin laminate windows have an appropriately high

Table 2.3: Strength properties (ultimate tensile strength, elastic modulus and strain at break) of Bulk HDPE, UHMWPE, and HMPE laminates, and the thickness required for a safety factor of 3, estimated using Eq. 2.2 and the lower of the reported ultimate tensile strengths.

Window Material	Ultimate Tensile Strength (MPa)	Elastic Modulus (GPa)	Strain at break (%)	Thickness for Safety Factor of 3 (mm)
HDPE	20 – 25	0.9	10	17.7
UHMWPE	22 – 40	0.7	30	13.6
HMPE Laminate (10 atm)	120 – 135	1.1 – 1.3	9 – 12	1.4

References. — Barkats et al. (2018); D'Alessandro et al. (2018)

margin, we conducted an empirical test of the laminate's strength, under the clamping boundary conditions that a deployed window would experience, which we discuss in the next section.

2.3 Safety Factor Measurement

Traditional tensile tests (like those reported above) only stress the woven fabric in a single direction (the direction of the pulling force). However, a three layer laminate will have HMPE fibers running in six separate directions over a circular aperture (each

layer is rotated 30 degrees from the one below it). A tensile test does not accurately capture the stress conditions of an actual window. The scaling of tensile tests to the full scale conditions thus remains uncertain: the maximum stress and safety factor calculations presented in §2.2 rely on the thin plate analytical relationship, which assume homogeneous material, and may not fully apply to composites. To mitigate this concern, we performed a full scale window failure measurement.

For safety, we performed this test hydrostatically. As water is nearly incompressible, minimal energy will be stored within the fluid as pressure is increased; energy is only stored within the deflected window and surrounding enclosure. Given the safety factors of 3 calculated in Table 2.3, we used standard utility water pressure (roughly 60 psi or 4 atm) as that was expected to exceed the ultimate tensile strength of these windows and produce a failure.

The BICEP Array frame and clamp were designed specifically to hold thin windows with a knurled pattern on the outer radius, with bolts running through the knurling to maximize concentration of force over relatively small points (Barkats et al. 2018). This knurled pattern on the frame and clamp works well but results in an unique window clamping boundary condition. We therefore used an actual frame and clamp for the hydrostatic test. The preexisting flange holes for circulating gas between the window and BOPP membrane were used to push the water into the chamber above the window. We sealed the water chamber with a 1/2" aluminum blanking plate and two 3/16" thick blended EPDM gaskets. The chamber was clamped with two sets of correctly torqued 1/4" bolts and C-clamps for additional clamping force[10], which are visible in a ring

[10]To protect the vacuum seal surface below the C-clamps, we placed HDPE squares as spacers between



Figure 2.5: Photo of the window at the highest pressure achieved during the hydrostatic test. The pressure gauge (center right) reads approximately 85 psi (5.7 atm), and there is water leaking out around the gasket seal on the left hand side.

around the window in Figure 2.5.

We reached 68 psi in the hydrostatic test using only the standard utility water supply. At that point we took a side-on picture of the deflected window, and estimated the deflection (and therefore radial strain) of the window (see §2.4 for a description of the calculations). The window had deflected approximately 75 mm. which corresponds to a radial strain of approximately 3%, well below the typical failure strain of 10%. In an attempt to produce a failure, we added compressed air to the system, increasing the pressure to 85 psi. At this pressure water began to leak around the gaskets, and the test was stopped. The end state of the window is shown in Figure 2.5, with the water pressure valve showing the final pressure on the right and the water leaking out around the gaskets on the left. Note that 85 psi is twice the predicted failure pressure for a

the C-clamps and the window frame.



Figure 2.6: Example photos of 'stress striping' behaviors in a tensile test sample and the post-hydrostatic test full scale window. [Left] Tensile sample example, with thumb for scale. There is a purple line added to aid the eye to the stripes. [Right] Full scale window after the pressure was relieved. There is a purple line added to aid the eye to the stripes.

safety factor of 3, a degree of uncertainty that emphasizes the importance of empirical testing of this novel structure.

After the stress was relieved from the window, and inspection revealed regions of the window that had white striping following what appears to be the bias direction of one layer of fabric (seen on the right of Figure 2.6). This is a similar pattern to what we saw on the tensile strength samples. I have termed them 'stress stripes,' at they appear to be stress concentrations around the regions of failure (the large white portion in the

sample strip on the left of Figure 2.6). It is typically seen along the bias direction of the tensile samples, which are likely the areas in the weave that experience the most strain. The stripes may be regions of delamination of the laminating LDPE or color change in the LDPE itself under high strain, such as those that occur when ripping an LDPE bag.

Given these results, we find that the safety factor of the thin windows is at least approximately 5.7 at sea level (1 atm) and 8.2 at South Pole (0.7 atm). Because the window will never experience a pressure over one atmosphere, we can be confident that the window will not fail under normal conditions. However, we must also explore other time-dependent factors which affect the mechanical properties of these window materials.

2.4 Creep

In addition to the initial elastic deflection, plastics will undergo a visco-elastic creep deformation under continuous load. The creep deformation is logarithmic with time; thus, most of the deflection occurs very quickly, and deflection out to longer time scales is simple to predict with a model, but not easily scalable from small scale measurements. We need to determine that the window will not creep into contact with the filters below (see Figure 1.9). The initial elastic deflection is described reasonably well by flat plate predictions (previously described in Eq. 2.1) from *Theory of Plates and Shells* (Timoshenko 1959).

We expect that the elongation of the window is azimuthally symmetric, and follows the shape of a parabola. Therefore, we can use the center deflection to predict the changed length along the arc of the parabola to find the radial strain (ie, the percent

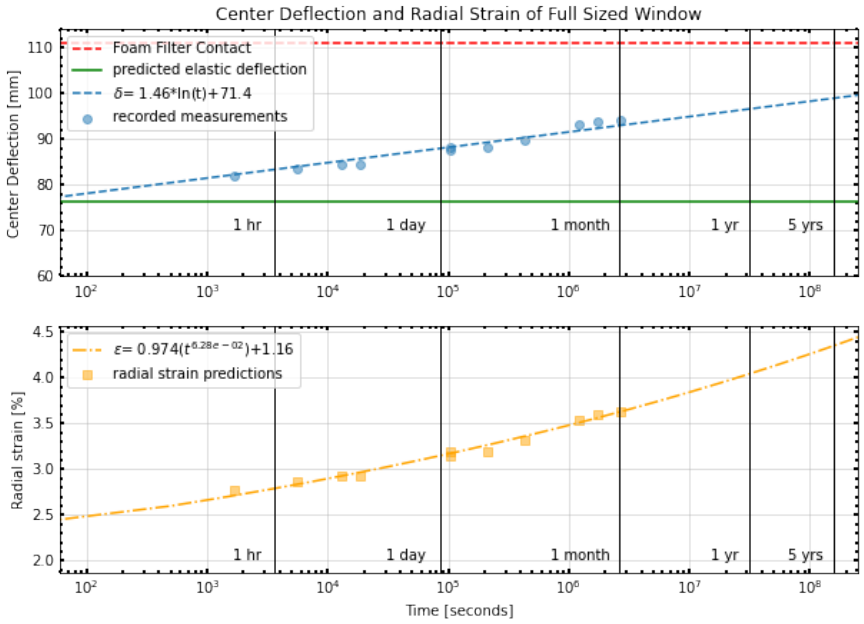


Figure 2.7: [Top] Measured center deflections of a window on a BICEP Array scale cryostat. The green solid line denotes the predicted elastic deflection from the plate equation Eq. 2.1 using a radius of 400 mm. The red dashed line denotes the deflection at which contact would be made with the foam infrared filters behind the window. [Bottom] Estimated radial strain using Eq. 2.3 and the measured deflections.

elongation). The relation is:

$$\begin{aligned} \text{Radial strain} &= (R_{\text{arc}} - R)/R \\ R_{\text{arc}} &= \frac{1}{2}\left[\sqrt{R^2 + 4\delta^2} + \left(\frac{R^2}{2\delta}\sinh^{-1}\left[\frac{2\delta}{R}\right]\right)\right], \end{aligned} \tag{2.3}$$

where R is the initial radius and δ is the center deflection. We can relate the estimated radial strain to a power law similar to that in Findley & Tracy (1974), obtained from observations of creep in polyethylene over sixteen years.

Figure 2.7 shows the measured deflection, fit deflection, and radial strain on a thin window on a BICEP Array scale cryostat over approximately a month. The radial strain is predicted to reach 4.5% after 5 years of operation, well below the failure strains of 9–12% (reported in §2.2). The window also remains out of contact with the foam filters below it.

2.5 Gas Permeation

As previously discussed, we pull a high vacuum[11] to limit convective heat transfer into the receivers. After a receiver is cooled it is valved off and the vacuum pump shut down to eliminate a source of electrical noise and vibration. We then rely on "cryopumping" (condensing residual gas on internal cold surfaces) to maintain vacuum through the observing season. If sufficient amounts of gas condense on the internal surfaces, the gas can move between surfaces (colloquially referred to as 'snowing'), potentially migrating to colder surfaces and increasing thermal load on the cryocoolers. It is therefore essential to ensure that the gas permeation rate through the window remains acceptably low.

Most gases can adsorb onto a polymer surface holding vacuum, random walk through the polymer, and desorb into the vacuum: this is the primary source of gas permeation[12] into the cryostat. Some gases are more likely to adsorb than others (water

[11]The distinction between rough vacuum (10^2 to 10^{-1} mbar), medium vacuum (10^{-1} to 10^{-3} mbar), high vacuum (10^{-4} to 10^{-7} mbar), and ultra-high vacuum (10^{-8} to 10^{-12} mbar) are somewhat arbitrary, but each require different techniques to accomplish (Ekin & Zimmerman 2007).

[12]Permeation is not the only source of gas entry into a cryostat; adsorbed gas onto inner surfaces can desorb over time, or trapped gas in internal polymers/between layers of MLI/caught in pockets

vapor is notoriously 'sticky'), and some materials are more permeable than others. The permeation rate is also dependent on the area on which a gas can adsorb. We use both small test chamber and full receiver pressure measurements to characterize the atmospheric gas[13] accumulation rate in a closed system. We then derive a normalized permeation rate (Q_{measured}) to compare between gas accumulation rates at different volumes as:

$$Q_{\text{measured}} = \frac{\Delta P_t}{\Delta t}\frac{V}{A} \tag{2.4}$$

where ΔP_t is the change in pressure over the change in time Δt, V is the volume of the closed system, and A is the area over which permeation occurs.

These rates can also be compared to other measured atmospheric permeation rates for HDPE, such as those found in Nguyen-Tuong (1993). At a given temperature, the permeation rate depends on the permeability of a material K, the pressure difference between the two volumes ΔP_v (for our purposes assumed to be atmospheric pressure at a given site), the thickness (h) and the area over which permeation occurs (A):

$$Q_{\text{predicted}} = \frac{K\Delta P_v}{hA}. \tag{2.5}$$

The fact that the thickness h is in the denominator is important to note: as we decrease

around screws can slowly escape into the main volume over time. These sources are finite, however, and permeation should dominate other sources of gas entry over long times at vacuum.

[13]Helium gas is typically used to test the vacuum tightness of cryostats. It is small and rare in the atmosphere, which makes it ideal to find and measure small leaks into a chamber using a mass spectrometer tuned to helium (a helium leak checker). Sometimes helium permeation rates are also used to set requirements on vacuum materials. I do not use helium permeation rates, as they are both difficult to measure (requiring placing a measured amount of helium gas on one side of a material) and difficult to scale to the air permeation that actually affects cryostats.

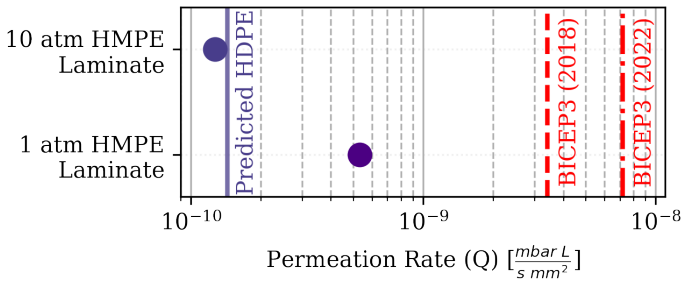


Figure 2.8: Measured permeation rates of small samples (circles) compared to the predicted rate with equivalently thick bulk HDPE from literature (solid line) and two limits on permeation from the unacceptable maximum rate of gas accumulation in two years; a marginally acceptable rate from BICEP3 in 2018 (dashed line) and and unacceptably high rate from BICEP3 in 2022 (dot-dashed line). The predicted HDPE value is calculated with Eq. 2.5, the permeation constant from Nguyen-Tuong (1993), and the laminate sample thickness. The two upper limit rates are determined through final warm pressure measurements (reported in Table 2.4) assuming that all the gas entered through the window over the time the BICEP3 cryostat was under vacuum. Both laminates are shown in Fig. 2.4. The high pressure (10 atm) laminate has a permeation rate closest to the bulk HDPE permeation rate.

the window thickness, we increase the expected permeation rate.

To help establish an acceptable maximum rate of gas permeation into our cryostats through the window, we use measurements of final warm pressures from two BICEP3 seasons that had cryogenic issues associated with excess gas accumulation. The receiver has a known volume of 905 L, and has a window open diameter of 730 mm (area 418×10^3 mm^2).

In 2018, a faulty o-ring seal at the window flange allowed gas to enter the cryostat

over the course of a little over one year. The gas accumulation was visible as temperature spikes on the 50 K shell. During the following summer season we warmed the receiver and measured the pressure inside the cryostat. The final pressure was 49 mbar after 359 days under vacuum. If all of that gas had instead permeated through the window, the window would have had a permeation rate of 3.4×10^{-9} mbar L s^{-1} mm^{-2}. The measured bolometer white noise did not increase significantly in 2018 compared to 2017, implying that this amount of gas accumulation did not adversely impact performance (see Chapter 5 for further discussion on BICEP3's white noise performance).

In 2022 we saw temperature variations again thought to be due to gas entry into the cryostat. These leaks were localized to a flange at the base of the cryostat. This rate of gas entry resulted in a significant increase in the measured white noise (again, see Chapter 5 for further discussion) for approximately 463 days. However, we cannot distinguish what fraction of the gas entered during the higher noise period and what fraction entered in the previous 626 days the receiver was under vacuum. Conservatively, we average the final warm pressure of 314 mbar over the entire period the receiver was under vacuum, recognizing that it is likely that the actual problematic permeation rate was higher than this average, and posit that a permeation rate through a thin window of 7.2×10^{-9} mbar L s^{-1} mm^{-2} will cause significant cryogenic and optical issues. For both the 2018 and 2022 cases, time at vacuum, final warm pressure and estimated upper bound window permeation rate are summarized in the bottom of Table 2.4.

To determine the actual permeation rate through the laminate window we tested two small scale samples (shown in the top of Figure 2.4) on a small vacuum vessel over a few days. Previously, we had noted an unacceptably high permeation rate in thin windows that were laminated at one atmosphere of pressure (Barkats et al. 2018; Eiben

Table 2.4: Measured upper bound permeation rates from thin windows on two receivers, one at sea level (PreSAT) and one at the South Pole (BA150), to compare to predicted window permeation rates and upper bound permeation rates associated with cryogenic issues.

Measured Upper Bound/Predicted Permeation Rates for Thin Window Cold Runs				
Receiver	Time at Vacuum [days]	Final Warm Measured Pressure [mbar]	Upper Bound Permeation Rate [a] $[\times 10^{-10}\ \frac{\text{L mbar}}{\text{s mm}^2}]$	Predicted Window Permeation Rate[b] $[\times 10^{-10}\ \frac{\text{L mbar}}{\text{s mm}^2}]$
PreSAT (Sea Level)	30	0.85	7.1	1.44
BA150 (South Pole)	341	2.1	1.5	1.21
Measured Upper Bound Permeation Rates Associated with Cryogenic Issues				
BICEP3 in 2018 (South Pole)	359	49	34	
BICEP3 in 2022 (South Pole)	1089	314	72	

[a] Estimated upper bound permeation rate through a window using Eq. 2.4, the reported time at vacuum, and final measured warm pressure, assuming that all receivers have the same internal volume (905 L) and window open diameter of 730 mm (area 418×10^3 mm^2).

[b] Predicted permeation rates of the thin windows on the receiver using Eq. 2.5 and permeation constant from literature (Nguyen-Tuong 1993).

et al. 2022). Therefore, we also compared the permeation rates of low pressure laminates and high pressure laminates, to confirm that higher pressure lamination improved the permeation rate. Figure 2.8 shows the results of the small scale permeation rate testing. The measured small scale permeation rates are compared to the predicted permeation rate for bulk HDPE of equivalent thickness (1.4 mm) using Equation 2.5 with the permeation constant found in Nguyen-Tuong (1993).

Both the measured and calculated window permeation rates in Figure 2.8 are significantly lower than the previously discussed problematic permeation rates estimated from two years with cryogenic issues on BICEP3 (red dashed and dot-dashed lines). However, the 1 atm laminate has a rate over three times higher than the 10 atm laminate; this is likely due to low laminate fill factor at low lamination pressure (as previously discussed in §2.2). The laminating LDPE is more effectively forced around the HMPE fibers at high pressure, resulting in better laminate consolidation and homogeneity and preventing gas from following the fibers to permeate through a thinner layer of LDPE. We will continue to use Eq. 2.5 and the Nguyen-Tuong (1993) permeation constant to predict permeation rates of windows, given the measured consistency between the 10 atm laminate rate and the predicted rate.

We have since run an observing season with a high pressure laminate window on BA150 at the South Pole (at altitude, with assumed 0.7 atm pressure difference) and warmed the receiver. We also measured the gas accumulation in another receiver (PreSAT) at sea level, though the run was over a shorter period of time (Petroff et al. 2024). Table 2.4 reports the time at vacuum and final pressures for each run. The final warm pressure can be used to estimate an upper bound permeation rate through a thin window, which can be compared to the predicted permeation rate for a given window, or

the upper bound window permeation rates associated with cryogenic issues.

Both of the receivers with thin windows cooled and remained cold successfully. These measured amounts of gas are also significantly below the previous problematic gas amounts. Roughly, we may estimate that the thin windows may allow a cryogenically problematic amount of gas into the cryostat after about twenty years, wherein the gas could be removed by warming, pumping out, and cooling again. The BICEP/*Keck* collaboration has very rarely kept a cryostat cold for longer than three consecutive years; permeation of gas through the window into the cryostat is unlikely to become a problem.

2.6 Environmental

There could be damaging forces arising from the environment around the window. People working above the cryostat could drop tools or other equipment. The environment at the South Pole is extreme, and may adversely impact window performance.

We have dropped a number of small, commonly lost items ranging from hex keys to screwdrivers (the heaviest object was 74 grams) on a window under rough vacuum (less than 5 mbar) from a height of approximately 60 cm in the lab environment (at room temperature). These drop tests saw no significant increases in leak rate or any other adverse effect, beyond the denting of the anti-reflection layer from the dropped objects. Additionally, during the 2024/2025 deployment season an aluminum bracket weighing approximately 0.5 kg was dropped on the BA220/270 window from a height of approximately 100 cm; the bracket punctured the BOPP membrane (which likely absorbed some amount of the kinetic energy) and dented the AR coat. The window

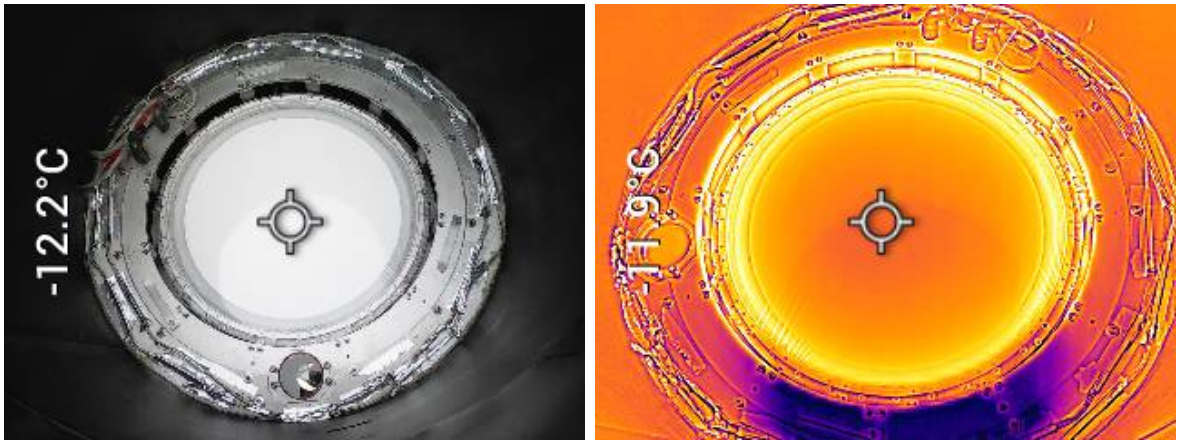


Figure 2.9: Pictures of the BICEP3 thin window taken with an IR camera in situ. [Left] A visible light image of the window on the interior of the forebaffle, with the noted temperature at the center cursor of -12.2°C. [Right] An infrared light image of the same window, with noted temperature at the center cursor of -11.9°C. Brighter colors denote hotter temperatures; the radial temperature gradient from the edge of the window to the center is visible.

seems to have survived the encounter.

I have also dropped a pound of dry ice on a window under vacuum and measured a large thermal gradient from 20°C to -78°C (293 to 195 K) across the window, again with no adverse effects to the vacuum performance.

We have also measured the equilibrium thermal gradient of both of the thin windows and a slab window at Pole during normal summer operations (when the cryostats were cold, room temperature nitrogen was circulating above the window, and with the forebaffles attached) with an IR camera, as shown in Figure 2.9. As the windows are all coated with a thick layer of PTFE[14] (which is known to be highly absorptive/emissive

[14]These are the anti-reflection coats discussed in the next chapter.

Table 2.5: Window surface temperature at different radial (R) points on cold receivers at the South Pole. Data taken by Marion Dierickx and Danielle Simmons with a hand-held IR thermometer: measurements are the mean and standard deviations on the four measurements taken in the cardinal directions.

Instrument	Center Temp (°C)	$\frac{1}{4}$R Temp (°C)	$\frac{1}{2}$R Temp (°C)	$\frac{3}{4}$R Temp (°C)	Edge Temp (°C)
BICEP3	-12±0.8	-12.3±0.5	-9±1.4	-3.5±1.3	-4.8±1.7
BA150	-19.6±1.3	-17.8±1.6	-14.5±2.6	-8.1±0.9	-1.0±2.8
BA30/40[a]	-17.6±0.6	-17.3±0.9	-15.9±2.5	-13.9±3.0	-5.6±4.0

[a]The BA30/40 window is a slab HDPE window.

to IR (BICEP2 and Keck Array Collaborations et al. 2015)), I expect that the camera was accurately capturing the true temperature of the windows. The BA150 window was slightly colder in the center, with an average temperature of -19.6±1.3°C, than the BA1 slab center temperature of -17.6±0.6°C. The BICEP3 window was the warmest at the center at -12±0.8°C. The different mount configurations may be the cause of the different center temperatures: BICEP3 is alone, while the two BA receivers share a mount.

The airflow rate between the BOPP membrane and the window is < 1 cubic foot per hour on BICEP3 based on nitrogen cylinder usage rates, which seems to be high enough to keep snow from accumulating on the membrane and the boundary of the window at an elevated temperature; thermometers on the window edge tend to hold between 270 K in deep winter and max at 285 K in the summer (-3°C to 12°C), which is significantly

warmer than the surrounding polar environment (temperature averaging around 235 K in the summer and 215 K in deep winter). The temperature gradient in the window is likely from radiative equilibrium with the environment (or some amount of convective cooling), as the window has cold optics below it and cold forebaffle and sky above it.

Chapter 3

Wide Bandwidth Anti-reflection Coats

Most of this thesis chapter originally appeared in the literature as

Eiben et al., *in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13102, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. J. Zmuidzinas & J.-R. Gao, 131022M*

Anti-reflection (AR) coats are a critical component of a transmissive optical system, reducing light lost to reflections and preventing spurious signal from reaching detectors. AR coats utilizing layered materials are typically quarter wavelength electrical lengths to produce destructive interference on reflection. Millimeter wavelength AR coats therefore exist in a unique macroscopic-but-still-thin regime; quarter wavelength optical lengths range from 0.3 mm to 3 mm for 300 to 30 GHz.

Instrumentalists designing AR coats for mm-wavelength optics must therefore source or produce materials that have the appropriate optical properties over the spectral bandwidth of their transmissive optics. Astronomical mm receivers are typically cryogenic, requiring (at least) vacuum windows and IR blocking filters, as previously discussed in Chapter 1 (Grimes et al. 2020; Bryan et al. 2018; Hui et al. 2018; BICEP/Keck Collaboration et al. 2022; Abazajian et al. 2019; Thornton et al. 2016; Iuliano et al. 2018). The CMB community also produces many refracting receivers with lenses, calling for AR coating solutions for lenses as well. (Hui et al. 2018; BICEP/Keck Collaboration et al. 2022; Abazajian et al. 2019; Thornton et al. 2016; Iuliano et al. 2018). Materials used for these transmissive optics are typically various polymers, such as many forms of polyethylene (PE), polypropylene (PP), polytetraflouroethylene (PTFE, commercial name Teflon), and nylon (Abazajian et al. 2019; Shitvov et al. 2022). Other common materials are alumina ceramic, which has been used in the past for lenses and filters, and silicon, also used for lenses (BICEP/Keck Collaboration et al. 2022; Dierickx et al. 2023; Iuliano et al. 2018; Thornton et al. 2016). Each material has unique challenges associated with producing and generating AR coats; alumina, for instance, has a high index of refraction (approximately 3.1) and a much smaller coefficient of thermal expansion (CTE) than most materials used to adhere AR coats[1] (Golec et al. 2020; Jeong et al. 2023). Plastic optics may have a better match to the CTE of the polymers used for hot melt adhesives, but due to their relatively low index (generally between 1.4 to 1.8) require even lower index materials to coat them, which may be difficult to source or produce (Lamb 1996; Dierickx et al. 2023; Eiben et al. 2022; Shitvov et al. 2022).

[1] Alumina is also notoriously difficult to machine, making it a poor candidate for machined AR surface structures (although, not impossible).

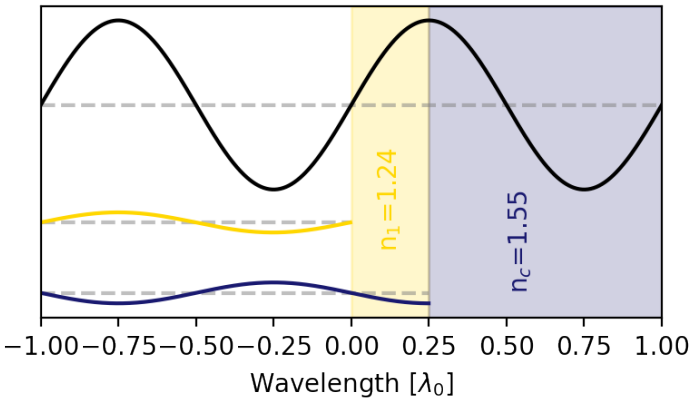


Figure 3.1: A cartoon representation of reflected waves off of a single layer AR coated material. A wave with wavelength λ_0 (black line) interacts first with the AR coat interface (yellow) resulting in a reflected wave with $\Gamma \approx 0.1$. Then the input wave interacts again with the core (blue) a quarter wavelength further on, resulting in another reflected wave with $\Gamma \approx 0.1$. Both reflected waves are 180° out of phase, and destructively interfere.

Other ways to eliminate reflection are possible, such as machined pyramids to smoothly transition light into the optic, or depositing intermediate index material directly onto the optic to smoothly increase the impedance into the optic (Golec et al. 2020; Jeong et al. 2023). Active research is being conducted in these areas, with exciting possibilities for future optics. For this chapter, I will focus on the parameter space provided by layered AR coats, as there remain optics where these newer solutions may require significant research and development effort (such as the thin vacuum window discussed in Chapter 2). Layered AR coats are used on nearly all transmissive optical components in BICEP Array receivers.

The layered AR coats for two mm BICEP Array receivers will be primarily discussed in this chapter. BA receivers have a large open aperture: the PE vacuum windows

require AR coats at least 70 cm wide to cover the illuminated area, and the filters and lenses require similar. Previous similarly scaled receivers deployed are BICEP3 (95 GHz), BA30/40 (dual band 30/40 GHz), and BA150 (150 GHz) (Hui et al. 2018; BICEP/Keck Collaboration et al. 2022). The high frequency BA receiver (220/270 GHz) is the highest frequency receiver (200–300 GHz) and deployed austral summer 2024/2025 to South Pole (BICEP/Keck Collaboration et al. 2024c). I will cover the AR coating design for (most of) the polymer optics in each of these receivers.

First, the let us discuss the theoretical ideal layered AR coats over a bandwidth and how achieved AR coats compare to the ideal. We will then cover the design of each of the BICEP Array AR coats in turn, starting with the lowest frequency bands to the highest. Then, an extended discussion on how high frequency AR coats are particularly challenging, a useful expanded PTFE (ePTFE) membrane for high frequency AR coats, and the design process for BA220/270 AR coats, with particular attention paid to the final BA220/270 thin window.

3.1 Ideal AR Coats

The goal of anti-reflection is to reduce the reflected power from the impedance change into an optic over the observed bandwidth. Light, being an electromagnetic wave, reflects when it encounters an abrupt change in the impedance of the medium it is traveling through. The impedance of a non-conductive medium is related to the permeability and permittivity of the medium, which goes like:

$$Z = \sqrt{\frac{\mu}{\epsilon}} = \sqrt{\frac{\mu_0 \mu_r}{\epsilon_0 \epsilon_r}} = Z_0 \sqrt{\frac{\mu_r}{\epsilon_r}} = Z_0 \sqrt{\frac{1}{\epsilon_r}} = \frac{Z_0}{n} \tag{3.1}$$

where Z_0 is the vacuum impedence, μ and ϵ are the absolute permeability and permitivity of the medium, while μ_r and ϵ_r are the relative permeability permittivity ($\mu_r = 1$ if the material is non-magnetic) and n is the refractive index. The index of refraction is directly related to both the relative permittivity and permeability by $n = \sqrt{\mu_r \epsilon_r}$, which simplifies to just $n = \sqrt{\epsilon_r}$ due to the non-magnetic nature of the medium. The refractive index can be complex (causing loss), but let's assume that it is lossless enough that the damping of the wave is negligible. The materials used for AR coats are thin enough that this is a reasonable assumption.

One can calculate the amplitude of the reflections of a plane wave with normal incidence with the reflection coefficient[2]:

$$\Gamma = \frac{n_1 - n_2}{n_1 + n_2} \tag{3.2}$$

It is a simple exercise to estimate the ideal AR coat for a single layer from the reflection coefficient: with a single layer AR coat, we have two impedance changes (two reflections) to work with. To get destructive interference of a wave, we'd like both reflected waves to be perfectly out of phase from each other (like the cartoon depicted in Figure 3.1). The travel distance of the light between both interfaces must equal some integer multiple of half the wavelength of the light. Since the reflection from the second interface travels the thickness of the coat twice, we need the travel distance to equal some multiple of a quarter wavelength of the light, such that the reflected wave travels

[2]Note that the reflection coefficient is the reflected wave amplitude which is useful for calculating the individual reflected waves of interfaces for AR coats. However, for our instruments we are most interested in reducing reflected power, which is Γ^2. Throughout most of this chapter we will be discussing reflections in fractional reflected *power*, either in percentages or decimal fraction [0-1].

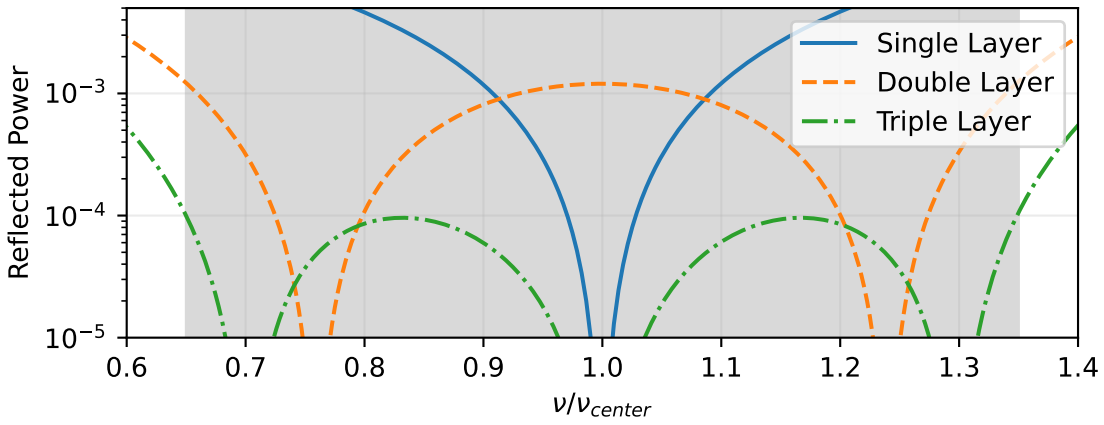


Figure 3.2: Reflections from ideal single, double, and triple layer anti-reflection coats for HDPE (n=1.55) over a fractional bandwidth of 0.7 (as required for the CMB-S4 MF) (Abazajian et al. 2019)

half a wavelength. The amplitude of each reflection also needs to be the same, such that they perfectly cancel. Using Eq. 3.2 and setting both reflected amplitudes (from both interfaces) equal to each other, we get:

$$\frac{1-n_1}{1+n_1} = \frac{n_1-n_c}{n_1+n_c} \tag{3.3}$$

$$n_c = n_1^2 \tag{3.4}$$

Therefore the ideal single layer AR index (n_1) is the square root of the core index (n_c). So if our core index is ∼1.55 (a typical polyethylene (PE) index) then the ideal single layer index would be $\sqrt{1.55} = 1.24$. Fully accounting for the lower order multiple reflection terms produces the same result.

The best possible broadband AR performance would be a very large number of quarter wavelength layers with small changes in index between adjacent layers. However, for practicality we require a minimal number of layers and therefore we need

to conduct careful optimization between the number of layers and desired performance. With a single layer it is possible to produce perfect destructive interference at a single wavelength, with two layers perfect destructive interference at two wavelengths, etc. This can be seen in Figure 3.2. The placement of those 'nulls' of destructive interference in frequency space sets the reflection profile over the band.

There are two design directives that one may take when exploring the ideal AR coat, related to the design goal. One may try to keep the reflections as flat as possible (the Binomial solution) or as low as possible on average over the band (the Chebyshev solution). For the wide bandwidths of our receivers, we choose to try to keep the band average reflection as low as possible over a top-hat band as the ideal case for our optics. Therefore, the ideal finite layered AR coat should move nulls to keep the maximum reflections even within the band. An identical problem is found in designing quarter wave transformers, as described in *Microwave Engineering* (Pozar 1997). We can find the idealized solution by using Chebyshev polynomials to identify the appropriate indexes for a given bandwidth (using the exact same approach that Pozar does with transformers).

3.1.1 Chebyshev Polynomials

Chebyshev polynomials are 'equal ripple' polynomials over a bandwidth. We can find the parameters of a multi-layer quarter wavelength AR coat by solving for the partial reflections off each layer. Chebyshev polynomials follow the formula:

$$T_n(x) = 2xT_{n-1}(x) - T_{n-2}(x), \tag{3.5}$$

where n is the order of the polynomial term T. The first three go like:

$$T_1(x) = x \tag{3.6}$$

$$T_2(x) = 2x^2 - 1 \tag{3.7}$$

$$T_3(x) = 4x^3 - 3x. \tag{3.8}$$

Chapter 5.7 of *Microwave Engineering* describes how one may utilize Chebyshev polynomials to produce even ripple reflections across a bandwidth (Pozar 1997). I will summarize the final design steps of an even and odd number of AR layers in the next section (the same procedures can be used for any higher order Chebyshev solution).

Example Chebyshev solution for double and triple layer AR coats

Starting with the formula laid out in Chapter 5.7 of *Microwave Engineering*, each solution requires initially calculating the bandwidth associated with the maximum ripple height (Pozar 1997). To calculate this, we iteratively test the bandwidth equation with an array of potential maximum ripple heights:

$$\sec\theta_m \approx \cosh\left[\frac{1}{N}\cosh^{-1}\left(\left|\frac{\ln n_c}{2A}\right|\right)\right], \tag{3.9}$$

where θ_m is the bandwidth in phase space, N is the number of layers, n_c is the core index, and A is the maximum reflected power ripple height.

For a double layer AR coat, we start with the polynomial:

$$2\Gamma_0\cos 2\theta + \Gamma_1 = A\sec^2\theta_m + A\sec^2\theta_m\cos 2\theta - A. \tag{3.10}$$

If we equate similar terms, we can solve for each of the partial reflections Γ:

$$2\Gamma_0 = A\sec^2\theta_m \tag{3.11}$$

$$\Gamma_1 = A\sec^2\theta_m - A, \tag{3.12}$$

which results in ideal indexes for each of the layers:

$$n_1 = Z_0 \exp\left[-\ln\frac{Z_0}{n_2} + 2\Gamma_1\right] \tag{3.13}$$

$$n_2 = Z_0 \exp\left[-\ln\frac{Z_0}{n_c} + 2\Gamma_0\right], \tag{3.14}$$

where Z_0 is the vacuum impedance.

Essentially the exact same procedure is followed for triple layer AR coats, with the initial polynomial starting as:

$$2\Gamma_0\cos 3\theta + 2\Gamma_1\cos(\theta) = A\sec^3\theta_m\cos(3\theta) + 3A\sec^3\theta_m\cos\theta - 3A\sec\theta_m\cos\theta. \tag{3.15}$$

The partial reflections become:

$$2\Gamma_0 = A\sec^3\theta_m \tag{3.16}$$

$$2\Gamma_1 = 3A(\sec^3\theta_m - \sec\theta_m). \tag{3.17}$$

Because of symmetry, we also assume that $\Gamma_3 = \Gamma_0$ and $\Gamma_2 = \Gamma_1$. Finally the indexes are:

$$n_1 = Z_0 \exp\left[-\ln\frac{Z_0}{n_2} + 2\Gamma_2\right] \tag{3.18}$$

$$n_2 = Z_0 \exp\left[-\ln\frac{Z_0}{n_3} + 2\Gamma_1\right] \tag{3.19}$$

$$n_3 = Z_0 \exp\left[-\ln\frac{Z_0}{n_c} + 2\Gamma_0\right]. \tag{3.20}$$

Table 3.1: Chebyshev single, double, and triple anti-reflection coat layer parameters on HDPE (n=1.55) for BICEP3 (ν_{center}=95 GHz, fractional bandwidth=0.25), and BA220/270 (ν_{center}=250 GHz, fractional bandwidth=0.4)

Layer	**BICEP3** (ν_{cen}=95 GHz, $\frac{\Delta\nu}{\nu_{cen}}$ =0.25)	**BA220/270** (ν_{cen}=250 GHz, $\frac{\Delta\nu}{\nu_{cen}}$ = 0.4)
Single	$n_1 = 1.245$, $t_1 = 634\mu$m	$n_1 = 1.245$, $t_1 = 241\mu$m
Double	$n_1 = 1.118$, $t_1 = 706\mu$m	$n_1 = 1.122$, $t_1 = 267\mu$m
	$n_2 = 1.386$, $t_2 = 569\mu$m	$n_2 = 1.382$, $t_2 = 217\mu$m
Triple	$n_1 = 1.058$, $t_1 = 746\mu$m	$n_1 = 1.061$, $t_1 = 283\mu$m
	$n_2 = 1.245$, $t_2 = 634\mu$m	$n_2 = 1.245$, $t_2 = 241\mu$m
	$n_3 = 1.465$, $t_3 = 539\mu$m	$n_3 = 1.461$, $t_3 = 205\mu$m

In Table 3.1, I report the found Chebyshev indexes and thicknesses for single, double, and triple layer AR coats on HDPE for BICEP3 (ν_{center}=95 GHz, fractional bandwidth=0.25) and BA220/270 (ν_{center}=250 GHz, fractional bandwidth=0.4). Note that the Chebyshev polynomials only change the index of a layer; the thickness changes to keep the layer at a quarter wavelength electrical length at the center frequency. For wider bandwidths, the multi-layer AR coats move the indexes slightly toward the $\sqrt{n_{core}}$ (the upper layer indexes go up, the deeper layer indexes go down). These impedance changes move the nulls away from the center frequency, therefore allowing a higher ripple.

The ideal Chebyshev indexes for common millimeter optical materials (such as bulk PTFE, nylon, alumina, etc) for common fractional bandwidths is documented in Appendix D. The above Table 3.1 is more 'complete' description of the ideal Chebyshev

solutions for these individual receivers, as it includes the required physical thickness to produce the quarter wavelength optical path.

3.1.2 Ideal AR Coats Compared to Achieved AR Coats

Practically, it is difficult to find materials that have precisely the properties required by the Chebyshev polynomials: for instance, in Table 3.1, the double layer Chebyshev solution for BA220/270 calls for a second layer with properties $n_2 = 1.382$, $t_2 = 217\mu$m. The index is possible to generate (potentially with a relatively high density sintered PTFE), but the thickness is difficult to acquire. Additional necessary adhesive layers will also change the effective electrical length at high frequencies. A typical adhesive layer such as 2 mil (50.8 μm) LDPE has an electrical length of 76.7 μm: equivalent to approximately 1/8 λ at 250 GHz (λ=1200 μm), which is not insignificant! Adhesion layers should be included when possible in the design process, as they may additionally help anti-reflection or may change the best solutions for other layers.

When attempting to identify potential AR membranes, we should have a way to estimate the index of a material from material properties commonly reported by manufacturers. For many PTFE membranes or polymer foams, manufacturers will report the filling factor or the density of the material. We can estimate the index of a given potential AR membrane with the density (or filling factor), and some knowledge of what the 'bulk' index is. There are two ways to go about estimating the membrane index, either with a simple linear relationship based on the filling factor (f) and densities:

$$f = \frac{\rho_m}{\rho_b} = \frac{n_m - 1}{n_b - 1}, \tag{3.21}$$

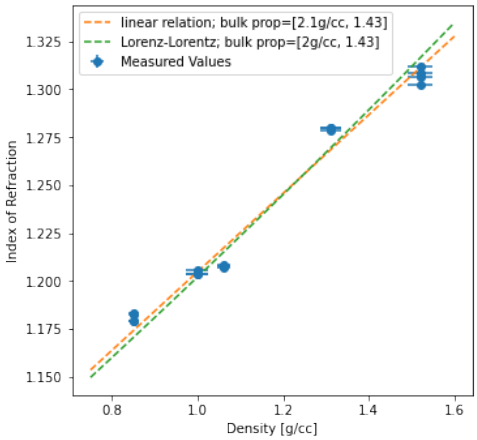


Figure 3.3: The Lorenz-Lorentz (Eq 3.23) and linear (Eq 3.21) relationship vs measured densities and indexes of various PTFE membranes. Uses the bulk material properties of $\rho_b \approx 2$ g/cc, $n_b \approx 1.43$ (Lamb 1996).

or with the Lorenz-Lorentz relationship, which goes like (Lorenz 1880; Lorentz 1881):

$$\frac{n^2-1}{n^2+2} = \frac{4\pi}{3}Na, \tag{3.22}$$

where N is the number density of the molecules in the material, and a is the polarizibility of the material. If we assume that the molecular number density and the polarizibility remain constant at different densities, we can rework the Lorenz-Lorentz relation to be a relation between the index of the bulk material and the different densities, like:

$$\begin{aligned} n_m &= \sqrt{\frac{2fC+1}{1-fC}} \\ C &= \frac{n_b^2-1}{n_b^2+2} \\ f &= \frac{\rho_m}{\rho_b}, \end{aligned} \tag{3.23}$$

where n_m is the membrane index (lower density than the bulk), n_b is the bulk index (for PTFE ~ 1.43 at 100 GHz (Lamb 1996)), and ρ_m/ρ_b are the membrane and bulk densities, respectively (the ratio of the densities, f, is also the filling factor). I have

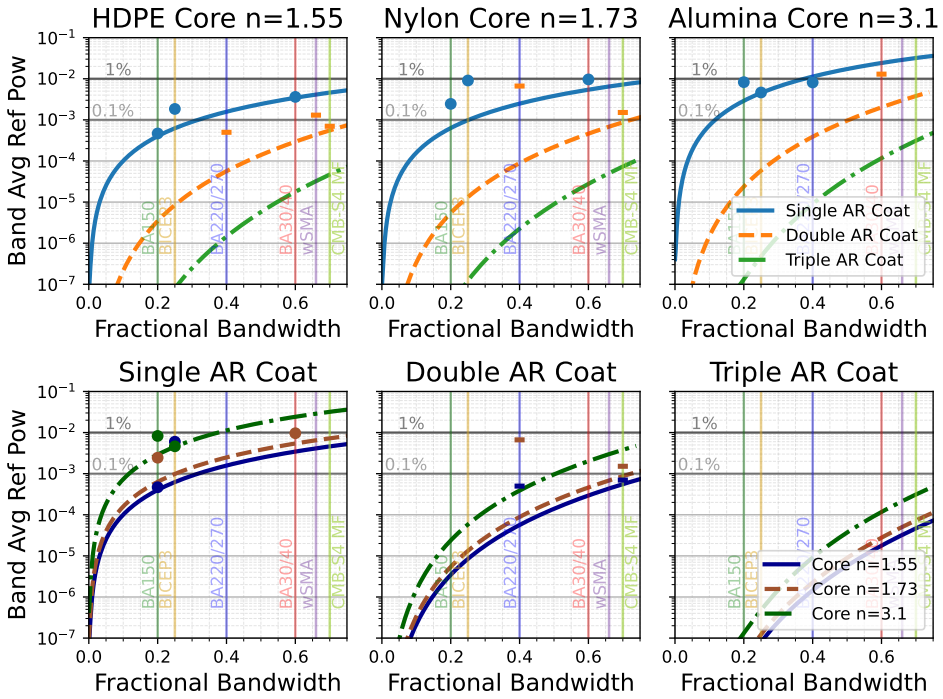


Figure 3.4: Average in-band reflections for ideal anti-reflection coats on infinitely thick cores of HDPE (n=1.55), nylon (n=1.73), and alumina (n=3.1). Points represent achieved AR coats for each material on each instrument where shape denotes rough equivalent number of layers. Circles are single layer equivalent AR coats, lines double layers.

tested the power of both relationships (the linear and the Lorenz-Lorentz) with different densities of membrane PTFE, and found a very good match, as shown in Figure 3.3.

We compare the achieved AR coats on recent optics in the current generation of BICEP/*Keck* instruments (and a few other instruments) to the ideal Chebyshev AR coats. Recent process improvements for the ePTFE for plastic optics have allowed for better control of the final AR layers produced, such that the single layer AR coats for BA150 produce reflections close to the ideal limit (Dierickx et al. 2023; Eiben et al. 2022). Some of the improved performance at wider bandwidths are related to new

materials and new design processes, which will be discussed further in §3.3. The high frequency (BA220/270) AR coats will be discussed in §3.4.

The three core materials in Figure 3.4 (HDPE, nylon, and alumina) are the three materials within the BA receivers that require layered AR coats. BA lenses and windows are made of PE, the 50 K filter is made of alumina, and the 4 K filter is made of Nylon 6 (Dierickx et al. 2023). One instrument (wSMA) is not a wide aperture CMB receiver but rather a wide bandwidth interferometer receiver (Grimes et al. 2020, 2024); we are collaborating with the wSMA team to make a high modulus polyethylene (HMPE) AR coated window for their new wider bandwidth. Further discussion of HMPE window optics (such as birefringence) will be covered in Chapter 4. The CMB-S4 Mid-Frequency (MF) optics (combined 95 GHz and 150 GHz bands) are expected to be tested within the PreSAT receiver (Petroff et al. 2024). The CMB-S4 MF AR coats in Figure 3.4 are proposed coats with known materials; this is discussed further in my conclusions (Chapter 7).

3.2 Low Frequency AR Coats (25–45 GHz)

The production run of the BA30/40 optics just predates my joining the Kovac lab (optics were shipped in early 2019, while I joined properly in late 2019). I did, however, revise the fits for the reflection measurements of the polymer BA30/40 AR coats (shown in Figure 3.5, and reported in Dierickx et al. (2023)). For completeness, let us briefly discuss the low frequency AR coats.

The low frequency BICEP Array bands have an un-intuitively large bandwidth:

Table 3.2: Generated AR coats for the current generation instruments in the BICEP/*Keck* collaboration, and CMB-S4 MF. BICEP Array (BA) bands should be similar to CMB-S4 bands (BA30/40 to CMB-S4 LF, BA220/270 to CMB-S4 HF). All AR layers' electrical lengths are summed and reported for each material, and those are used to estimate the equivalent number of quarter wavelength layers.

Instrument	**Center Freq [GHz]**	**Material**	**Quarter Wavelength [mm]**	**Summed Actual AR Electrical Length [mm]**	**Estimated Equivalent Layers**
BA150	150	HDPE	0.5	0.51	1.02
		Nylon		0.63	1.27
		Alumina		0.44	0.88
BICEP3	95	HDPE	0.79	0.62	0.78
		Nylon		0.78	0.97
		Alumina		0.79	1.00
BA220/270	250	HDPE	0.3	0.58	1.93
		Nylon		0.72	2.41
		Alumina		0.39	1.28
BA30/40	35	HDPE	2.14	1.99	0.93
		Nylon		2.15	1.01
		Alumina		4.32	2.02
CMB-S4 MF	125	HDPE	0.6	1.15	1.91
		Nylon		1.3	2.17

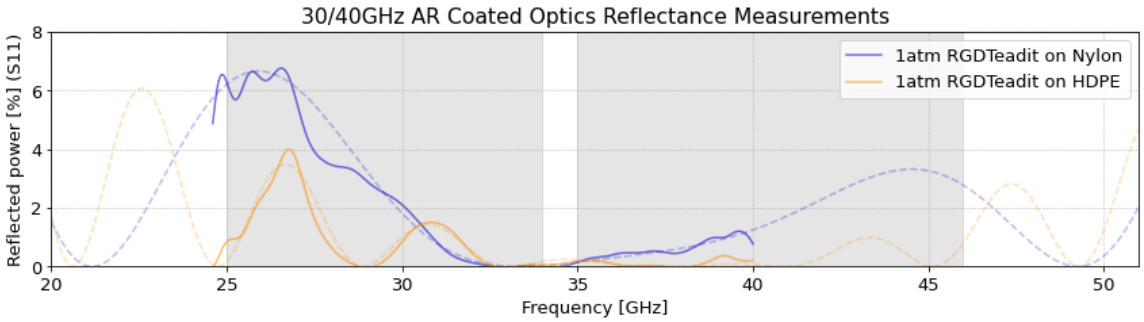


Figure 3.5: Calibrated Vector Network Analyzer free space S11 reflection measurements for the RGD Teadit AR coats on polymer optics for BA30/40. Solid lines are calibrated VNA data, dashed lines are fits to the data.

though the full dual band only extends from 25 to 46 GHz, the fractional bandwidth, due to the low frequency, is three times as large as the BA150 band. This limits the lowest possible single layer reflection band-averaged power significantly, as seen in Figure 3.4.

We used 1.6 mm (1/16") thick Teadit ePTFE (descussed further in §3.3.1) for both the nylon filter and the HDPE lenses in BA30/40. The measured parameters for these AR coats (found through fits to the VNA measurements of optical samples shown in Figure 3.5) are $n_1 = 1.231 \pm 0.002$ and $t = 1.618 \pm 0.001$ for the HDPE, and $n_1 = 1.226 \pm 0.002$ and $t = 1.632 \pm 0.001$ (Dierickx et al. 2023). Given that these were the first 30/40 GHz AR coats produced by the BICEP/*Keck* collaboration, these AR coats perform well. The thicker AR coat allows for a very reasonable tolerance around the ideal quarter wavelength electrical length, which also, at Teadit densities, luckily requires physical thicknesses close to a common inch fractional thickness. However, in retrospect, the BA30/40 AR coats would be improved by having more layers than a single coat.

For the higher frequency bands we are not as lucky; the quarter wave electrical lengths are not close to a common inch fractional thicknesses, and the thinner required

layer are also more difficult to manufacture. We must, then, get more creative.

3.3 Mid Frequency AR Coats (85–170 GHz)

Previous mid and high frequency receivers in the BICEP/*Keck* collaboration have used sintered PTFE (sPTFE) from Porex[3]: the widths of Porex rolls are maximum 33 cm (13 in). Utilizing Porex would require at least two seams to fully cover a BA aperture, which is difficult to work over curved surfaces and prevent potential polarizing effects.

There are other non-Porex sPTFE manufacturers who may be able to produce the mid-frequency AR coats. The relatively low densities required for single layer AR coats, however, are more easily achievable with expanded PTFE (ePTFE), and were previously used for the HDPE/UHMWPE windows on BICEP3 (BICEP/Keck Collaboration et al. 2022). We have also found another interesting trait of ePTFE that sPTFE lacks: ePTFE is compressible.

3.3.1 Teadit: Compressible PTFE

Teadit is an expanded polytetrafluoroethylene (ePTFE) membrane. ePTFE is stretched biaxially from a resin: this produces very small nodules of highly crystalline material surrounded by fibrils of amorphous polymer strands (Ebnesajjad 2016). Teadit is typically used as a gasket material for corrosive fluid pipes, and it comes in a few standard thicknesses and densities. It is expanded in sheets 150 cm (59") square, which is twice as wide as is required for BA optics, so we typically order quarter sheets. However,

[3] `https://www.porex.com/`

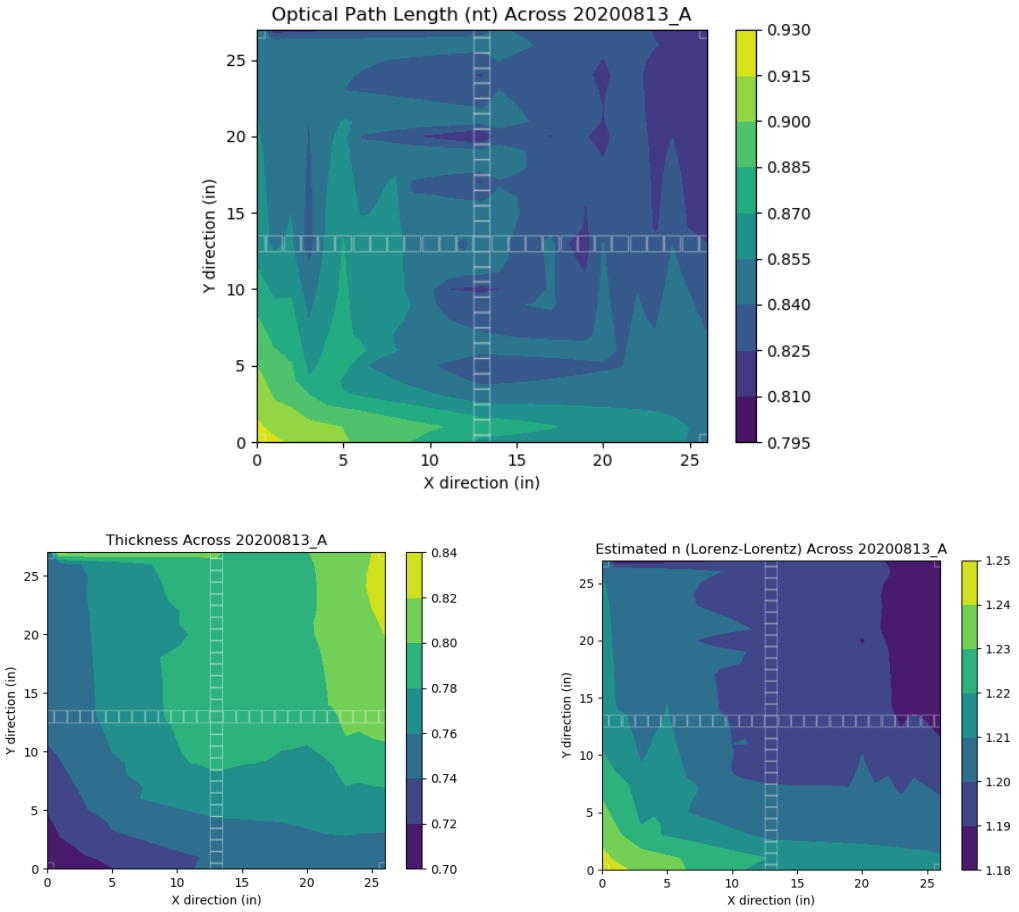


Figure 3.6: [Top] Optical path length (nt, measurements shown in subfigures below) over a sheet of Teadit. White squares denote where inch squares were cut and thickness and density measurements were taken. [Bottom Left] Thickness measurements over a sheet of Teadit. Spaces between measurement squares are linearly interpolated between the three closest x, y and corner measurements. [Bottom Right] Index estimates from physical densities and the Lorenz-Lorentz relation/Eq. 3.23. Spaces between measurement squares are linearly interpolated between the three closest x, y and corner measurements.

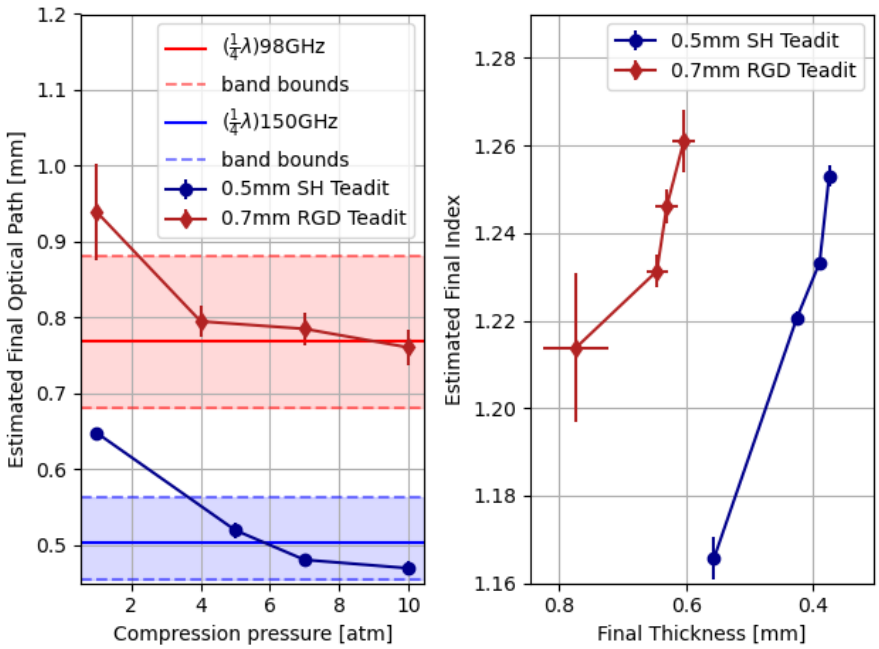


Figure 3.7: [Left] Compression pressures and associated estimated final optical path of the two formulations of Teadit used for AR coats in BICEP3 (95 GHz) and BA150 (150 GHz). [Right] Final thicknesses and estimated final indexes (estimated from physical densities and the Lorenz-Lorentz relation/Eq. 3.23) of the two formulations of Teadit on the left.

as ePTFE is stretched biaxially (stretched on one axis, then the perpendicular axis), the full sheet exhibits a mounding behavior, where the center of the sheet is thicker and less dense than the edges. An example of this behavior on a single sheet of RGD Teadit is shown in Figure 3.6. This is problematic for AR coats, as electrical length (or optical path length) differences will impact the performance of the coating.

Teadit makes two different formulations of their ePTFE gasket sheets: SH and RGD. SH Teadit is currently (as of writing this thesis in March 2025) the only formulation listed on the manufacturer's website. It is slightly lower density than RGD Teadit.

RGD Teadit therefore produces indexes closer to the ideal index of a single layer HDPE AR coat (n≈1.24) out of the box, which is why it was used for previous AR coats on PE optics in BICEP3 and BA30/40. However, RGD must be special ordered from the manufacturer and may not be as available currently as it was in the past.

When ePTFE is heat compressed above a glass transition temperature (a second order phase transition) and cooled while compression is still applied, it will remain in its compressed state. I suspect that above the glass transition temperature of PTFE the polymer is more susceptible to permanent deformation and therefore the compression gets 'locked in'. We use this to tune the compressed properties; as the pressure is increased, the thickness of the sheet will go down and the density will go up. The optical path length (or the electrical length) of the material tends to asymptote at high pressure. However, by carefully choosing the starting material properties, we can fine-tune the final properties with pressure. This gives near unprecedented control over the final AR layer.

Figure 3.7 shows the optical path length compression curves of two types of Teadit; 0.5mm SH Teadit used for BA150 polymer AR coats and 0.7mm RGD Teadit used for the BICEP3 thin window. Teadit compression also reduces the variability across the sheet (as seen in the error bars on the properties in Figure 3.7), which is very helpful in parameter tolerance control.

3.3.2 Using Teadit

As the optical path length scales with the compression pressure, that is the knob we can tune to generate the best possible AR coat for a PE optic. Determining the best

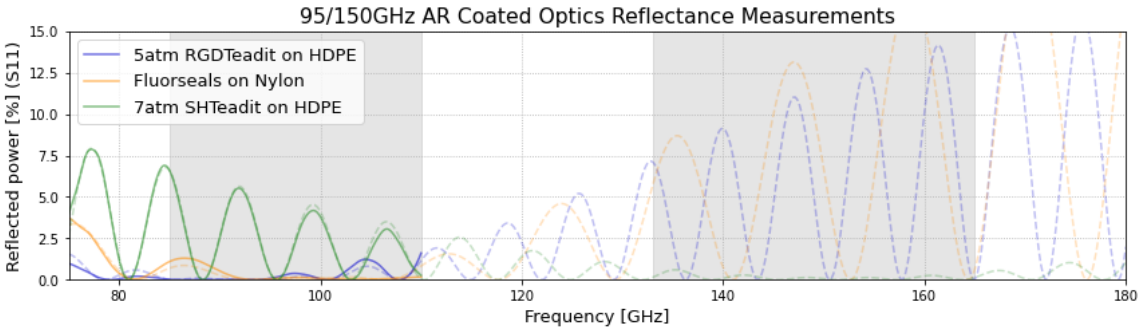


Figure 3.8: Vector Network Analyzer free space S11 reflection measurements of AR coated optics samples for 95 and 150 GHz optics. Solid lines are calibrated VNA data, dashed lines are fits to the data.

pressure to compress the single layer AR coating for our optics is simple for the infinite core case (a useful model for designing an AR coat for a variable thickness core like a lens): the ideal pressure is obvious from the left side of Figure 3.7. For the BA150 lenses we aimed for 7 atm compression, which, in sample AR coated HDPE free-space VNA measurements, produced very low reflections (as seen in Figure 3.8).

For our nylon filter (with an index of refraction significantly higher than HDPE at around 1.73), we used the same SH Teadit, but compressed at 10 atm (Lamb 1996). This produced an AR coat with a higher index (n≈1.27), but not the ideal $\sqrt{1.73} = 1.31$. However, given that we only have one nylon filter, the slightly non-ideal AR coat (producing a band average reflection of 0.42%) was determined to be acceptable.

Thin windows, however, add an extra dimension of complexity. As previously described in Chapter 2, we can control the thickness of the window, which we can use to essentially produce another AR layer. However, windows need to have integral number

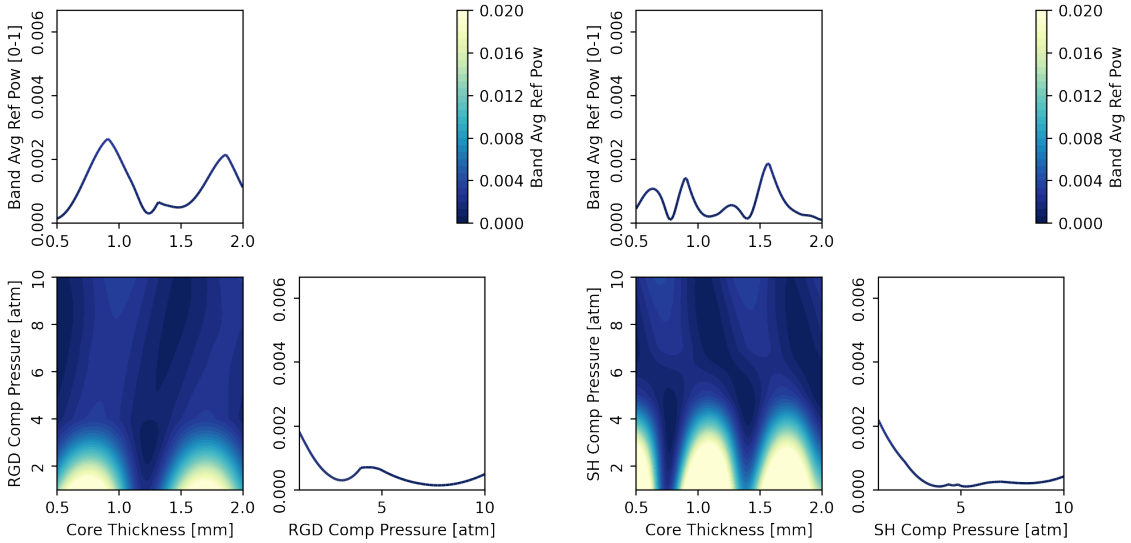


Figure 3.9: Design corner plots for thin window AR coats on BICEP3 (95 GHz, left plot with RGD Teadit) and BA150 (150 GHz, right plot with SH Teadit). Parameters optimized over are the window core thickness and Teadit compression pressure, with actual optical parameters derived from the linear interpolations shown in Figure 3.7.

half wavelength electrical lengths instead of quarter wavelength[4].

We now have two parameters to optimize over for the mid-frequency thin windows: the compression pressure and the window core thickness. We have developed a design tool to help visualize the parameter space; the design corner plots shown in Figure 3.9. We iteratively test all combinations of the parameters (in this case, the compression pressure of the single layer AR coat and the thickness of the thin window for BICEP3 on the left of Figure 3.9 and BA150 on the right of Figure 3.9) to find the combinations

[4]In a traditional ideal AR layer both interfaces have the same sign impedance jump (from low-to-high index or high-to-low index). This results in the same sign reflection at both interfaces (remember Eq.3.2). But in a window (or any core material) the two interfaces will always result in opposite sign reflections, requiring in a larger path length to produce destructive interference.

Table 3.3: Design, measured and modeled AR coat parameters for the BICEP3 and BA150 deployed thin windows. Thicknesses were measured with calipers, modeled thicknesses and indexes are fit parameters to VNA measurements, such as those in Figure 3.10.

Receiver	Core Thickness [mm]		AR Thickness [mm]		AR Index	
	Design	Measured	Design	Modeled	Design	Modeled
BICEP3	1.45	1.38	0.63	0.61	1.25	1.25
BA150	1.20	1.18	0.39	0.40	1.23	1.23

that produce the lowest band average reflection.

The BICEP3 replacement thin window was measured with the free space VNA in two bands (in band WR 10 and out of band WR 6.5). The signal is coupled to free space with a rectangular gain horn. The beam is then collimated with a 90° off axis parabolic mirror, with the beam waist placed where two posts allow us to hold large optical materials up to the flat surface the posts provide. We normalize the measurement with a short measurement of a flat aluminum plate at the measurement surface (see Figure **??**). Time gating the S11 (reflection) signal (calculated by our HP 8510C Network Analyzer) allows us to remove spurious reflections from the system, such as those from the horn. We found that window to have an AR coat very close to the ideal parameters as outlined in Table 3.2. However, the slight discrepancy between the actual and designed core thickness does push the nulls away from the ideal, resulting in larger than desired reflectivity on the lower edge of the band and a band average reflection of 0.08%. The final BA150 window better matched all of the desired AR parameters, thus achieving a better average in band reflection of 0.04%.

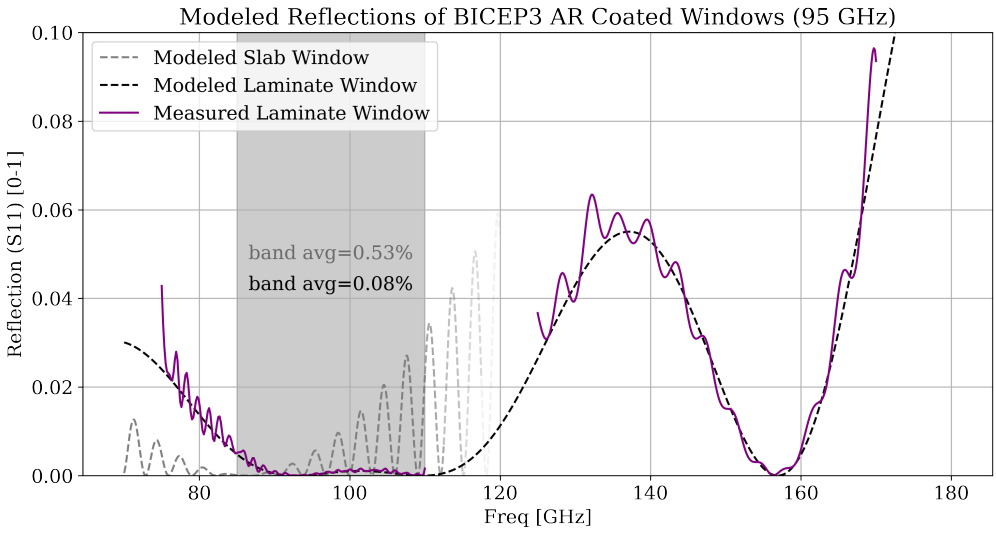


Figure 3.10: Measured VNA S_{11} reflections of the BICEP3 thin window (solid purple lines, both in-band WR 10 and out of band WR 6.5), and modeled reflections of the laminate (dashed black line) and the slab window (dashed gray line). The average in-band reflection of the AR coated slab (deployed on the receiver (BICEP/Keck Collaboration et al. 2022)) is estimated to be 0.53% while the average in-band reflection of the AR coated laminate is estimated to be 0.08%. The in-band and out-of-band measurements on the laminate window show excellent consistency with the modeled reflection.

3.4 High Frequency AR Coats (200–300 GHz)

As shown in Table 3.2, the high frequency receivers (BA220/270 and wSMA) have the thinnest quarter wavelength. This presents a tougher challenge to find materials that can have consistent thicknesses and densities (indexes of refraction) over the area of an optic. The area is particularly challenging for BA220/270 (200–300 GHz), which has an clear aperture diameter of over 70 cm (Hui et al. 2018). The wSMA has a much larger fractional bandwidth, so though the aperture size is much smaller, multi-layer AR coats

are required to stay below a reasonable reflection (Grimes et al. 2020).

Single layer AR coats could work for BA220/270 but would require strict tolerances on both the index and the density to keep the single null centered. A reasonable tolerance could be plus or minus 5% off the ideal single layer parameters. For HDPE with an ideal single layer ePTFE coat, the thickness would be 240 ± 10 μm and the density ρ=1.320 ± 0.066 g/cc (equivalent to n=1.245 ± 0.013). Those 5% tolerances can change the band average reflection by 25%. All relevant parameters (thicknesses, densities/indexes) are difficult to achieve with traditional high frequency AR materials to the tolerances desired. And while it is possible to source sPTFE at wider widths than Porex provides, the thicknesses required for high frequencies are not practical to skive[5] to the tolerances required.

Multi-layer AR coats are more resistant to parameter tolerances and can generate lower reflections than single layer AR coats. For these reasons, we decided to attempt multi-layer AR coats for BA220/270.

3.4.1 DeWAL: Stackable ePTFE

DeWAL ePTFE membranes (hereafter referred to as DeWAL) can be produced with densities "ranging from 0.2 to over 1.0 g/cc" (equivalent indexes of 1.035 to 1.18), thicknesses "from 0.001 to 0.010 in" (25.4 to 254 μm), and widths up to 30 in (76.2 cm) (DeWAL 2023). These thicknesses and densities are unusually low, and I am unaware of

[5]Sintered PTFE (sPTFE) is formed by sintering—heating carefully to only melt the outer surface—powder to form a low density cylinder called a billet. The billet is then skived—spun very fast on a special machine, then shaved with a sharp edge—to produce a thin membrane.

a competitor that produces membranes with a similar combination of properties. The wider widths of DeWAL seem to be limited to the lower density and thinner material, but crucially this ePTFE can be stacked.

When DeWAL is stacked and compressed at temperatures above 50°C, the stacked layers will bond. This bonding process works best at pressures above 3 atm (∼45 psi). If the membrane has already been processed (run through a pressurized temperature cycle, with a bulk PTFE release layer helping prevent adhesion to other surfaces), subsequent attempts at bonding will fail[6]. Shown in Table 3.4 are 90° peel test strengths of various processed forms of DeWAL stacks (some DeWAL only, some testing adhesion strength). We call DeWAL that has never been compressed 'Naive' and DeWAL that has been run through a compression cycle at least once 'Pre-compressed'. We also report testing of adhesion layer peel strengths: both hot melt adhesives tested have been used as adhesives on cold optics in BA in the past. Low density polyethylene (LDPE) is typically used as the adhesive for other PE optics, such as the vacuum window and the lenses. E100 was used as the adhesive for the BA2 alumina filter AR coat. As shown in Table 3.4, the Naive DeWAL has peel strengths an order of magnitude above the Pre-compressed DeWAL. The adhesion layers have a peel strength an additional order of magnitude above the Naive.

Similar to other ePTFE AR membranes that we have used, DeWAL changes thickness (and therefore density) when compressed at pressure (Dierickx et al. 2023;

[6]This caused the bubbled AR coat on the second attempt at a science grade BA220/270 window, as attempts to add layers to the too thin stack peeled up when the window deflected. SN#20231001 is deployed as a non science grade back-up window, as the bubbled layers of DeWAL had to be peeled off and the subsequent AR performance is terrible (see Figure 3.14).

Table 3.4: Peel strengths of various kinds of DeWAL compressed together.

DeWAL Layer Recipe	**Peel Strength [g/cm]**
2-layer of Naive compressed at 1atm 135°C	1.96–3.05
2-layer of Naive compressed at 3atm 135°C	3.9–5.9
2-layer of Naive compressed at 10atm 135°C	5.9–7.9
2-layer of Pre-compressed (3 atm) at 1 atm 135°C	0.25–0.5
2-layer of Pre-compressed (3 atm) at 1atm 135°C compressed together with a Naive layer of DeWAL between	0.15–0.61
2-layer of Pre-compressed (3 atm) at 1atm 135°C compressed together with 1mil LDPE between	60–80
2-layer of Pre-compressed (3 atm) at 1atm 135°C compressed together with 1mil E100 between	30–40

Eiben et al. 2022). Its lower initial density, however, means that even compression at 10 atm does not get the density very high. For example, HDPE's index is around $n_c = 1.55$, so the ideal single layer index is $n_{AR} = \sqrt{1.55} = 1.24$, while the highest DeWAL index measured is ∼1.2 (see Figure 3.11). DeWAL therefore requires at least an additional set of layers to push the AR electrical lengths closer to increments of quarter wavelength.

Luckily, the low index and thickness provide an advantage in a multi-layered coat. DeWAL is so low index at the lowest compression it can act almost like a very thin air gap, which allows for 'pseudo-layers' of alternating high and low index materials. Given that typical adhesion materials, such as LDPE, have much higher indexes, the

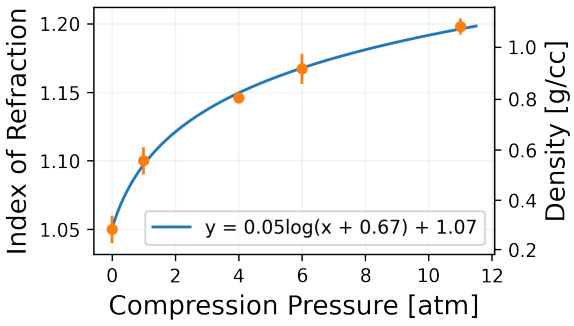


Figure 3.11: DeWAL density and estimated index at different compression pressures. I estimate the index from the density using the Lorentz–Lorenz relation, assuming that the bulk index and density of PTFE is 1.43 and 2.2 g/cc, respectively. (Polder & van Santeen 1946; Lamb 1996)

combination of DeWAL and the required adhesive can generate a good AR coat together. To accomplish this we must explore a large parameter space to find the right combination of layer parameters that we believe will produce the most robust AR coat. The next sections explore how I used DeWAL to design AR coats for BA220/270

3.4.2 Using DeWAL

BA220/270 Cold Polymer Optics

BICEP Array 220/270 GHz (BA220/270) is the highest frequency receiver in the BICEP Array. The frequency band covers from 200 to 300 GHz, which also makes it one of the broader fractional bandwidths covered in Figure 3.4. The receiver has three main optical materials that require AR coats: polyethylene (PE) lenses and a window, a nylon filter, and an alumina filter. We will focus mainly on the AR coats for the PE optics, as the

nylon filter is a very similar process, while the alumina filter is completely different[7].

We assume that the index of high density polyethylene (HDPE) is on the higher end of previous measurements at around 1.55[8] (typical measurements put HDPE indexes between 1.53 to 1.55) (Lamb 1996; Shitvov et al. 2022). This increases the ideal indexes for AR layers slightly, which makes it slightly more difficult to use DeWAL (a low index material). The best solutions are typically pretty similar for the small range of likely HDPE indexes, and therefore if the true index of HDPE is slightly lower than we assume, the AR coat should still work well.

The initial full roll of DeWAL we acquired had a central 'mound' of thicker area across the width (see Figure 3.16 and discussed further in § 3.5): this limited the minimum thickness at which we were willing to use stacked DeWAL. If we used too few layers then we would not be able to rotate the thickness profile enough to average out the mound. Therefore, we resolved to using an initial thick layer of stacked DeWAL (AR1) and some other material for the second layer (AR2).

From the parameter variation plots (as in Figure 3.12), there was a clear portion of the high index parameter space that would work for $n_{AR,2}$, all the way up to the index of bulk PTFE (n$\sim$1.43) (Lamb 1996). Skived sheets of bulk PTFE are readily available, if potentially difficult to laminate. Subsequently, we scuffed up the surfaces of the PTFE sheets prior to lamination as much as possible with scour pads. I was unaware, at the time, of an internal review within our group that found that coarse sandpaper works

[7]Also, the alumina AR coats traditionally have not been made by the Harvard group for BICEP.

[8]This is a reasonable assumption to make as the cold optics will actually have a higher index than at room temperature due to contraction and the Lorenz-Lorentz relation/Eq.3.23.

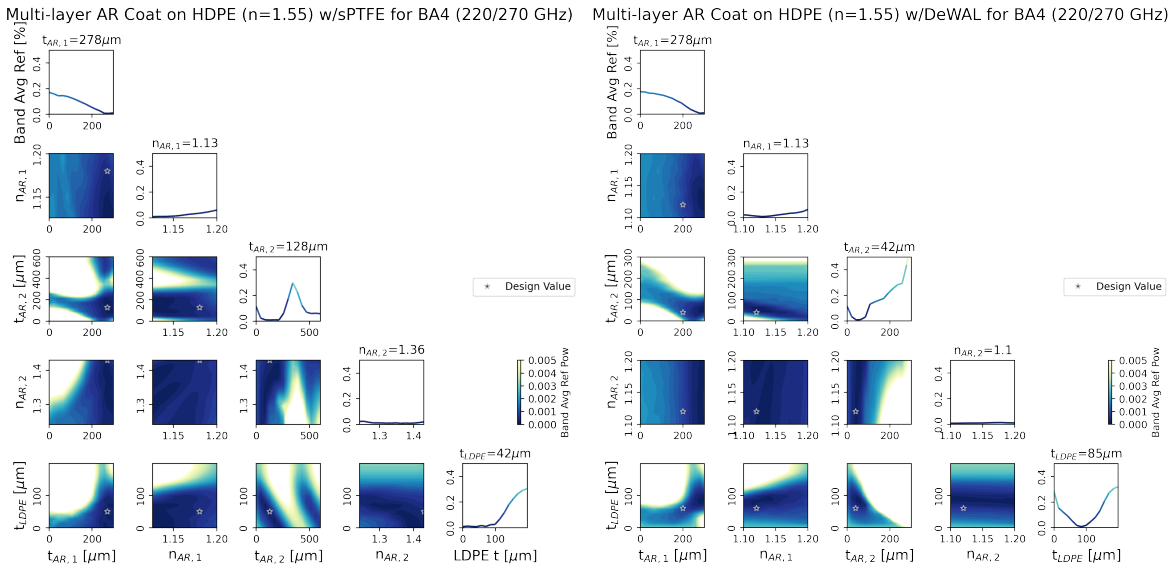


Figure 3.12: [Left] Design corner plot for multilayer High Frequency (HF) HDPE optics using DeWAL, sPTFE, and LDPE layers (BA220/270, 200–300 GHz). Layers go: AR layer 1 ($t_{AR,1}$,$n_{AR,1}$) stacked DeWAL, LDPE adhesion layer (t_{LDPE}), AR layer 2 ($t_{AR,1}$,$n_{AR,1}$) is sintered PTFE, LDPE adhesion layer (t_{LDPE}), core. This design was used for the BA220/270 lenses; stars denote the chosen design parameters. [Right] Design corner plot for multilayer High Frequency (HF) HDPE optics using only DeWAL and LDPE layers (BA220/270, 200–300 GHz). Layers go: AR layer 1 ($t_{AR,1}$,$n_{AR,1}$) stacked DeWAL, LDPE adhesion layer (t_{LDPE}), AR layer 2 ($t_{AR,1}$,$n_{AR,1}$) stacked DeWAL, LDPE adhesion layer (t_{LDPE}), core. This design was used for the BA220/270 thin window; stars denote the chosen design parameters. Parameters not shown are optimized to produce the lowest possible band average reflection.

best for roughening bulk PTFE, which we should use in the future.

We optically tested the AR coating for the lenses before laminating them to the lens surface. These measurements were taken with a Vector Network Analyzer with a Virginia Diodes WR6.5 head. The signal is coupled to free space with a WR6 rectangular gain horn. The beam is then collimated with a 90° off axis parabolic mirror, with the beam waist placed where two posts allow us to hold large optical materials up to the flat surface the posts provide. We normalize the measurement with a short measurement of a flat aluminum plate at the measurement surface (see the top of Figure 3.13). Time gating the S11 (reflection) signal (calculated by our HP 8510C Network Analyzer) allows us to remove spurious reflections from the system, such as those from the horn.

The top of Figure 3.13 is a picture of a measurement with this set up in progress of the AR coating for a lens. Figure 3.13 shows the actual VNA measurements of the lens AR coat; at this point, all layers are adhered together, except for the final LDPE adhesion layer to the lens. Therefore, the stack is designed to be: DeWAL ($t_{AR,1}$=0.27, $n_{AR,1}$=1.18), LDPE adhesion (t_{LDPE}=0.05 (2 mil), n_{LDPE}=1.51), and bulk PTFE ($t_{AR,2}$=0.127 (5 mil), $n_{AR,2}$=1.43), as shown in the bottom left of Figure 3.13. We also plotted models with three potential thicknesses of the DeWAL layer (dashed lines), to test the tolerances of our stack. In the bottom right f Figure 3.13 we project the modeled AR layers, with an infinitely thick HDPE core, into the design band. Also, we take averages of the reflected power in band and find that this multi-layer AR coat is expected to produce band average reflections between 0.05% and 0.27%; this is consistent with the predictions from the design corner plots in Figure 3.12.

The design process of the nylon filter AR coats is, as stated before, very similar

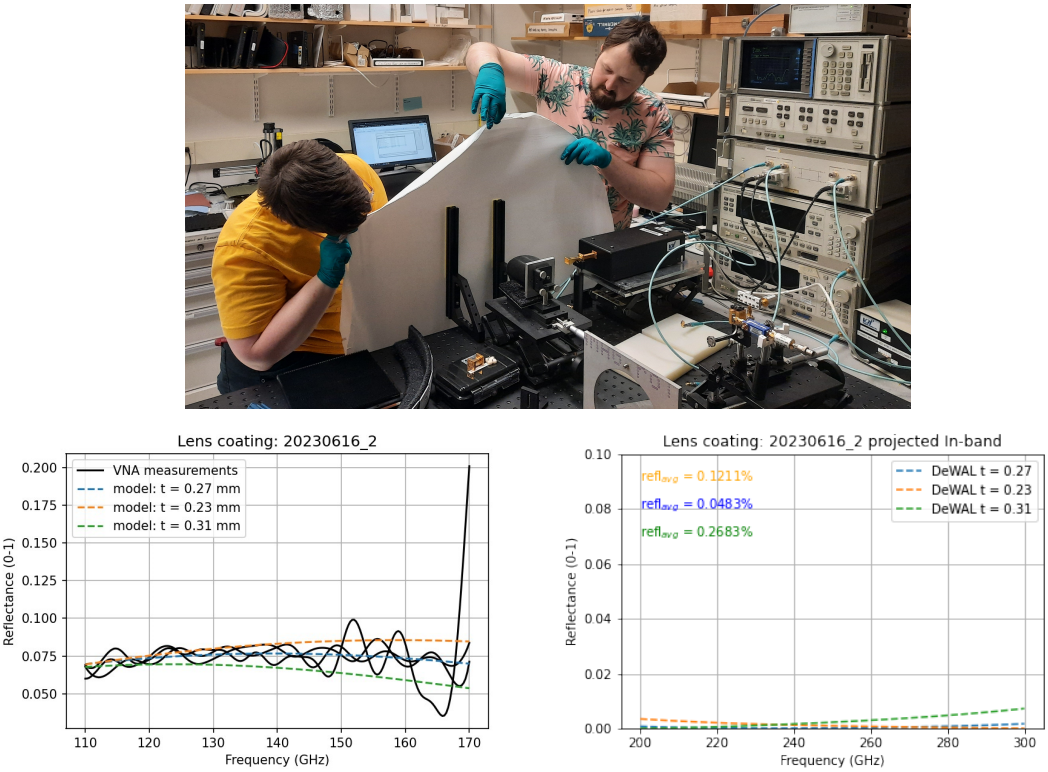


Figure 3.13: [Top] Picture of the out of band measurement of the lens AR coat. The AR coat is held up to the measurement surface by two people (left Miranda Eiben, right Brodi Elwood). Signal from VNA (right most stack of electronics) is sent to a WR6.5 head (center right black box) coupled to free space with a rectangular horn (gold, attached to head). The signal is collimated by a right angle parabolic mirror, the back of which can be seen to the left of the horn. Full scale optics samples are held at the referenced measurement plane with two extruded aluminum posts (center black posts). Photo taken by Matthew Miller. [Bottom Left] Out of band measurements of a BA220/270 lens AR coat taken with the VNA with three possible DeWAL (AR1) layer thicknesses. The three modeled thicknesses for the DeWAL layers are 0.23mm, 0.27mm, and 0.31mm (dashed lines). The other layer parameters are assumed to match the physical measurements taken pre-lamination. [Bottom Right] The same parameters in the models from the left are modeled on an infinitely thick core of HDPE to judge efficacy of AR parameters.

to the HDPE, but with one crucial difference. The ideal middle layer index (n_2 ∼1.50) in a two layer AR coat (see Table D.3) is actually very close to the index of LDPE[9] (n_{LDPE} ∼1.51)(Lamb 1996). As such, we can actually use our traditional hot melt adhesive as an AR layer itself, which we did.

The final design parameters of the BA220/270 nylon AR coat were a DeWAL outer layer 250 μm thick with an index of 1.15, and an inner layer of LDPE 200 μm thick with an index of 1.51. However, the stack of DeWAL after compression was too thick at 380 μm: we removed a layer of DeWAL from each side to make an outer layer 362 μm thick. This improved the performance from the initial band average reflection of 1.3% to 0.9%, which was deemed acceptable. We likely could have taken another layer off to further reduce the in-band reflections.

BA220/270 Thin Window

The initial thin window for BA220/270 used a very similar design of AR coating as the other cold PE optics. In retrospect this was a mistake: the bulk PTFE was too stiff to deform with the window, and during initial pump down the AR coating delaminated. Addionally, the bulk PTFE wrinkled, and caused a large wrinkle in the center of the window; this allowed an unacceptable amount of gas into the cryostat, and prevented it from cooling.

We had to return to the drawing board. At this point, we had acquired a new roll of

[9]Polypropylene is also close to the ideal index (n_{PP} ∼1.50), but we have had difficulties with bi-axially oriented polypropylene (BOPP) contracting atop optics when melted and therefore not working well out of the box as a hot melt adhesive (Shitvov et al. 2022).

DeWAL, which had much better thickness consistency across the width (see Figure **??**). This allowed us to explore using only DeWAL for a multi-layer AR coat, as thinner layers could be made with more consistent thickness across the entire span.

As DeWAL is significantly more flexible, lower density and thickness than bulk PTFE, the inner pseudo-layer could be composed of a very thin layer of DeWAL and LDPE. The exploration of this parameter space is shown in Figure **??**. However, the new roll of DeWAL, probably due to the thinner initial thickness, takes significantly more layers to build up to the thicknesses required for an outer layer.

If we miss the desired specification outer layer thickness, we may be able to compensate with the other AR layers' parameters (for example, increasing the thickness of the inner DeWAL layer and decreasing the LDPE thickness) to still produce a decent band average reflection. We cannot—as we learned with SN#20231001—add more layers to the outer layer.

During the second window creation campaign, we had to iterate on the DeWAL flattening procedure, as the thinner DeWAL roll came with significant wrinkles that (initially) had to be flattened by hand. Once we developed the stretching rig (shown in Figure 3.15 and the procedure is discussed in Appendix C), DeWAL lay-ups became significantly easier, but this initial procedure difficulty led to another production mistake.

When producing the AR layers for SN#20231001, I made the inner layers first, as they required fewer layers of DeWAL, and the lay-up procedure was time-consuming and annoying. We then made the outer layer and completely missed the desired thickness. We attempted to add more layers of DeWAL to the outer layers, but when taking quality assurance deflection and gas entry rate measurements the added DeWAL layers on the

Table 3.5: Design, measured and modeled AR coat parameters for the second and third BA220/270 thin windows. Thicknesses were measured with calipers, modeled thicknesses and indexes are fit parameters to VNA measurements, such as those in the right side of Fig. 3.14. The modeled parameters below were used to plot the models in the left side of Fig. 3.14.

Window	t_c [μm]		t_1 [μm]		n_1		t_{LDPE} [μm]		t_2 [μm]		n_2	
SN#	Des	Mea	Des	Mod	Des	Mod	Des	Mod	Des	Mod	Des	Mod
20231001	1250	1110	280	190/ 250[a]	1.14	1.18	90	90	50	50	1.20	1.20
20240722	1250	1170	263	204	1.13	1.17	76	61	60	34	1.18	1.20

[a]The top (concave) outer AR layer has had a layer peeled off to removed bubbled sections, so there are two different outer thicknesses on this window.

top surface peeled up in large, uneven bubbles. Despite attempts to save the added layers, they would not bond. After exploring peel strengths of various processed layers of DeWAL (discussed in Table 3.4), we have determined that we only have one shot to make a bonded multi-layer stack of DeWAL. Therefore, on the third and final thin BA220/270 window (SN#20240722) we made the outer layers first. When we again undershot the desired outer layer specification, I could then compensate (somewhat) with the inner layer parameters, producing a significantly better band average reflection.

We measured the out-of-band reflection with the WR 6.5 head and set up described previously, but the amplifier in the head was damaged at the time. As such, the measurement at high frequency is not very reliable, and time-gating was not applied to

the data in Figure 3.14. We managed to take some in-band VNA measurements on the center of the deflected window, but these measurements are highly variable thus not depicted here; they were useful to determine that there is a reflection peak in band at approximately 230 GHz, which is depicted as a vertical green line in Figure 3.14. Due to the complexities of these datasets, I fit the corresponding modeled window (dashed blue line) by hand, attempting to match:

1. the frequency of the measured null at approximately 135 GHz,
2. the amplitude of the reflection below 135 GHz (which seemed to be less impacted by the damaged WR 6.5 VNA head),
3. the measured in-band peak at approximately 230 GHz,
4. the combined measured final thickness of the pseudo-inner layer (combined AR_2 layers of DeWAL and LDPE) being approximately 95 μm,
5. and the measured final thickness of the outer DeWAL layer being approximately 204 μm.

The final fit AR parameters are reported in Table 3.5 for the reflection profile shown on the right side of Figure 3.14. The goal design is compared to each of the two windows' reflections (both SN#20231001 and 20240722) in the left side of Figure 3.14.

3.5 Problems with Using ePTFEs

As one may expect from developing new techniques with a new material, we have encountered some difficulties with using ePTFEs (including both Teadit and DeWAL)

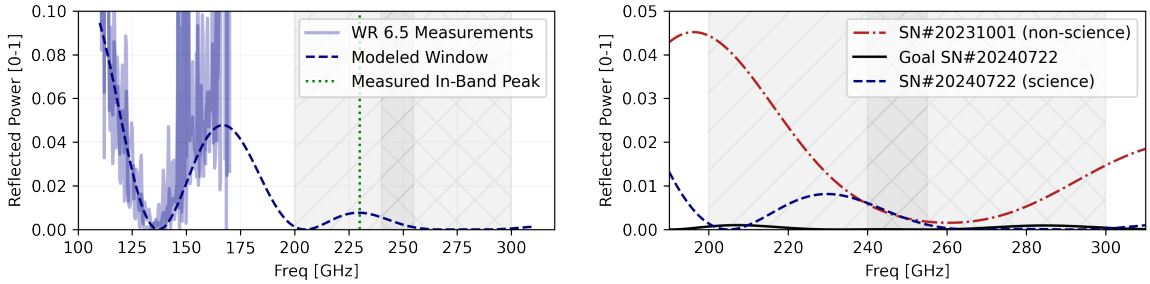


Figure 3.14: [Left] Out of band measured reflections of the science grade BA220/270 window (SN#20240722) with the damaged WR 6.5 VNA head (solid blue). Location of the in-band peak is noted with a dotted green line. Modeled reflection off the window denoted by the dashed blue line. [Right] Models of the goal window design (black solid line), final SN#20231001 (non-science) reflections (red dot-dashed line), and final SN#20240722 (science) reflections (blue dashed line).

for AR coats. I describe a few of the problems we have encountered here, such that the reader may be aware of the potential issues that must be overcome to use these materials. These problems do not detract from the usefulness of such low density, low thickness and wide materials for single and multi-layered AR coats but are meant to advise readers what they may be getting into when using them.

Some of these issues are specific to one form of ePTFE, some to both. I attempt to be clear which issue corresponds to which.

Bonding process (DeWAL)

As shown in Table 3.4, the bonding process between layers seems to only work between two Naive layers; when attempting to add additional layers to an already compressed stack, either Pre-compressed or new Naive layers will not bond as well as two layers of



Figure 3.15: [Left] DeWAL stretching rig, which flattens a sheet of DeWAL before it is smoothed over a flat plate. This process helps prevent wrinkles and bubbles from forming under a sheet, which must be smoothed out by hand. [Right] The AR coating of the Short Keck (SK) window, with the top layer of DeWAL peeled up. The top layer got caught on the release layer and came up when the release layer was removed.

Naive DeWAL will bond to each other. Even 1 atm compressed Naive layers exhibit 10 times stronger bonding than Pre-compressed layers with each other.

We are unsure why the bonding seems processing specific: it is possible that the heated compression changes the surface texture or the receptiveness of the polymer chains to making weak bonds with other chains. However, practically this means that a stacked layer of DeWAL membranes must be made in one compression cycle. If additional electrical length is required to meet reflection requirements, they must be added with hot melt adhesive layers.

Static (Both)

Most ePTFEs are incredibly prone to accumulating surface static. This generally means that, unless handled in clean room environments, the AR layers will attract dust in

large, visible quantities. The static can be mitigated somewhat with ionizing fans, but thankfully small bits of dust and fibers do not seem to adversely affect AR performance. Copious amounts of dust may impact the bonding performance of DeWAL, however, so static should be mitigated as much as possible.

The static does introduce difficulties in handling cut sheets of DeWAL, both big and small. Because of the membrane's incredibly low density, the static can generate enough force to move the DeWAL. If holding a long sheet up by an end and letting it hang down, it will be attracted towards a person's body. If laying small sheets onto another static prone surface, such as the bulk PTFE release layer, the DeWAL may shift, or lift off the surface. It is also difficult to blow an ionizing fan over uncompressed DeWAL due to the air flow potentially lifting up the material.

Therefore, handling long sheets are best accomplished with two people or a rig, such that four corners of the sheet can be secured and prevented from moving (see Figure 3.15 for a picture of our DeWAL stretching rig and Appendix C for instructions on how to lay down DeWAL layers). Additionally, we have found that 92-600 Nitrile texture-less gloves work best when smoothing bubbles and wrinkles by hand. Other types of gloves tend to catch on the DeWAL, causing it to bunch up and wrinkle.

Release layer (DeWAL)

When compressing the DeWAL, or adhering it to another surface, there must be a bulk PTFE release layer above the layers. DeWAL tends to stick to other release layers (such as polyester). Even bulk PTFE can generate enough static force to peel up the top layer of DeWAL; release must be done carefully, in small increments, preferably with an

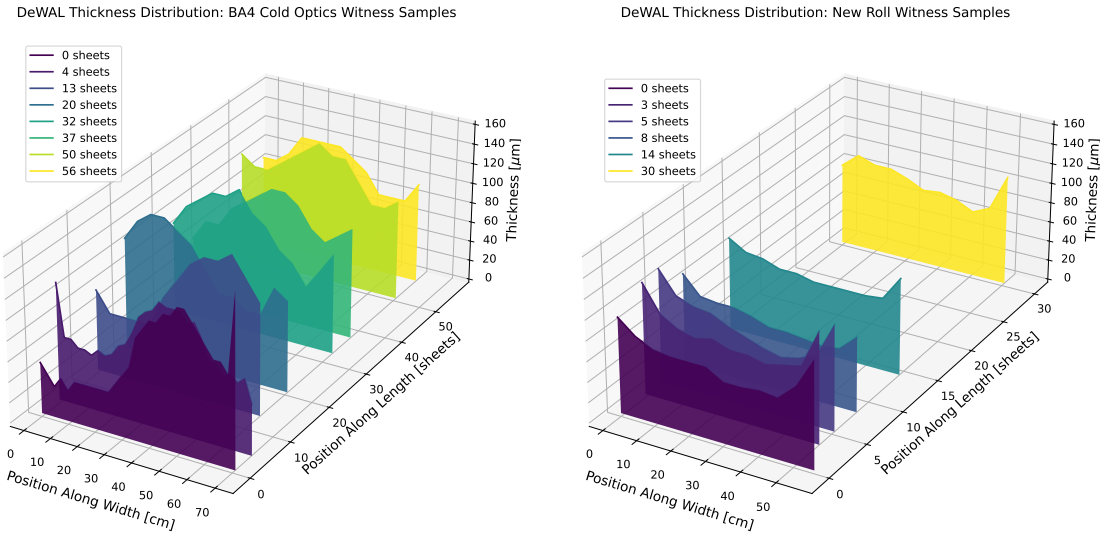


Figure 3.16: [Left] Thickness distribution across width over the roll of DeWAL used for the BA4 cold optics. Length along the roll is measured in 'sheets', which are roughly 75 cm lengths cut off the roll. [Right] Thickness distribution across width over the roll of DeWAL used for the final two thin windows and wSMA window. Length along the roll is measured in 'sheets', which are roughly 75 cm lengths cut off the roll.

ionizing fan running over the peel surface. See Figure 3.15 for an example of an AR coat that had a DeWAL layer peel off with the release layer. I suspect that the layer came up with the release layer because of improper cooling after adhesion to the slab window: the bottom of the slab was still warm, and therefore there was uneven contraction of the release layer, which pulled on the DeWAL layer.

Thickness consistency (Both)

Teadit, due to the mounding behavior shown in Figure 3.6, can have very large thickness and density variation across the sheet. Compression mitigates this effect somewhat, but

the gradient can still be measured. The final compressed thickness is correlated with the initial thickness.

DeWAL comes in rolls at least 200 feet long. We then cut the appropriately scaled 'sheets' off these rolls and lay the layers onto a flat surface with a release layer, before transferring that stack into our optics autoclave[10] (Eiben et al. 2022). We typically rotate the stack between layers of DeWAL, to somewhat evenly distribute any thicknesses profile across the stack. However, this was complicated somewhat by the initial roll having a changing thicknesses profile across the roll.

In the left of Figure 3.16, we show how the thickness varied across the initial roll of DeWAL, which was used for the BA220/270 cold optics. As shown, the 'mound' of the thickest section moved back and forth along the width of the roll. This presented a problem with making the final stacked AR coating a consistent thickness. The latest roll had much better thickness consistency across the roll, but was about half as thick initially as that center mound. The thickness distribution across the new roll is shown on the right Figure 3.16, which was used for the BA220/270 window and wSMA windows (described in § 3.4.2). The production company has been very receptive to attempting new techniques to match desired parameters, so I hope that the next rolls will only improve their usefulness.

[10]Instructions for operating the autoclaves are in Appendix B.

Chapter 4

Thin Window Optics

Portions of this thesis chapter originally appeared in the literature as

BICEP/*Keck* Collaboration et al., *arXiv e-prints, arXiv:2410.12089*

The slab vacuum windows were previously modeled to be the single highest contributor to photon noise in BICEP/*Keck* series of instruments (Barkats et al. 2018; BICEP/Keck Collaboration et al. 2022). However, there was not a lot of direct evidence to support this assumption: on the receivers we had previously taken some instrumental loading measurements, and tested the optical impact of removable external optics on the receiver (namely the forebaffle). Individual optical component characterization was limited to room temperature free space measurements, such as the reflection measurements described in Chapter 3 and scattered power measurements described in Corrigan (2019).

Having now characterized the window mechanically in Chapter 2, and designed the window AR coats (which takes an understanding of both mechanical and optical

properties) in Chapter 3, we now turn our attention to optical characterization of the window. We will quickly cover the HMPE window index of refraction measurements that were taken as a given in Chapter 2 in §4.1.1. Of particular concern is the potential polarization effect of laminates at our frequencies (§4.1.2), as this may impact the polarization response of our sensitive polarimeters. Finally we shall report the in-situ measured change in optical power associated with windows on the BICEP3 telescope three different ways (§4.2).

4.1 Optical Tests and Modeling at Home

4.1.1 Index of Refraction

The index of refraction of the laminate windows is of particular interest for two reasons: first, we must have an expected core index of refraction to impedance match AR coats, and second, the index of a material is dependent on the density or fill factor of the material and therefore the lamination quality of the laminates.

As previously described in Chapter 3, we have a Vector Network Analyzer (VNA) which we can amplify and multiply to millimeter-wave bands with various amplifier heads. The VNA can record the reflectivity off a large sample (when properly calibrated) by coupling the signal to free space with a rectangular horn, then is collimated with a parabolic 90° off-axis mirror. Samples are placed at the beam waist of the resultant beam, and normalized with a 'short' (aluminum plate) assumed to have perfect reflection. We can then fit a reflection model to the resultant reflected power. Given we have a good constraint on the physical thickness from mechanical measurement, the fit is primarily to

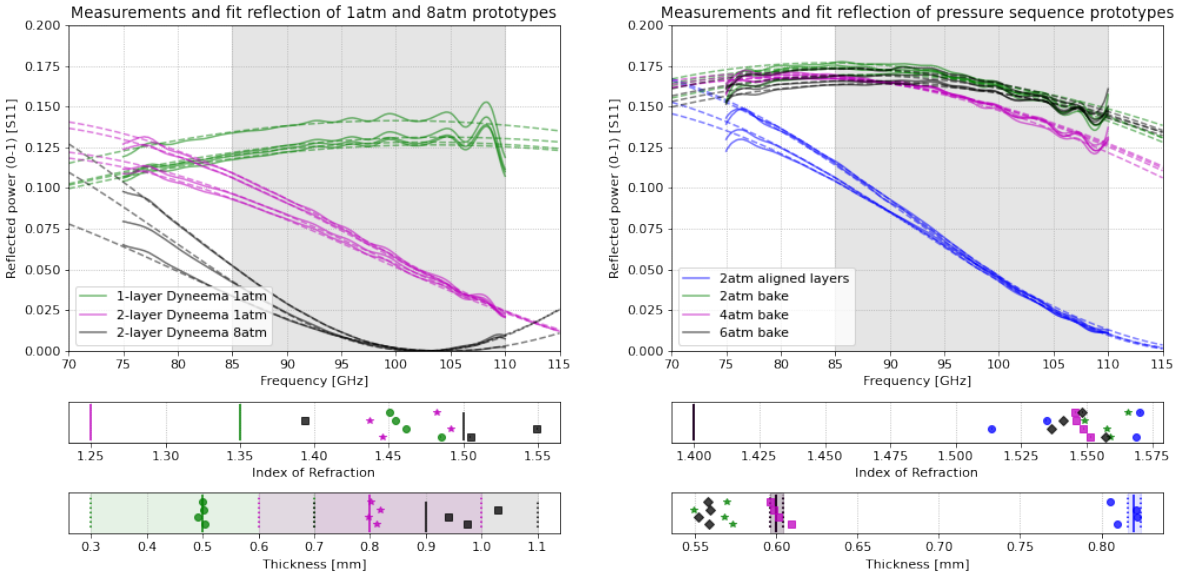


Figure 4.1: Free space VNA WR 10 measurements of HMPE window cores baked at different pressures. [Right] 1 atm in green and magenta, 8 atm in black [Left] 2 atm in blue and green, 4 atm in magenta, and 6 atm in black. Top panel depicts the measured S_{11} reflected power (solid lines) and associated fits (dashed lines). The bottom panels depict the fit parameters for each of the measurements (points) and the initial guessed parameters (solid lines) of index and thickness of the samples.

estimate the index of refraction for either the core or the core and AR coat, depending on the sample.

Recall that the index of a material is related to that material's density or fill factor (the Lorenz-Lorentz relation/Eq. 3.23). For a laminate composite window, the quality of the lamination is also dependent on the fill factor[1] and therefore another method to estimate fill factor is by measuring the index of refraction.

[1]Recall the relation between lamination quality/fill factor and strength in §2.2 and the increased permeation rate associated with poor lamination in §2.5.

In Barkats et al. (2018), the low fitted value of the index of refraction for laminate windows (n≈1.35) was attributed to a low fill factor of 30–40%: this is partially true, but those index measurements were later found to be affected by systematic error. If there is a difference in angle between the short measurement plane and the sample plane, the normalization of reflected power of a sample will be biased down. As the dominant constraint on index[2] from the S_{11} measurements is sourced from the amplitude of the reflected power (recall Eq. 3.2), a systematic reduction in power results in a systematic reduction in fitted refractive indexes.

I remeasured the samples discussed in Barkats et al. (2018) (these are the 1-layer and 2-layer Dyneema 1atm on the left of Figure 4.1) with proper systematics correction, and found much higher indexes . Of greater interest, however, was the estimated fill factor from fitted refractive indexes, and the fitted refractive indexes' dependence on lamination pressure. As such, the fits to various pressure bakes are reported in Figure 4.1. The 1 atm pressure bakes (the same samples that were previously reported on in Barkats et al. (2018)) are found to have lower indexes of refraction (n∼1.45) than the bulk index of polyethylenes, but only corresponding to a fill factor of 85%. The fitted refractive indexes for fully consolidated higher pressure bakes tend to be similar to the bulk indexes of polyethylenes, between 1.536<n<1.557 (Lamb 1996; D'Alessandro et al. 2018; Elwood et al. 2024); the range of indexes is due to repeatability in the S_{11} measurements. The high fill factors found from indexes of laminates baked at high pressure is corroborated with the density measurements, as reported in Table 2.2.

[2]The dominant constraint on the thickness of a sample from S_{11} measurements is peak/null placement in frequency, though those are also degenerate with the fitted refractive index because in actuality the peak/null placement is dependent on electrical length.

4.1.2 Birefringence

The BICEP/*Keck* Collaboration is aiming to measure the polarization of a cosmological signal to very high precision. The B-mode angular power spectra are expected to be multiple orders of magnitude below the temperature anisotropies. Due to the highly aligned structure of the HMPE fibers within the woven fabric at macroscopic scales, there is a legitimate concern that the woven structure we are incorporating into these windows may induce a polarized response, reducing our sensitivity to primordial polarization states. An instrumental polarized effect, such as birefringence, from the window would induce correlations that our normal analysis self-calibrates for and eliminates (Keating et al. 2013; BICEP/Keck Collaboration et al. 2021). Nevertheless, we constrain any potential polarizing effect from the laminate window using both the benchtop measurements described in this section and direct measurements on the BICEP3 receiver in BICEP/Keck Collaboration et al. (2024a).

Recall the Stokes parameters from Equation 1.15; Stokes Q (linear horizontal or vertical polarization), U (linear $\pm$ 45 deg polarization), and V (circular polarization). Given that our orthogonal linear antennas couple only with the Q polarization, our detectors should not see any U signal (BICEP2 and Keck Array Collaborations et al. 2015). This assumes that our linear antennas are perfectly orthogonal, and that our instrument does not rotate the polarization of the input signal. Birefringent optical materials have a "slow" and "fast" axis, which will retard the phase of incident polarized light, potentially transforming a hypothetical pure Q signal into some other combination of these Stokes parameters. Birefringence depends on the difference in index and path length (i.e., electrical length) between the slow and fast axis, and the orientation of the

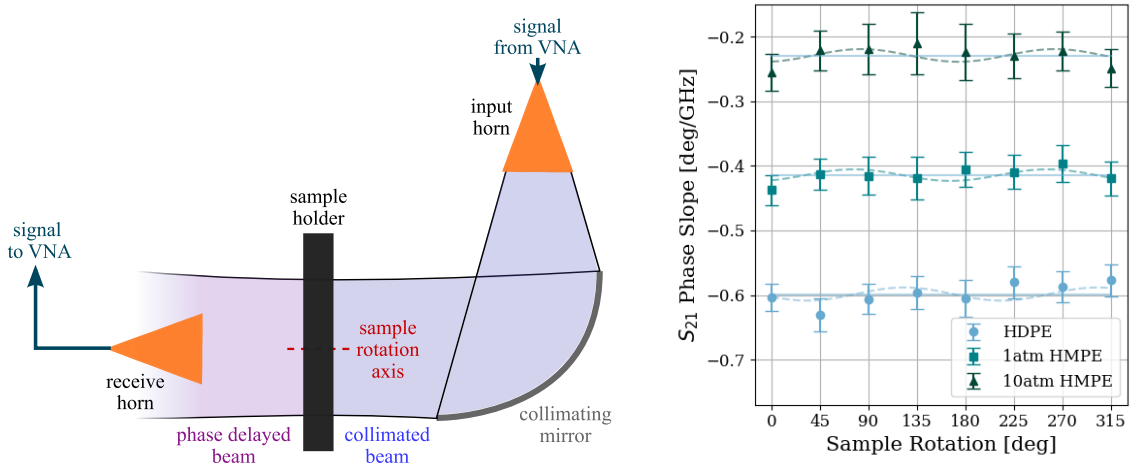


Figure 4.2: [Left] Diagram of phase offset measurement (not to scale). Millimeter wave signal comes from the VNA and is coupled to free space through the input horn and collimated by the parabolic mirror. The collimated beam then interacts with the sample in the sample holder orthogonal to the propagation axis, where the phase will be delayed by the sample's electrical length. The receive only horn then couples the signal back into the VNA (S_{21} phase measurement). [Right] Fitted phase slopes for S_{21} phase offset measurements of 0.79 mm (1/32") thick HDPE and the two HMPE laminates, one baked at 1 atm (0.68 mm thick) and the other baked at 10 atm (0.4 mm thick). If the materials are not birefringent then the phase slope should not vary significantly with sample rotation. If the materials exhibit uniaxial birefringence, we would expect the same phase slope on measurements 180 deg rotated from each other: we have fit both a flat line (solid lines) and a $\sin(2\theta)$ wave (dashed lines) to the data to test the maximum amplitude of the possible phase offsets for each material.

axes to the incident polarized light. There are two potential impacts that a birefringent window could have on our receiver's polarized transfer function:

1. Rotation of polarization, or introducing a U beam.

2. Reduced polarization efficiency, or rotation from Q partially into U and V.

If there is a birefringent effect from the difference in index between the HMPE fibers and the surrounding LDPE, the windows are thin enough that there should not be a long enough path to produce a significant phase offset. However, it is also possible that the HMPE weave induces other polarization effects, such as polarized scatter, that in practice act like a birefringent material. Additionally, as discussed previously in §4.1.1, the lamination quality (fill factor) has an effect on the index of the laminate. The lower pressure laminates have a lower fill factor and therefore lower index; the HMPE weave in lower pressure laminates is even more visible due to the poor lamination (see Figure 2.4). If birefringence or other polarized effects were to occur, it is more likely to be measurable in a low pressure laminate where gaps between the HMPE and LDPE are more prevalent.

However, multi-layer laminates—such as the deployed windows—have three layers of HMPE weave rotated 30 degrees offset from each other, as discussed in Appendix A. Multiple weave orientations may obscure the relation between the weave and the phase offset. To better understand the polarization properties of the laminated fibers, we test single layer HMPE laminates to reduce potentially confounding additional angles.

We can compare between a poorly laminated (1 atm) single layer laminate, a well laminated (10 atm) single layer laminate, and a thin extruded bulk HDPE sample.

Extruded bulk HDPE windows have been observed to exhibit birefringent properties at far IR frequencies, so we use it as a detection upper limit (Paine & Turner 2013). To measure the transmission (S_{21}) phase we used the previously described S_{11} VNA set up with an additional receive-only port using an even harmonic mixer to downconvert the received signal. This setup is shown in the left of Figure 4.2. When we introduce a material to the beam of our VNA, it will introduce a phase offset due to the change in electrical length of the path between the horns that varies as:

$$\phi = \frac{2\pi n h}{\lambda}, \tag{4.1}$$

where ϕ is the phase offset (in radians), n is the index of refraction, h is the thickness of the material, and λ is the wavelength of the incident light. The electrical length is the product of the index and the thickness of the material. Therefore, changes in the index at different orientations will result in a different offset of phase, and a different electrical length.

We centered a small rotating stage between the horns, such that the rotation axis was centered with the beam, then took phase measurements on 45 degree increments. We then fit the slope of the phase offset over frequency (fit slopes for each rotation measurement and material shown in the right Figure 4.2), which is directly related to electrical length. The repeatability of the phase was 3 degrees: this is very similar to the calculated fit uncertainties shown in at the right of Figure 4.2.

If the materials are not birefringent, then we expect no relationship between the sample rotation angle and the fitted phase slope. If the material exhibits uniaxial birefringence (one fast axis and one slow axis orthogonal to each other), then we expect the same change in electrical length on measurements rotated 180 degrees from each

other; in other words, S_{21} phase slopes fitting a sine wave with a fit phase offset (x), amplitude (A), and amplitude offset (y) over the sample rotation (2θ):

$$\phi = Asin(2\theta + x) + y \tag{4.2}$$

I fit both a flat line and a $\sin(2\theta)$ wave to the measurements to test the null hypothesis. The sine wave fit has two more degrees of freedom than the flat line, and therefore the $\delta\chi^2$ between the flat fits and the $\sin(2\theta)$ fits would have to exceed the extra degrees of freedom to be significant, which they do not.

We might expect more than one fast or slow axis induced by the HMPE weave in the laminate. However, these measurements cannot constrain more than a uniaxial birefringent effect. We would also expect that extruded HDPE would act more like a traditional uniaxial birefringent material—as we suspect that the extrusion should induce a preferred orientation of polymer crystals along one axis in the manufacturing process—which should be captured by these measurements (Paine & Turner 2013).

It is worth noting, as an aside, that the birefringence measured in Paine & Turner (2013) was dependent on temperature, and the largest phase shifts occurred at or below 280 K. Therefore the HDPE we measured at room temperature (∼300 K) lab may not have been birefringent. On the other hand, the thermal birefringence detected in the Paine & Turner (2013) HDPE window is associated with an HDPE glass transition temperature; as parts of the polymer chain "freeze" below the transition temperature[3],

[3]Semi-crystalline polymers such as polyethylene have at least two phase transitions associated with either the crystalline (the melting temperature) or the amorphous region (the glass temperature). The polymer is still a solid between the glass transition and the melting point, but it will behave differently; below the glass transition, for instance, a polymer becomes more brittle.

the coefficient of thermal expansion changes and therefore so does the relationship between index of refraction and temperature (Wietzke et al. 2011). One may naturally question if we should be concerned about a similar thermal birefringent effect in thin windows, particularly as the thin window equilibrium temperature was found to be at 260 K (well below room temperature) in Table 2.5, and may vary in different weather conditions. I am not concerned about thermal birefringence in the HMPE laminates for two reasons:

1. The glass transition temperature of LDPE (the majority of the polyethylene in laminate windows) should be lower than the glass transition temperature of HDPE due to the higher amorphous fraction of the molecules, by potentially 10 to 70 degrees (Badr et al. 2000; Wietzke et al. 2011). The HMPE fibers are highly crystalline: if HMPE has a glass transition temperature it should be at a higher temperature than HDPE (Dyneema 2008). It is unlikely that the window's temperature moves through the LDPE or the HMPE glass transition temperature during normal operation, but if it does, I do not think that birefringence will be induced for the following reason:

2. The HDPE thermal birefringence in Paine & Turner (2013) was associated with an extruded HDPE slab (extrusion would induce some alignment of the polymer crystals and therefore a preferred direction of polymer chains), while our LDPE should have no preferred orientation. We fully melt the LDPE to laminate the HMPE fibers. The LDPE polymer orientation within the laminate should be fully randomized, if it even existed in the films in the first place.

As the HDPE was also the thickest of the samples tested, and it would therefore

induce the largest phase difference and the greatest chance of detection. There is no strong evidence for a detection of birefringence in the extruded HDPE sample. As such, we set an upper limit on birefringence in the laminate windows at the maximum level of the fit $\sin(2\theta)$ amplitudes. The maximum fit amplitude in electrical length is approximately 0.008 mm for any of the materials, or approximately a 0.02 change in index for the thinnest laminate.

An upper limit change of 0.02 in the index is supported by the typical repeatability of S_{11} index of refraction fits to laminate windows (fitted refractive indexes range between 1.536 to 1.557, as discussed in §4.1.1). If we calculate the impact of such a birefringent window on a fully polarized incident Q beam, we can put constraints on those potential impacts. We can take the combination of window parameters that will produce the worst phase retardation: the thickest deployed window (1.4 mm on BICEP3), and the highest frequency a deployed window is designed for (170 GHz, the upper edge of the BA150 band) to constrain the worst case scenario. The window (or any birefringent optic) can be modeled with a Mueller matrix ($M_{\rm window}(\phi)$) such as:

$$M_{\rm window}(\phi) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -\cos\phi & -\sin\phi \\ 0 & 0 & \sin\phi & \cos\phi \end{bmatrix} \tag{4.3}$$

$$\phi = 2\pi(n_s - n_f)h/\lambda,$$

where ϕ is set by our worst case parameters, the thickness (h) of 1.4 mm, the fast and slow axis indexes are set have a 0.02 difference ($n_f = 1.536$ and $n_s = 1.556$), and λ is the wavelength of interest. This matrix will apply a rotation depending on the eigenstate (orientation) of the birefringent window relative to the input polarization. We

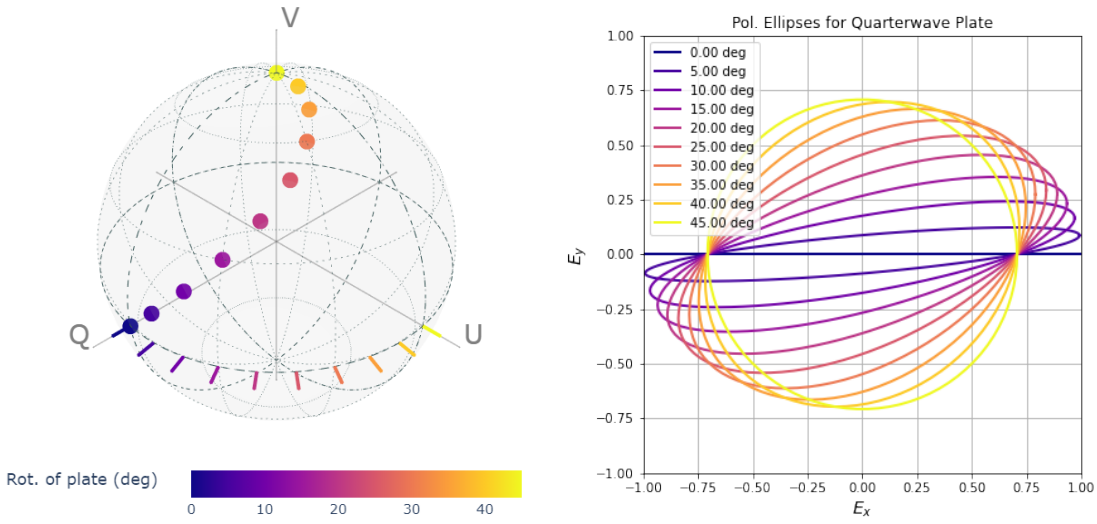


Figure 4.3: Example Poinecaré sphere [left] and polarization ellipses [right] of a quarter-wave plate oriented from 0 to 45 degrees. As the rotation of the plate increases (lighter colors) the incident linear Q polarized wave is further and further transformed into a circular V polarized wave and the orientation of the linear polarization is mixed between Q and U until it is fully transformed to V at a plate rotation of 45 degrees. The plate rotation of 22.5 degrees transforms the incident Q wave furthest into U.

then multiply the matrix by an input Stokes vector (for our purposes, purely +Q) to generate an output Stokes vector of polarization states transformed by the hypothetically birefringent window.

An extreme example of this transformation for a quarter wave plate ($\phi = \pi/2$) is shown in Figure 4.3[4], with two useful polarization state visualizations: a Poinecaré sphere and polarization ellipses. Poinecaré spheres are the polarization Stokes parameters plotted in Cartesian coordinates; they are useful for polarization transformation

[4]Both plots are generated using the open source python package py_pol: `https://py-pol.readthedocs.io/en/master/index.html`

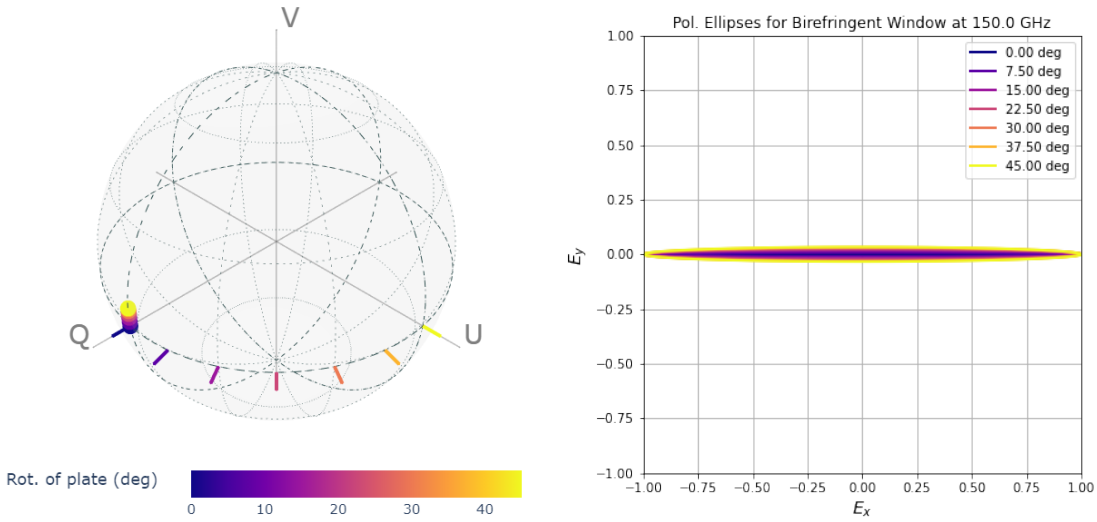


Figure 4.4: Poinecaré sphere [left] and polarization ellipses [right] of a maximally birefringent ($n_f = 1.536$ and $n_s = 1.556$, h =1.4mm) thin window oriented from 0 to 45 degrees at 150 GHz. As the rotation of the plate increases (lighter colors) the incident linear Q polarized wave is further and further transformed into a circular V polarized wave, until it is fully transformed at a plate rotation of 45 degrees. The plate rotation of 22.5 degrees transforms the incident Q wave furthest into U.

visualization as the transformation of the Stokes vector through an optic will rotate on the sphere over the relative orientation axis and by the phase retardation. So for a quarter wave plate (a $\pi/2$ phase retarder) with incident Q polarization, as the rotation of the plate gets further away from the Q direction (lighter colors), the $\pi/2$ phase rotation on the sphere moves further up the sphere.

Polarization ellipses also help visualize the polarization transformation through an optic, but one more grounded in the actual E field orientation. Perfectly linearly polarized light will appear as a flat line; Q is aligned with the x or y axis, while U is at 45 degrees. Perfectly circularly polarized radiation (pure V) traces a perfect circle. The

ideal quarter wave plate in Figure 4.3 fully transforms the incident $+Q$ radiation into $+V$ when orientated at 45 degrees.

The Poinecaré sphere and polarization ellipses for the maximally birefringent thin window at 150 GHz is shown in Figure 4.4. The window acts essentially like a phase retarder with a small phase offset ($\phi_{\rm window} \ll \pi/2$), with a similar 'figure eight' shape traced on the Poinecaré sphere from rotations at all 180 degrees of possible orientations as the quarter wave plate—though on a much smaller scale, due to the smaller phase offset through the window[5].

There are two eigenstates (orientations of the window relative to the linear polarization) that produce the worst rotation for our purposes. At 22.5 degrees the window will rotate the incident Q beam the furthest into U, and at 45 degrees the Q beam will be the furthest rotated into V. The calculated rotations of a hypothetical birefringent window are shown in Table 4.1 at these two worst orientations. It is highly unlikely that the windows are oriented at the worst eigenstates, but these calculations can be taken as the upper limits. The maximum power possibly transformed into a U beam is 0.2% at 170 GHz, significantly less than the previously constrained BICEP2 and Keck Array U beams with powers of 0.8% in the same band (BICEP2 and Keck Array Collaborations et al. 2015). At higher frequencies the transformation into U is worse, but at about level of the previously measured U beams. This effect, however, should be a uniform rotation of the monopole, which would be taken out during the beam fitting process.

[5]In Figure 4.3 and 4.4 only 45 degrees of possible orientations are shown: the next 45 degrees would trace the inverse shape on the other half of the northern hemisphere on the sphere, and the next 90 degrees would mirror the northern hemisphere trace to the southern hemisphere.

Table 4.1: Stokes parameters and polarization efficiency (ϵ) through a worst case laminate parameters are a combination of the thickest deployed window (1.4 mm) oriented at the worst eigenstates (22.5 and 45 deg), and calculated at the edges of the BA150 and BA220/270 bands (170 and 300 GHz, respectively).

	170 GHz		300 GHz	
Stokes	22.5 deg	45 deg	22.5 deg	45 deg
Q (S1)	0.997	0.994	0.991	0.982
U (S2)	0.003	0.000	0.009	0.000
V (S3)	0.076	0.107	0.133	0.188
ϵ	0.003	0.006	0.009	0.018

The maximum reduction of polarization efficiency (ϵ) is estimated by the difference of the full intensity (combined polarization of the polarized Stokes parameters) to the out Q parameter:

$$\epsilon = \sqrt{Q^2 + U^2 + V^2} - Q = 1 - Q. \tag{4.4}$$

Using Eq. 4.4, we can see the polarization efficiency reduction remains very low at 170 GHz (0.3%), roughly half of the previously constrained BICEP3 polarization efficiency at 95 GHz (BICEP/Keck Collaboration et al. 2022). Fully characterizing the polarization efficiency of a receiver requires a full end-to-end polarization measurement, which is very difficult to accurately conduct (Cornelison et al. 2022; BICEP/Keck Collaboration et al. 2024b). We ran another absolute polarization calibration on BICEP3 with a thin laminate window in the 2024/2025 summer season, allowing the change in polarization angle of detectors to be directly compared. Work on analyzing the data

taken is ongoing.

Though a hypothetically birefringent window results in negligibly uniform rotations of polarized light to first order, a higher order change in the polarized transfer function of the telescope would not be well characterized by these estimates. In BICEP/Keck Collaboration et al. (2024a) we explored the effects of how the thin window changed the residual beams in BICEP3, and therefore the estimated temperature to polarization leakage, and found no significant increase in temperature to polarization leakage.

4.2 Optical Characterization on BICEP3

In the austral summer of 2022–2023 we deployed two science grade laminate windows to the South Pole. The BA150 thin window is the primary science-grade window deployed on the receiver, but on BICEP3 the thin window was replacing the previous bulk UHMWPE window: this opened up opportunities to run comparative studies on window optical performance. We could both measure the window with BICEP3 before it was installed, and compare receiver states before and after the window was replaced.

The anticipated decrease in instrumental optical load from decreasing the window thickness is a primary driver for developing thin laminate windows, as shown in Figure 2.1. We can test the difference in optical load with our receivers themselves, as they are highly sensitive and offer the most direct measure of receiver response. We call such measurements "radiometric," as they directly measure the change in optical load with the receiver.

We will discuss the window optical characterization on BICEP3 in this section and

the next chapter. As the legacy instrument, it provides a bevy of data with different well-documented instrument states to compare (BICEP/Keck Collaboration et al. 2022). BA150 has only ever operated with a thin window and has struggled in its first two seasons with excess optical noise we believe to be associated with a blue leak; both of these facts make it difficult to disentangle the window's impact on optical performance. BICEP3, on the other hand, has now had three different windows in the past eight years of operation: an HDPE slab window (31.8 mm thick) for 2016 through 2018, an UHMWPE window (24.6 mm thick) for 2019 through 2022, and an HMPE window (1.4 mm thick) for 2023 through to the present.

We shall explore three different measures of the optical loading on our detectors associated with the window: first, relative insertion radiometric loading measurements of windows on the exterior of the receiver; second, comparing between two total instrumental optical load measurements, one with a slab window on the receiver and one with a laminate; third, comparing the regular CMB campaign calibration optical load measurements between seasons.

We have explored some other thin window related effects on BICEP3, particularly on beam shape and residuals (mainly the T-to-P leakage study reported in BICEP/Keck Collaboration et al. (2024a), and the reduction in beamwidth reported in Giannakopoulos et al. (2024)). I did not conduct the main portions of those studies, however, and will not discuss them further.

4.2.1 Measured Material Optical Load

As we deployed the thin window to a receiver with a window already installed, we could take external relative loading measurements with that window prior to installation. We could also compare the thin window load to the spare 31.8 mm UHMWPE slab. These are inserted material optical load measurements.

Our detectors are transition edge sensors (TES), which voltage-bias a superconducting element—either titanium or aluminum depending on the measurement—on the superconducting transition edge. We measure optical power variation by compensating with electrical power. A given detector will have a designed total saturation power, and changes to the applied optical power will change the margin to that saturation power (BICEP/Keck Collaboration et al. 2022). To gain a direct measure of the power on a detector we can vary the applied bias voltage out of the transition, then step down in voltage until the detector is superconducting. These measurements are known as partial load curves (PLCs); partial because we only explore a range around the predetermined bias voltage. The detector saturation power is:

$$P_{\rm sat}(T_{\rm bath}) = G_c T_{\rm bath} \frac{(T_{\rm oper}/T_{\rm bath})^{\beta+1} - 1}{\beta + 1} \tag{4.5}$$

The module average values from the instrument paper put the $G_c = 40$pW/K, $T_{\rm oper} = 500$ mK, $T_{\rm bath} = 280$ mK, and $\beta \sim 2$. Note that the detector saturation power is heavily dependent on the bath temperature of the focal plane.

PLCs allow us to quickly probe the difference in optical power between two instrument states, such as adding and removing an optical component in front of the receiver, assuming that there is no other change to the applied optical or electrical power. If one is careful in taking a variety of reference measurements (such states with and

without the absorptive forebaffles), one can attempt to determine not only the optical loading from transmission loss in the material, but also any additional optical loading from scattered power associated with the material.

First, we take reference measurements of the receiver without a material (R). We subsequently add the absorptive forebaffles to the receiver and perform the same measurement (R_{fb}). Each of these two measurements are repeated with a material placed on top of the receiver (M, M_{fb}). The excess load from the material itself and the excess scattered power associated with the added material can therefore be derived by $(M - R)$ and $[(M_{fb} - M) - (R_{fb} - R)]$ respectively. Histograms of measured powers per detector for the BICEP3 HMPE window and the UHMWPE spare slab window are reported in Figure 4.5.

As expected, the thin 1.4 mm HMPE window causes significantly less optical loading on the detectors compared to the 31.8 mm (1.25") UHMWPE window. However, the magnitude of the difference is not as large as expected. If the loss tangent (Eq. 5.5) is the same in both polyethylene windows the difference in thickness between windows would result in a factor of 10 difference in optical power. With the radiometric measurements we only measure a factor of 6.6 difference in optical loading. It is likely that the two polyethylenes have different loss tangents, or effectively different window system (combined window and AR coat) loss tangents (discussed further in §5.3). This is supported by the differences in noise equivalent temperature (NET) between years, as discussed in §5.2.

As previously described in Chapter 2, the forebaffle (shown in Figure 1.9 for BICEp3 and Figure 1.11 for BICEP Array) terminates scattered light exiting the receiver on an

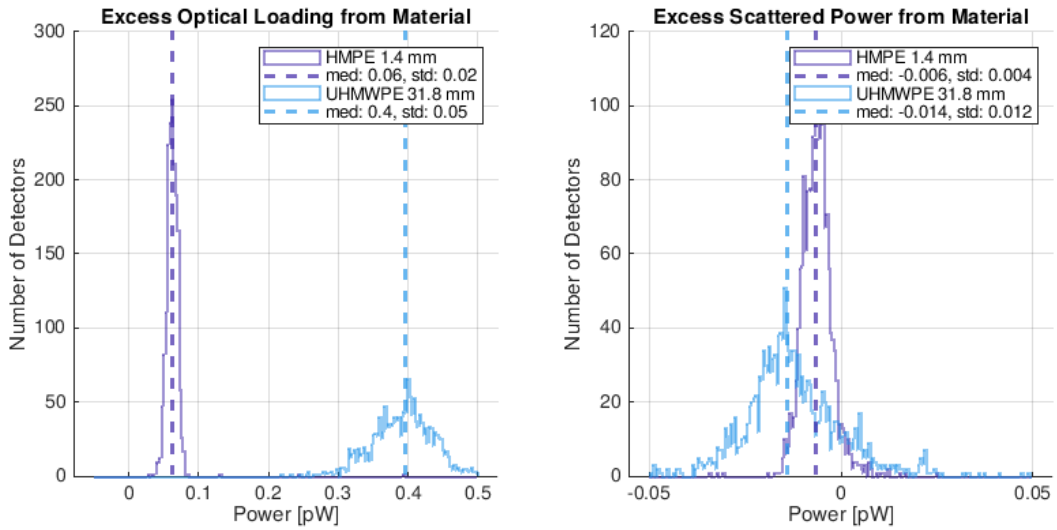


Figure 4.5: [Left] Histograms of excess loading per detector from two different windows placed above BICEP3. [Right] Histograms of scattered power to the forebaffle associated with placing the material in front of the BICEP3 receiver. Blue solid line is a spare UHMWPE 31.8 mm (1.25") thick window for BICEP3; dashed line denotes the median excess load from the slab. Purple solid line is the replacement HMPE 1.4 mm thick window for BICEP3; dashed line is the median excess load from the laminate.

absorber of known temperature. By taking PLC measurements of the materials with and without the forebaffles on the receiver, we can directly measure the excess in optical load associated with the scattered light through these windows. Shown in the right side of Figure 4.5 the measured scattered power associated with the window materials is very small, and for both materials within two standard deviations of zero. These measurements put upper limits on the scattered fraction from these materials, at roughly 4% for the UHMWPE and 10% for the HMPE laminate. The amplitude of these upper limits on scattered power from window material have been corroborated by benchtop scatterometer measurements reported in Corrigan (2019).

These radiometric measurements provide a powerful cross check of the excess load from warm optics in front of the receiver. We can use this information as a supplement for optical modeling, and to provide estimates of optical material properties, such as in Figures 2.1 and 5.4. Independent measurements of material optical load are particularly useful for estimating the receiver's photon load, which is the dominant contributor to the total per-detector white noise, discussed in the next chapter.

4.2.2 Change in Instrument Optical Load

The expected change in instrumental optical load by replacing the thick windows with thin windows was the primary driver of the thin window project. The reduction in optical load has a direct relation with the photon noise on the detectors, and thus our sensitivity. As such, we sought to characterize the change in instrumental optical load and optical noise. The next chapter will discuss the measured change in optical noise on BICEP3, but first we shall characterize the change in optical power.

Elevation vs. Loading

When we take a partial load curve at any given elevation, the measured saturation power (which is the margin to the absolute detector saturation power) depends on both the power coming from the sky (mainly the atmosphere which varies with elevation, $P_{\mathrm{sky}}(el)$) and a constant load associated with the instrument (P_{const}). This is shown as a cartoon in Figure 4.6, and numerically like:

$$P_{\mathrm{meas}}(el) = P_{\mathrm{const}} - P_{\mathrm{sky}}(el) \tag{4.6}$$

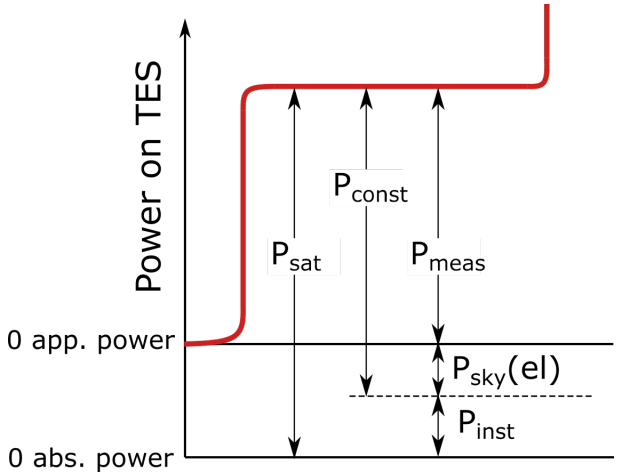


Figure 4.6: Cartoon representation of various named powers on a light TES detector. $P_{\rm sat}$ is the absolute saturation power of the detector which is dependent on detector properties and bath temperature (Eq. 4.5), $P_{\rm meas}$ is the power measured by PLCs (Eq. 4.6 and 4.9), $P_{\rm sky}(el)$ is the optical power from the atmosphere (Eq. 4.7), $P_{\rm inst}$ is the optical power from the instrument, and $P_{\rm const}$ is the difference between the instrument load and the detector saturation power which should remain constant (Eq. 4.8). Note that there is an offset between 0 applied power and 0 absolute power which is dependent on the optical load on the detectors from the sky ($P_{\rm sky}(el)$) and the instrument ($P_{\rm inst}$).

We should like to disentangle these two sources, as the loading from the atmosphere will change over time as the zenith temperature changes, making measured saturation powers taken at different times difficult to compare between directly.

However, on a clear day the sky loading should be directly related to airmass—and therefore the elevation—with a simple function:

$$P_{\rm sky}(el) = \frac{P_{\rm sky}(\text{zenith})}{\sin(el)} \tag{4.7}$$

While $P_{\rm const}$ is the saturation power remaining from $P_{\rm sat}(T_{\rm bath})$ (the absolute total

saturation power without instrument loading, determined by detector parameters and bath temperature, Eq. 4.5) if we account for the instrument loading $P_{\rm inst}$ that goes like:

$$P_{\rm const} = P_{\rm sat}(T_{\rm bath}) - P_{\rm inst} \tag{4.8}$$

Deriving the absolute saturation power ($P_{\rm sat}$) is a bit of a fraught process however, requiring separate analysis of actual dark runs, where no power optical power is deposited on the detectors; if we can assume that the absolute saturation power has not changed (which should be correct, as we have not changed any detector modules on BICEP3 since January of 2017, and the bath temperatures should be held at the same value), then any change in $P_{\rm const}$ *should* be due to a change in $P_{\rm inst}$. To account for detector parameter variation, we directly compare detector to detector. Therefore, if we take a series of PLCs at different elevations we should be able to model the loading from the sky, and only have two parameters to fit to per detector, with a simple equation:

$$P_{\rm meas}(el) = P_{\rm const} - \frac{P_{\rm sky}(\text{zenith})}{\sin(el)} \tag{4.9}$$

However, remember, the $P_{\rm meas}(el)$ is not just dependent on sky loading, or the detector parameters or instrument loading: it's also dependent on the temperature of the focal plane (Eq. 4.5, 4.8). If the focal plane temperature varies during the measurement, the measured saturation power will also vary; this occurred during the 2023 elevation vs. loading measurement described below.

The elevation vs. loading measurements took two series of PLCs at set elevations; first up from 45 degrees to 85 degrees, then back down to 45 degrees. The focal plane temperatures during the 2023 measurement stabilized during the initial upswing, so I only fit to the downswing in that year (see Figure 4.7).

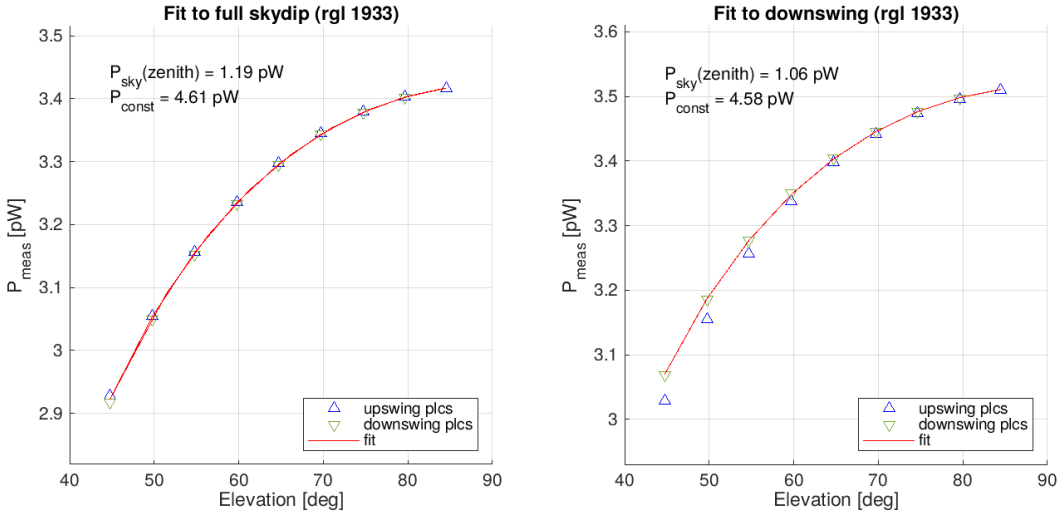


Figure 4.7: Example per detector fits to the PLCs taken at set elevation intervals for the 2017 (left) and 2023 (right) elevation vs loading schedules. Note that in 2023 we only fit to the downswing data, due to focal plane temperature instability during the upswing.

We fit a $P_{\rm const}$ and a $P_{\rm sky}$ to the set of $P_{\rm meas}(el)$ of every 'really good light[6]' detector: then we can directly compare the fit values per detector between the two measurements, one in 2017 with an HDPE slab window, and one in 2023 with a thin laminate window. If we have a change in either $P_{\rm const}$ or $P_{\rm sky}$, we expect a coherent offset in those powers per detector.

The 2017 data were taken in January (when the sun was up at the South Pole) while the 2023 data were taken in May (after sundown). We do see a coherent reduction in $P_{\rm sky}$ in 2023 in the right side of Figure 4.8; this is likely due to atmospheric zenith temperature being lower in May 2023 compared to January 2017. We also see a slight coherent reduction of $P_{\rm const}$ from 2017 to 2023, as expected with the thin window.

[6]'Really good light' (rgl) detectors are detector pairs determined to be well behaved in beam shape/pointing and SQUID tuning.

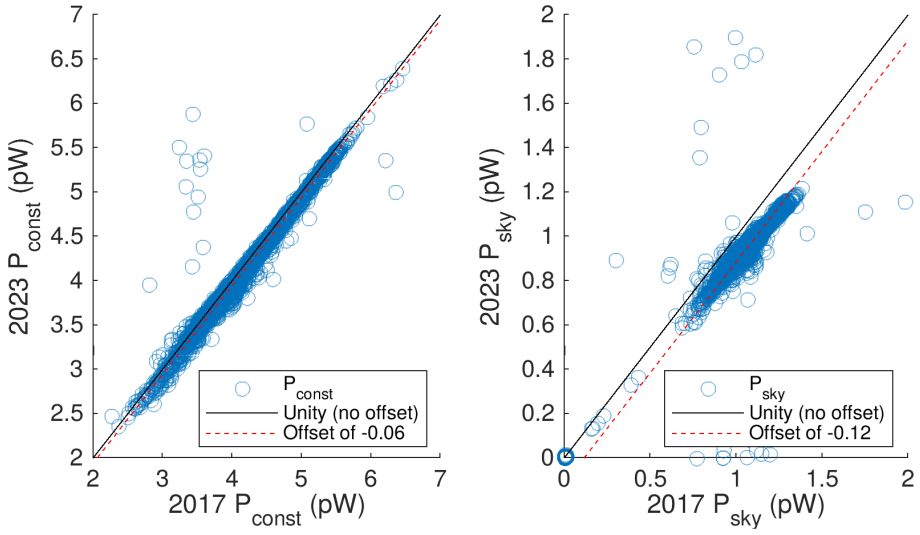


Figure 4.8: Comparison of the fit per detector $P_{\rm const}$ and $P_{\rm sky}$ between 2017 and 2023. No change between years (unity) is represented as a black solid line, while the median linear offset is plotted as a red dashed line. The median offset (2017 $P_{\rm const}$-2023 $P_{\rm const}$) is 0.06 $\pm$ 0.48 pW, while the median offset (2017 $P_{\rm sky}$-2023 $P_{\rm sky}$) 0.12 $\pm$ 0.14 pW.

However, it is not as large as one may expect from the radiometric measurements. There is the possibility that these data were not analyzed well, or there were other complexities at play. The next section explores a separate measure of the change in $P_{\rm const}$, which uses regular calibration data instead of a dedicated schedule.

Yearly Calibration Data

During our regular CMB observations in BICEP3, we take two PLCs per scanset (approximately an hour of time ordered data), just before and after the elevation nods that bracket the azimuth scans. We used to use these PLCs to do a relative gain

calibration[7], but they could also be used to probe $P_{\rm const}$ changes, which is of particular interest for years when a elevation vs loading schedule was not done.

As the regular schedule PLCs are always taken at the same elevation they should always be looking through (approximately) the same airmass[8]: however, weather will change the $P_{\rm sky}$ between PLCs. Over the course of an observing season we might expect that the sky temperature will vary about the some mean value. If we take a median over a season's worth of per scanset PLCs the differences between different weather conditions should be averaged out, and then we can directly compare the per detector $P_{\rm const}$ changes between seasons.

Dark detectors are TESes on the modules that are not connected to antenna, and therefore should not pick up optical loading. There are 80 in BICEP3's focal plane, though about half are currently not reading out. If, say, accumulated gas could impact the conductance of all of the detectors in focal plane, then the dark detectors would also be impacted, but crucially, *should* not see the changed optical loading. However, dark detectors can be influenced by direct island coupling, where power is directly absorbed by the TES island[9]. Direct island coupling would also affect light TESes, so a change in

[7]Recently we have discussed reducing the number of PLCs taken during regular observations to increase the number of scansets possible during a CMB schedule; the interested reader should check the current CMB schedule scripts for PLC cadence if they wish to conduct a similar loading study on future data.

[8]Different receiver rotation (deck) angles will vary the actual airmass that a detector would be looking through during any given regular calibration PLC. Given that I am about to propose take a median over the entire season to (hopefully) average over weather this should also average over the slight differences in airmass between different deck angles.

[9]The TES islands are all exposed to incident radiation, as they and every detector component (the

dark detector loading would also be seen in the light detector loading.

Light detectors are connected to antenna, and should be more strongly influenced by optical power changes. We could attempt to correct for sky loading using the regular elevation nods to then recover an estimate of instrumental loading from the light detector calibration PLCs: further research could be attempted with the yearly BICEP3 PLC data for 2017 through 2023.

I found the P_{meas} for all detectors over all tags in all relevant observing seasons (2017-2023). I then took the median P_{meas} per detector for the entire observing season[10], and compared the median P_{meas} of the dark and light between years. If the instrumental optical load (or the average sky optical load) changed between years, we should see a coherent shift in the measured margin to all detectors' saturation power.

To get a representative value for the offset, I take a subtraction per detector between seasons, then take the median of that difference. The median offset value is plotted on all per detector scatter plots (Figures 4.9 and 4.10) as a red dashed line. In general, the linear offset fits the main population of detectors well, which indicates that the shifts in power between seasons are a coherent effect. We will discuss the optical changes between years as this median offset value.

---

island, the microstrips, the antenna, etc) are machined onto a flat chip. Other radiation coupling techniques, such as feed horns, leave the islands less exposed and hypothetically less likely to directly couple with incident radiation.

[10]We generate an observing season 'taglist' which incorporates all scansets that pass cut parameters for an observing season. Observing seasons typically last from early March to late October, but start/end dates can vary for a variety of factors.

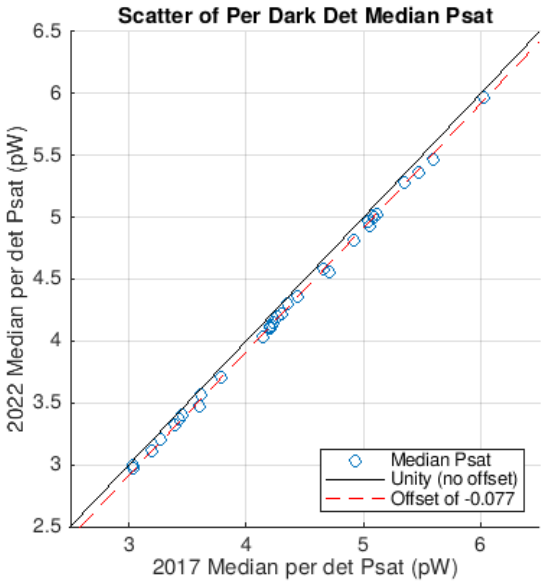


Figure 4.9: Comparison of median yearly per dark detector measured P_{sat} between 2017 and 2022. 2022 has the only significant difference in median yearly dark detector measured P_{sat}. The median yearly dark detector offset between 2022 and the other years ranges between ∼0.058-0.087 pW, and is over 2 times higher than the next largest median yearly dark detector offset (2021-2023, ∼0.024 pW)

Most of the offsets between the median per season dark P_{meas} are small ($\geq$0.025 pW), which provides some evidence that the focal plane and direct island coupling are relatively stable between years. The only year with a consistently large offset in median dark P_{meas} compared to other years is 2022 (shown compared to the baseline season (2017) in Figure 4.9), which may indicate that the condensed gas in the receiver and subsequent increased optical temperatures either increased the out of band power enough to change the direct island coupled power on the TESes, or the gas had some other, subtler effect.

The measured saturation power offset in light detectors between years is likely close

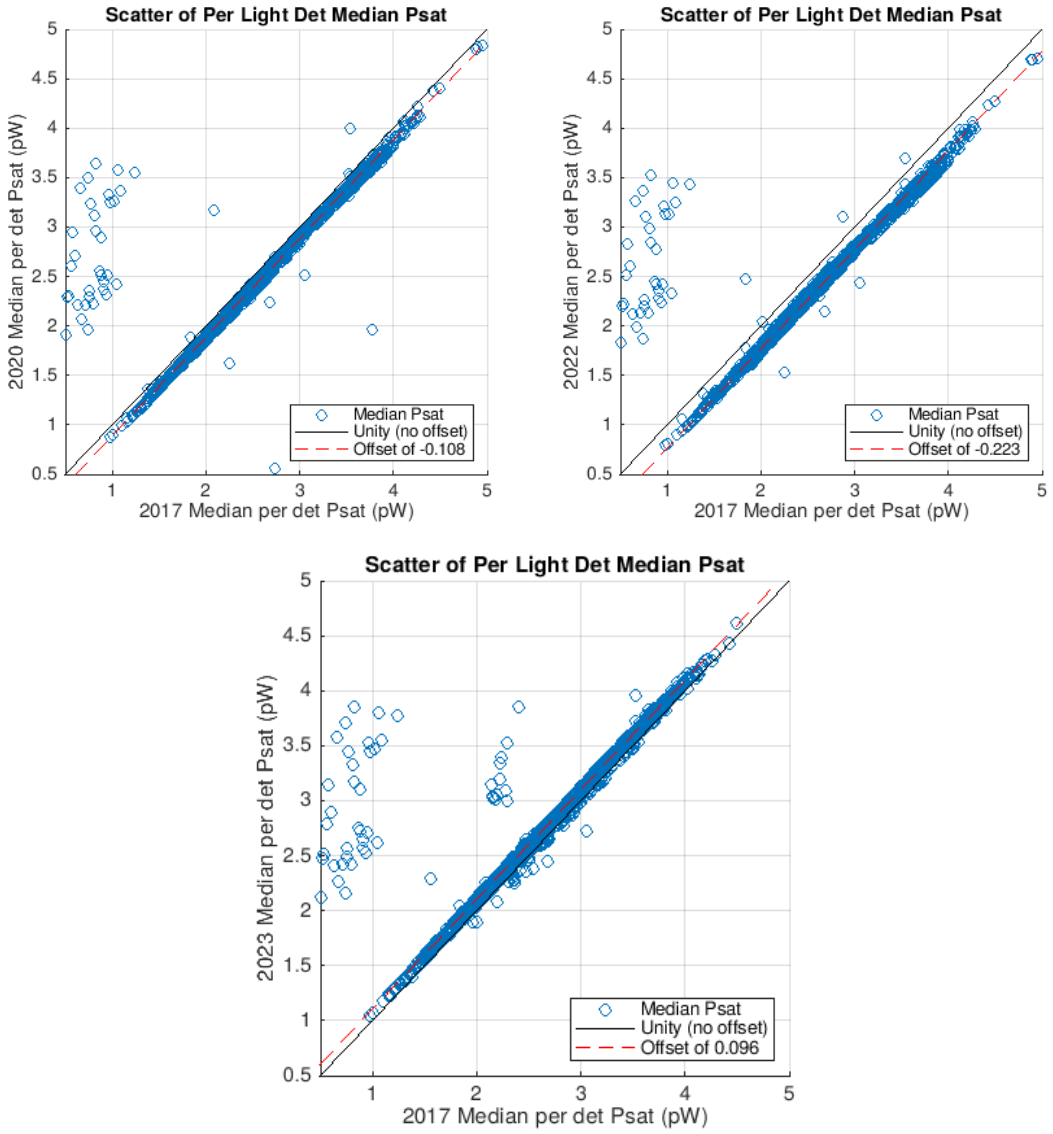


Figure 4.10: Comparison of median yearly per light detector P_{meas} between 2017 and three other representative seasons. For illustrative purposes, a black solid unity line is shown, and a red dashed line shows the median difference in per detector loading. In 2017 (the baseline comparison) BICEP3 had an HDPE window, in 2020 [top left] BICEP3 had a UHMWPE window, in 2022 [top right] BICEP3 had elevated optics temperatures, and in 2023 [bottom] BICEP3 got a thin window.

to the actual change in P_{const}, and therefore the change in the instrumental load. In Figure 4.10, I show scatter plot comparisons of median per light detector offsets for three different representative seasons: for the baseline season (2017) BICEP3 had an HDPE window 31.8 mm thick, in 2020 the receiver had a thinner slab UHMWPE window 24.5 mm thick, in 2022 the receiver optics temperatures were elevated due to gas entry (as previously described in §2.5), and in 2023 the receiver got a thin HMPE window 1.4 mm thick. Comparing years with no expected change in instrumental loading (as the optical load from the instrument should remain stable) in Table 4.2—such as the first entry comparing 2018 to the baseline 2017 or 2019 through 2021—we do not see a large coherent offset in the light detector median P_{meas}. A negative offset in light detector P_{meas} means that there was an increase in load (the later year had lower P_{meas}, meaning there was less margin to the absolute P_{sat} because there was increased P_{inst} or P_{sky}). There is small variability between stable instrument configurations: the offset between 2017 to 2018 is -0.033 pW, and the offset between 2019 and 2020 is -0.035 pW. This may be due to changes in instrumental load or changes in the average sky loading between observing seasons: some years (such as 2020) had extended seasons that lasted from February to late November, which may bias the average sky loading for that season up from other seasons that only average over deep winter sky temperatures.

The very general assumption that the mean P_{sky} does not vary significantly between seasons may not work perfectly, but does allow us have a quick independent cross check our other measurements of optical noise and optical load. Every season after 2017 has a larger median optical load than 2017 (a negative offset in Table 4.2), except for 2023. This is a good indication that the positive offset of P_{const} from the elevation vs load measurements in the previous section is correct. Both of the measured change in

Table 4.2: Relative offset in loading compared to 2017 of per season median per light detector P_{meas}. For every year except 2023 the offset is negative, indicating that most years have a higher median optical load than what was applied to the detectors in 2017.

	2018	2019/2020/2021	2022	2023
Offset Median Light P_{meas} (pW)	-0.033	-0.073, -0.108, -0.101	-0.223	0.096

instrument load are not as large as one may naively expect from the difference in relative material loading reported in §4.2.1: the material loading estimates put the inserted optical load differences between a thick slab UHMWPE window and a laminate window to be ∼0.28 pW, while the actual measured difference in instrumental load between an equivalently thick slab HDPE window and a laminate window is ∼0.06 to 0.096 pW, with very large uncertainties.

In the next chapter we will explore how the instrumental optical load changes between years impacted the measured noise, and how both the radiometric measurements described here and the measured noise may be explained with an instrument optical model.

Chapter 5

Instrumental Noise in BICEP3

Portions of this thesis chapter originally appeared in the literature as

BICEP/*Keck* Collaboration et al., *arXiv e-prints, arXiv:2410.12089*

Our ability to measure a diffuse background signal is limited by our sensitivity. A higher sensitivity measurement requires a higher signal to noise ratio; for our background limited measurements of a very small signal, any change in instantaneous noise has a high impact on the instantaneous sensitivity of the measurement. The lower the instantaneous signal to noise ratio, the longer we must observe the sky to average down the noise. Therefore, the noise on our detectors directly corresponds to mapping speed.

Both the thin window and anti-reflection projects are meant to increase our sensitivity in different ways: the AR coats reduce power lost to reflections off internal optics, while the thin windows reduce the total photon noise on the detectors. There are three major noise sources in a CMB telescope: readout noise (dependent on readout electronics), phonon noise (dependent on detector properties), and the dominant photon

noise (emission from astrophysical sources, the atmosphere, and the telescope). In this chapter, we will briefly discuss the main sources of noise in our detectors (§5.1). We will then pay particular attention to the noise in BICEP3; as the one of the longest running receivers in BICEP/*Keck*, it provides a wonderful testing ground for noise models. The measured and modeled noise within BICEP3 (§5.2 and §5.3) will be discussed, paying particular attention to how the three different windows on BICEP3 appear to have changed the noise within the receiver.

5.1 Receiver White Noise

5.1.1 Optical (Photon) Noise

We carefully design our telescopes to be 'photon noise limited': i.e., photon noise from optical components or astrophysical sources should be the dominant noise source on our detectors. Photon noise is directly related to the optical power $P_{\rm opt}$ on a detector, which in turn can be estimated from a model of the sky and optical elements in the instrument like:

$$P_{\rm opt} = \int_0^\infty [\sum_{i=1}^{N_{\rm elem}} p_i(\nu)] S(\nu) {\rm d}\nu \tag{5.1}$$

where the power spectral density p_i is the thermal emission at the detector from the ith optical element, accounting as described below for the intervening elements. The product of the detector response function $S(\nu)$ and the sum over all p_i is integrated over frequency to obtain the total loading. The power spectral density of an individual optical element is dependent on a number of properties, such as the element's temperature T_i, emissivity $e_i(\nu)$, scattering within the optic, the power on the element (dependent on the throughput

efficiency above that element, η_j), etc. For a given optical element, the dominant source of optical power from that element is generally the emissivity/absorptivity (e_i) in combination with the element's temperature: a transmissive optical component acts essentially like a gray-body[1]. Since the final optical power seen by a detector is integrated over the bandpass, the optics' temperature and absorption loss properties are highly degenerate. We can use a simple model with each component's emissivity and temperature to calculate the efficiency of all optics between a given element and the focal plane $[\eta_{i+1}(\nu), ..., \eta_{N_{elem}}(\nu), \epsilon_1(\nu)]$:

$$p_i(T_i, [\eta_{i+1}(\nu), ..., \eta_{N_{elem}}(\nu), e_i(\nu)]) = \prod_{j=i+1}^{N_{elem}} \eta_j(\nu)[e_i(\nu)B(T_i, \nu)] \tag{5.2}$$

For a diffraction-limited, single-moded polarimeter, the frequency and temperature dependent power spectral density is the Planck function:

$$B(T, \nu) = \frac{h\nu}{\exp[\frac{h\nu}{kT}] - 1} \tag{5.3}$$

The frequency dependent emissivity $e_i(\nu)$ of a given transmissive optic is directly related to the power absorbed (known as Kirchhoff's Law of thermal radiation):

$$e_i(\nu) = 1 - (\exp[-\alpha_i(\nu)t]) \tag{5.4}$$

where α_i is the material's absorption coefficient, and t is the thickness of the material. For a low loss dielectric material (which most of our optics are), the absorption coefficient is directly related to the complex relative permittivity[2] of a material, $\epsilon = \epsilon' + j\epsilon''$.

[1] A gray-body only partially absorbs and thermally re-emits radiation; it's essentially a fractional black-body. An object with partial opacity will only absorb and emit some fraction of incident power. If an element is, say, 1% absorptive and is at 300K, it will emit a spectrum like a 3K black-body.

[2] Historically, complex relative permittivity was called the dielectric constant.

The transmission loss, or the related absorption coefficient, are directly related to the complex part of the relative permittivity of a material. They describe the attenuation of light by the dielectric, and thus the loss the light-wave experiences as it is damped by the dielectric. For low loss materials, the loss tangent is related to the absorption coefficient or relative permittivity of a material:

$$\tan\delta = \frac{\epsilon''}{\epsilon'} = \frac{\alpha c}{2\pi n \nu} \tag{5.5}$$

where α is the absorption coefficient, c is the speed of light, n is the real part of the index of refraction and ν is the frequency of the light. I will characterize our optics' absorptivity with its associated $\tan\delta$ throughout this chapter. The permittivity of a material varies with temperature, and the $\tan\delta$ at millimeter wavelengths of many polymers decreases with cryogenic cooling, reducing the absorption and emission of a polymer optic in-band (Frank et al. 1977; Schnabel 2014). However, the reduction in optical emission at lower temperature is driven primarily by the lower physical temperature rather than the lower loss (Lamb 1996).

Having now defined a model of the amount of power expected to be emitted by each instrumental optical element, we must return our attention to the photon noise on our detectors. At short wavelengths (photon occupation numbers $p_n \ll 1$) the photon noise follows a Poisson distribution; essentially, background photons are particles, arriving separately and randomly (a.k.a. shot noise) with the expected distribution given by the square root of the number of photons received ($\propto \sqrt{P_{\rm opt}}$). At long wavelengths (photon occupation numbers $p_n \gg 1$), the photon acts less like a set of individual particles and more like waves, producing fluctuations in power proportional directly to $P_{\rm opt}$. At millimeter wavelengths, however, the average photon occupation number is actually very

close to unity[3] which means we cannot choose one model of photon noise over another, and must account for both (Zmuidzinas 2003). We do this with a quadratic sum:

$$\mathrm{NEP}_{\mathrm{photon}} = \sqrt{2\int_0^\infty [h\nu S(\nu)\sum_{i=1}^{N_{\mathrm{elem}}} p_i(\nu) + (S(\nu)\sum_{i=1}^{N_{\mathrm{elem}}} p_i(\nu))^2]\mathrm{d}\nu} \tag{5.6}$$

where $\mathrm{NEP}_{\mathrm{photon}}$ is the noise equivalent power (NEP) from photons on a detector.

The majority of the rest of this chapter (§5.2 onward) discusses the optical model for BICEP3, and how our uncertainties on our elements' optical parameters can strongly influence the predicted $\mathrm{NEP}_{\mathrm{photon}}$, and thus our final noise equivalent temperature (NET).

5.1.2 Other Sources of Noise

Bolometer Thermal Carrier (Phonon) Noise

The TES island (the bolometer) has very thin legs connecting it to the thermal bath, so noise can arise from fluctuations in heat flow between the bath and the island. This noise is dependent on the thermal conductance from the island to the bath like:

$$\begin{aligned} \mathrm{NEP}_{\mathrm{phonon}} &= \sqrt{4\ k\ F_{\mathrm{link}}\ T_{\mathrm{oper}}^2\ G} \\ F_{\mathrm{link}} &= \frac{\beta+1}{2\beta+3}\frac{1-(T_{\mathrm{bath}}/T_{\mathrm{oper}})^{2\beta+3}}{1-(T_{\mathrm{bath}}/T_{\mathrm{oper}})^{\beta+1}} \\ G &= P_{\mathrm{sat}}(\beta+1)\frac{T_{\mathrm{oper}}^{\beta}}{T_{\mathrm{oper}}^{\beta+1}-T_{\mathrm{bath}}^{\beta+1}} \end{aligned} \tag{5.7}$$

where T_{oper} is the TES operating temperature, T_{bath} is the bath temperature, k is the Boltzmann constant, G is the thermal conductance from the island to the bath, and

[3]The photon occupation number is found with $p_n = \exp[h\nu/kT]$. The occupation number at 3 K and 100 GHz ($\lambda \sim 3$ mm) is 1.6 ($p_n \simeq 1$).

F_{link} is a factor depending on the thermal conduction index β of the link (Mather 1982). In BICEP3, the thermal index for silicon nitride is assumed to be $\beta \sim 2$, and the operating and bath temperatures are 500 mK and 280 mK, respectively (BICEP/Keck Collaboration et al. 2022). The thermal conductance is also dependent on the detector absolute saturation power P_{sat}, which is designed to be 5 pW (in actuality, this varies detector to detector).

Johnson Noise

Johnson noise is the thermal noise within a circuit generated by the thermal agitation of charge carriers within the line. For the TES circuit, the Johnson noise is:

$$\text{NEP}_{\text{Johnson}} = \sqrt{4kT_{\text{oper}}R_{\text{TES}}\left(\frac{I_0}{L_{DC}}\right)} \tag{5.8}$$

where L_{DC} is the DC loop gain (assumed to be 20), R_{TES} is the resistance through the TES ($\sim 3\mu\Omega$ in BICEP3), I_0 is the current through the TES $\left(I_0 = \sqrt{\frac{P_{\text{opt}}}{R_{\text{TES}}}}\right)$, and T_{oper} is the operating temperature ($\sim 500mK$ in BICEP3).

Modern TES's are carefully designed to reduce the Johnson noise to be an order of magnitude below other noise sources (Galeazzi 2011; BICEP/Keck Collaboration et al. 2022). For example, at 95 GHz the phonon noise might be expected to be at order 1×10^{-17} W/$\sqrt{\text{Hz}}$ while the Johnson noise would be expected to be at order 1×10^{-19} W/$\sqrt{\text{Hz}}$. As such, we will be ignoring the Johnson noise when aggregating the total expected per detector noise, though Johnson noise should not be ignored entirely when considering detector design and characterization (Galeazzi 2011).

However, there is another resistor within the TES bias circuit: a shunt resistor parallel to the TES, which also produces Johnson noise (though we have traditionally

treated this specific resistor's Johnson noise separately from the above Johnson noise across the TES resistor):

$$\mathrm{NEP}_{\mathrm{shunt}} = \sqrt{4kT_{\mathrm{bath}}R_{\mathrm{shunt}}\left(\frac{I_0}{L}\right)} \tag{5.9}$$

where R_{shunt} is approximately 3 mΩ in BICEP3 and $L = \frac{L_{DC}}{L_{DC}-1}$.

This shunt noise is still sub-dominant to the photon and phonon noise, but generally an order of magnitude above the Johnson noise. BoloCalc[4] considers all Johnson noise sources to be negligible due to high loop gain, though only considers the Eq. 5.8 version to make this claim (Hill et al. 2018)[5]. Instead, BoloCalc considers what it calls 'Readout Noise' (covered in the next section), which functionally produces an equivalent small sub-dominant-but-not-completely-negligible noise source to the shunt noise in BICEP3. As such, I use BoloCalc's formulism with BICEP3 parameters, which reproduces an equivalent total per detector noise as methods using $\mathrm{NEP}_{\mathrm{shunt}}$, despite arriving at the small additional noise source by two different methods[6]. It is possible in BoloCalc to account for additional TES related noise sources by manually setting F_{link}, if needed.

[4]BoloCalc is the open source python library used for the sensitivity models discussed in this thesis. `https://github.com/chill90/BoloCalc`

[5]This calculation is not actually run through in the published BoloCalc paper; it is covered in the User Manual found on the git page: `https://github.com/chill90/BoloCalc`

[6]For a BICEP3-like configuration, the calculated $\mathrm{NEP}_{\mathrm{shunt}}$ is $\sim 1.8\times10^{-18}$ W/$\sqrt{\mathrm{Hz}}$, while the $\mathrm{NEP}_{\mathrm{read}}$ is $\sim 1.2 \times 10^{-18}$ W/$\sqrt{\mathrm{Hz}}$. These values are similar and an order of magnitude below the photon and phonon noise, so I treat them as equivalent. This may not hold true for another receiver.

Readout Noise

The current generation TES's are read out using superconducting quantum interference devices (SQUIDs) as amplifiers. The TES's are voltage-biased at the transition edge, and a change in power on the TES corresponds to a change in current at the SQUID. Therefore, the read out noise scales with both the voltage across the bolometer $V_{\rm bias}$ and the amplifier noise in the SQUID (typically reported as NEI). The bias voltage is determined from the TES resistance and the bias power[7], which in turn is the difference between the detector absolute saturation power $P_{\rm sat}$ and applied optical power $P_{\rm opt}$:

$$V_{\rm bias} = \sqrt{R_{TES}(P_{\rm sat} - P_{\rm opt})} \tag{5.10}$$

Therefore, the readout noise is:

$$\mathrm{NEP}_{\rm read} = \sqrt{(P_{\rm sat} - P_{\rm opt})R_{\rm TES}} \times \mathrm{NEI} \tag{5.11}$$

In the BICEP3 instrumentation paper we reported the SQUID/readout NEI to be 3 pA$\sqrt{\rm Hz}$, and the design $P_{\rm sat}$ to be 5 pW (BICEP/Keck Collaboration et al. 2022).

5.1.3 Noise Equivalent Temperature

We assume that all the NEPs (Eq. 5.6, 5.7, and 5.11) add in quadrature on a given detector. However, the actual fraction of power coupling to the detectors is configuration dependent: things like the conductance to a TES island will vary across a focal plane and the amount of optical power filtering through the optics is dependent on what optics

[7]In the previous chapter I referred to this value as $P_{\rm meas}$, as this is the margin to the absolute saturation power.

are in the receiver. NEP's are not particularly comparable between instruments, nor are they directly applicable to estimate sensitivity to temperature fluctuations, nor useful to estimate survey time. Therefore we must scale the found per detector NEP to an applicable external temperature. Two common signal references are the temperature of a 1 K Rayleigh-Jeans source or the CMB (a 2.725 K Rayleigh-Jeans source). As we are CMB instrumentalists, the most common choice is to scale to the CMB:

$$\mathrm{NET}_{\mathrm{det}} = \frac{\sqrt{\mathrm{NEP}^2_{\mathrm{photon}} + \mathrm{NEP}^2_{\mathrm{phonon}} + \mathrm{NEP}^2_{\mathrm{read}}}}{\sqrt{2}(dP/dT_{\mathrm{CMB}})} \tag{5.12}$$

where the conversion factor is:

$$\frac{dP}{dT_{\mathrm{CMB}}} = \int_0^\infty \left[\prod_{i=1}^{N_{elem}} \eta_i(\nu) \frac{1}{k_B} \left(\frac{h\nu}{T_{\mathrm{CMB}} \left(\exp\left[\frac{h\nu}{kT_{\mathrm{CMB}}} \right] - 1 \right)} \right)^2 \exp\left[\frac{h\nu}{kT_{\mathrm{CMB}}} \right] \right] B(\nu) \mathrm{d}\nu \tag{5.13}$$

Essentially, a detected fluctuation in power P is equivalent to a fluctuation in CMB temperature (T_{CMB}), and the conversion between is calculable based on the optical model we define with Equations 5.1 through 5.5.

5.2 Measured Noise in BICEP3

Each of our scansets (time-ordered azimuth rastered scans of the sky) contain both the signal and the noise of our measurement. We can get an estimate of the noise in the measurement by taking a Fourier transform of the raw pair-differenced data and then taking an average over the resultant noise spectrum. This noise must then be calibrated

out of raw detector values[8] using external publicly available temperature maps for reference and calibration (BICEP/Keck Collaboration et al. 2022; Planck Collaboration et al. 2020d). After calibrating our per scanset noise per detector, we take a median over all the scansets in a year to get a median per detector yearly $\mathrm{NET}_{\mathrm{det}}$.

Detector parameters (such as those relevant to $\mathrm{NEP}_{\mathrm{phonon}}$) vary across the focal plane, but should remain stable over time. Therefore there should be a consistent distribution of measured $\mathrm{NET}_{\mathrm{det}}$s over the detector population, which we do see between years (see Figure 5.1). Optical changes should result in a bulk offset in P_{opt}, and therefore a coherent offset in measured $\mathrm{NET}_{\mathrm{det}}$ across the focal plane. We take a median of the $\mathrm{NET}_{\mathrm{det}}$ per year to use as a representative measured NET value for that year, which are also the center bar in each of the boxes in Figure 5.1. There are three occasions when the NET change between years was greater than 1.5%[9]: between 2018 and 2019, 2021 and 2022, and 2022 and 2023.

5.2.1 Uncertainty on Absolute Calibration

The calibration of our measured $\mathrm{NET}_{\mathrm{det}}$ are dependent on external maps, and therefore uncertainty on our measured $\mathrm{NET}_{\mathrm{det}}$ are dependent on the variance between our 'real' maps (m_{real}), and the reference and calibration maps (m_{ref} and m_{cal}, respectively). That

[8]Internally, we typically refer to the raw detector output as either "feedback units" (FBU) or "analog to digital units" (ADU) interchangeably, though they actually do refer to slightly different conversion processes.

[9]The significance of the change in measured NET is dependent on the uncertainty on the absolute calibration.

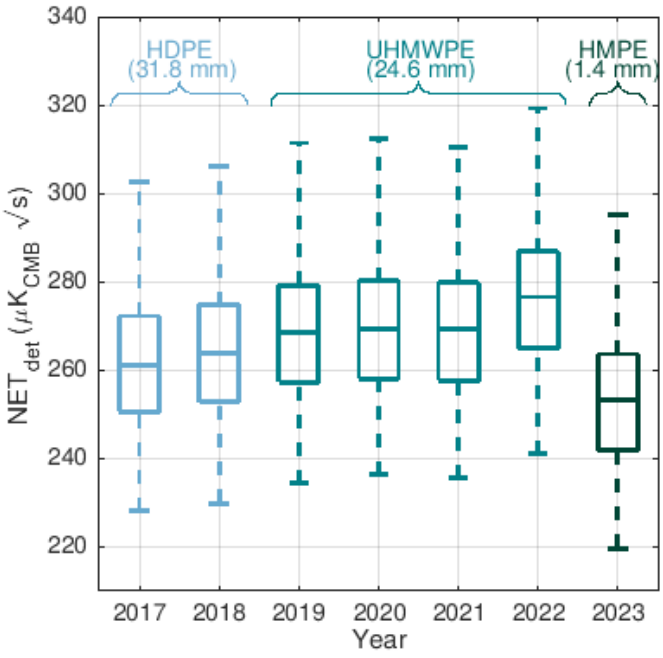


Figure 5.1: BICEP3 median white noise (NET) distribution over all detectors per year between 2017 and 2023. Per detector per scanset (roughly 50 minutes of time ordered data) NETs are calculated, then a median per detector over the entire year's scansets are shown, with 25th (q_{25}), 50th, and 75th (q_{75}) percentiles shown in the box, to demonstrate average noise of the receiver between years. Whiskers cover the remainder of the distribution, excluding outliers, defined by $q_{25/75} \pm (q_{75} - q_{25})$. Color corresponds to the kind of window on the receiver: light blue is HDPE 31.8 mm (1.25"), blue is UHMWPE 24.6 mm (0.97"), and dark green is HMPE 1.4 mm. There is a significant change in noise between three observing seasons: 2018–2019, 2021–2022, and 2022–2023.

variance goes like:

$$\mathrm{var}\left(\frac{\langle m_{\mathrm{cal}} \times m_{\mathrm{real}}\rangle_\ell}{\langle m_{\mathrm{real}} \times m_{\mathrm{ref}}\rangle_\ell}\right) \approx \left(\frac{\alpha_\ell}{C_\ell}\right)^2 \left(\mathrm{var}\left(\langle sn^*\rangle_\ell\right) + \mathrm{var}\left(\langle nn_r^*\rangle_\ell\right)\right) \tag{5.14}$$

where α_ℓ is the abscal, C_ℓ is the signal cross spectra, $\langle sn^*\rangle_\ell$ is the signal crossed with the noise, and $\langle nn_r^*\rangle_\ell$ is the reference map noise crossed with the real map noise. We expect the reference map (*Planck*) noise to be negligibly small and the actual magnitude of the abscal should not be relevant, so in practice, the actual uncertainty on the abscal becomes:

$$\mathrm{var}\left(\frac{\langle m_{\mathrm{cal}} \times m_{\mathrm{ref}}\rangle_\ell}{\langle m_{\mathrm{real}} \times m_{\mathrm{ref}}\rangle_\ell}\right) \approx \frac{\mathrm{var}\left(\langle sn^*\rangle_\ell\right)}{C_\ell^2} \tag{5.15}$$

All of the ℓ bins shown in Figure 5.2 have fractional absolute calibration uncertainties below 0.005 (0.5%), but the uncertainties for the range over which we take the abscals (ℓ ∼72-177) is less than 0.002 (0.2%). The real maps are integrated over a three year data set, so we may expect the yearly abscal uncertainty to be higher than 0.2%. However, the real dataset covers 2016–2018, and 2016 is known to have had significantly higher noise than 2017 and 2018[10]. Even if the yearly uncertainty is a factor of three higher than these coadded uncertainties, that would put the uncertainties at well less than 0.5% in the relevant bins (this is equivalent to ∼1 $\mu K\sqrt{s}$ noise) making the greater than 1.5% differences we are seeing between three seasons significant.

All three of these significant changes in noise are correlated with optics changes. These changes were:

[10]In 2016 BICEP3 had metal mesh filters where the HD30 foam filter stack is currently (recall Figure 1.9), which caused a significant amount of excess scattered power. Most of the metal mesh filters were removed during the 2016–2017 austral summer, except for the low-pass edge filters just above the detector tiles, and one just below the 50 K alumina filter.

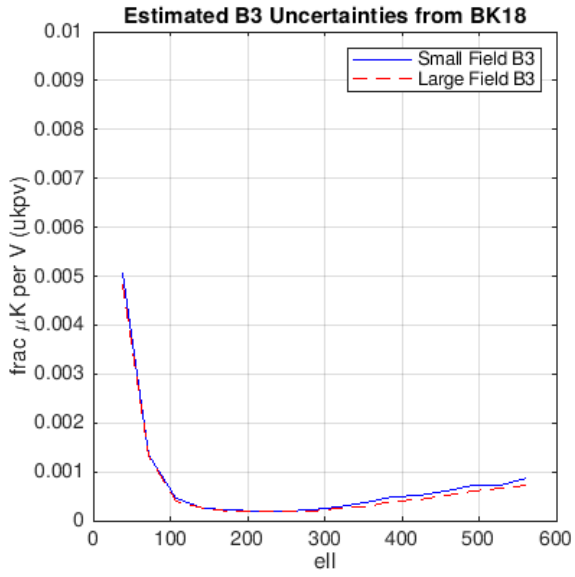


Figure 5.2: Calculated estimates of the abscal uncertainties for BICEP3 from BICEP/Keck Collaboration et al. (2021) signal and noise cross-spectra using Eq. 5.15 across all BICEP/*Keck* ℓ bins and both BICEP3 fields (the small field is the field covered by BICEP2 and *Keck* Array, the large is the full BICEP3 FOV). The ℓ bins used for absolute calibration are 72-177, which remain well below 0.002 for both fields.

1. Between 2018 and 2019: A window change. The window thickness decreased (31.8 mm to 24.6 mm), and the material changed (HDPE to UHMWPE). The median per detector yearly NET increased 1.7% (263 $\mu K\sqrt{s}$ to 268 $\mu K\sqrt{s}$).

2. Between 2021 and 2022: High optics temperatures. A leaky flange led to many cryogenic optics' temperatures being elevated. The median yearly NET increased 2.7% (269 $\mu K\sqrt{s}$ to 276 $\mu K\sqrt{s}$).

3. Between 2022 and 2023: A window change. The window thickness decreased (24.6 mm to 1.4 mm), and the material changed (UHMWPE to HMPE laminate). The median per detector yearly NET decreased 8.6% (276 $\mu K\sqrt{s}$ to 252 $\mu K\sqrt{s}$).

5.3 Modeled Noise in BICEP3

We have always made estimates of noise seen by our detectors: a significant fraction of the calibration measurements we take are dedicated in some fashion to this issue, from measured detector resistances to forebaffle on/off measurements. However, we have never had as long or as sensitive an instrument as BICEP3. The only optical changes on BICEP3 since 2017 have been window changes, allowing an exciting experimental comparison between different window configurations.

Between the radiometric material insertion data reported in Figure 4.5 and the median yearly NETs reported in Figure 5.1, we have a solid set of measurements that can be used to test our models. I use median detector parameters reported in BICEP/Keck Collaboration et al. (2022) for the other sources of noise, which were also transcribed in the previous section. As a cross check, I also compare the final updated model to the median yearly light detector P_{meas} reported in Table 4.2 and the related change in P_{const} (§4.2.2), but I did not attempt to fit to those measurements, given their uncertainties.

In this section we shall explore the previous BICEP3 optical models (which were used to produce loading estimates in the BICEP3 instrumentation paper, BICEP/Keck Collaboration et al. (2022)), what optical model parameters may be inaccurate in that model, and how an updates to those parameters may explain the measured NETs and radiometric measurements.

Table 5.1: Table of parameters in the BICEP3 loading spreadsheet (circa 2018), including calculated loading.

Optic	Temperature [K]	$\tan\delta$ [e-4]	Thickness [mm]	Estimated Optical Load [pW]
Forebaffle	230			0.136
Window	280	2.22	31.75	0.69
Foam Filter	280			0.097
Alumina Filter	55	1.2	10.0	0.046
Obj Lens	5.5	1.2	20.9	0.046
Apt Stop	5.5			0.078
Upper Nylon Filter	6.0	1.13	5.0	0.001
Field Lens	5.5	1.2	26.9	0.011
Lower Nylon Filter	5.5	1.13	9.5	0.002

5.3.1 Previous Model

I report the optical parameters used in the 'old model' (optical model used for BICEP/Keck Collaboration et al. (2022)) in Table 5.1 and that model's predictions for the measured values (both NETs and inserted material radiometric loads) in Table 5.2. Note that the old model predicts a reduction in noise from 2018 to 2019 (the loss properties are assumed to be the same between HDPE and UHMWPE, so as the window thickness went down the load should as well), and a very large reduction in noise with the thin window in 2023 (again, with the same loss properties the optical load should reduce significantly). The model also overestimates the optical load from the inserted

Table 5.2: Measured NETs and insertion optical loads compared to the predicted NETs and optical loads from the old model outlined in Table 5.1.

	2017/2018 NET $[\mu K\sqrt{s}]$	2019/2020/2021 NET $[\mu K\sqrt{s}]$	2022 NET $[\mu K\sqrt{s}]$	2023 NET $[\mu K\sqrt{s}]$	31.8 mm UHMWPE [pW]	1.4 mm HMPE [pW]
Measured	261, 263	268, 269, 269	276	252	0.39	0.06
Meas. Uncer.	± 1	± 1	± 1	± 1	± 0.05	± 0.01
Old Model	266	256	256	222	0.711	0.060

UHMWPE slab radiometric load.

These discrepancies between the model and the measurements point to two potential issues with the old model: that we may be overestimating the fractional load from the window (because the thin window did not reduce the noise as much as predicted) and that the polyethylene losses between different grades may not be as consistent as we would assume (because of the increase in NET in 2019/2020/2021 and the smaller than expected difference between UHMWPE slab and HMPE optical load).

Let us discuss where I believe the old model may have inaccurate optical parameters.

Forebaffle

Over the last few summer seasons at Pole we have upgraded our forebaffle load measurement strategy[11]. Previous forebaffle measurement campaigns suffered from focal plane temperature instabilities (remember that the detector $P_{\rm sat}$ is dependent on the bath temperature, as in Eq. 4.5). These temperature instabilities were seemingly related to vibrational loads traveling through the receiver, likely sourced from removing and replacing the forebaffles above the telescope. Changes to forebaffle procedure—such as extra care taken in placing the forebaffles, and monitoring the focal plane temperature during the procedure such that PLCs were taken only when the temperatures stabilized—resulted in significantly higher measured forebaffle loads than were previously reported in BICEP/Keck Collaboration et al. (2022). Recent (as of late 2024) measurements with aluminum lining placed inside the forebaffles have corroborated the high forebaffle load at approximately 0.34 pW on BICEP3. This is approximately twice as high as was previously reported and modeled.

[11]The 2022/2023 season had a very careful forebaffle measurement with particular attention paid to focal plane temperature instabilities (recall that $P_{\rm meas}$ is dependent on the bath temperature of the focal plane, so temperature fluctuations can complicate PLC optical loading measurements). The 2024/2025 season introduced aluminum lining forebaffle load tests, which approximate removing the forebaffle by inserting an aluminum sheet within the forebaffle (reflecting scattered power to cold sky instead of coupling to the warm absorptive forebaffle).

Window Loss

The HDPE window loss estimate is well supported by other measurements of PE loss[12]. However, note how the old model predicts a radiometric load for the UHMWPE slab of 0.711 pW in Table 5.2, when the actual measured load was 0.39 pW. The material insertion radiometric measurements reported in §4.2.1 seem to point to a discrepancy in the bulk PE loss: if the HMPE window and the UHMWPE slab had the same loss properties, the change in optical load between them should have been a factor of 10, not a factor of 6.

Additionally, the measured NET *increased* when the window thickness *decreased* (31.8 mm to 24.6 mm), and the material changed (HDPE to UHMWPE). This points to some amount of variability between either the loss properties of these slabs, or some kind of coupling change between windows (the window aperture did increase slightly when we swapped windows, to allow for a 'BA-style' window clamping mechanism). Either way, if we attempt to solely explain the NET differences between years with the window loss properties, then the UHMWPE loss must be higher than the HDPE loss, and the HMPE loss must be even higher still.

Alumina Optics

A comprehensive review of alumina $\tan\delta$'s are not the purview of this thesis, so I will only give a short summary here. It is difficult to compare many alumina loss measurements not only because of measurement strategy differences between studies,

[12]Barkats et al. (2018) has a good review of how PE loss may vary with frequency, and 2.22×10^{-4} is well inline with those estimates at 95 GHz.

but also due to the formulation and structure of what different manufacturers refer to as 'alumina.' It is therefore difficult to determine which properties are actually associated with the alumina used in BICEP3 from literature.

However, I do want to note that many studies (and internal measurements) find an alumina $\tan\delta$ upwards of 2×10^{-4} (Lamb 1996; Inoue et al. 2014; Nadolski et al. 2020). As such, we should not take the relatively low alumina loss of 1.2×10^{-4} initially used as gospel; there is ample evidence that the true loss may be higher. I will also note that I (and others) have a tendency to hold the alumina loss constant between the both the 50 K filter and 4 K lenses: there is evidence that alumina loss may change between those two temperatures[13]. However, it is difficult to estimate how strong an influence temperature may have on our alumina loss throughout BICEP3 without independent measurements of the exact materials used at different temperatures, so I will continue to keep the loss consistent for both the alumina filter and lenses.

Nylon Filters

The nylon loss tangent reported in Table 5.2 is 1.13×10^{-4}, which is significantly lower than any other reported nylon loss tangent I have found. Lamb (1996) puts the nylon loss at *minimum* 8.8×10^{-4} (at 100 GHz and 300 K), while many (both internal and external) studies report losses orders of magnitude higher. We have some reason to

[13]The primary evidence of the loss in alumina changing with temperature comes from an ASTRON 191 report from 2019 showing a shift from $5.3\pm1.2\times10^{-4}$ at 77K to as low as $2.5\pm0.1\times10^{-4}$ at 3K. Nadolski et al. (2020) does find a significant shift in $\tan\delta$ from $49\pm6\times10^{-4}$ at 300K to $3\pm1\times10^{-4}$ at 77K, though that does not indicate that an additional large change in loss occurs at lower temperatures.

suspect that the nylon loss does drop significantly at low temperatures (potentially going from order 100×10^{-4} to 10×10^{-4}), but not below the loss of alumina (Afsar 1987; Birch et al. 1981; Halpern et al. 1986). As such, the nylon loss is likely much higher than what was assumed previously, by at least a factor of five.

Additionally, nylon is well known to be highly hygroscopic, and exposure to the atmosphere[14] for extended periods can allow enough water to enter the polymer to significantly change the absorptive properties. Therefore, nylon filter loss has the potential to vary considerably even between filters in our own instruments, simply due to different atmospheric exposure histories. It is also difficult to directly compare absorption loss measurements in the wetter climes outside the South Pole. The only comfort in this issue is that once the filters are cold, whatever water is adsorbed should remain, and it is unlikely that the nylon loss varies significantly between years.

Optics Temperatures

We have recently taken measurements of the windows on cold receivers at the South Pole, as reported in Table 2.5. In summary, we found that the windows were actually at around 260 K, not 280 K, in their centers.

All the optics temperatures deserve a bit of scrutiny. Photon power from cold optics

[14]The atmosphere at the South Pole is incredibly dry, so it is unlikely that significant amounts of water vapor enter the nylon filters at the South Pole, but pre-deployment receiver testing and shipping is likely to expose the nylon to high humidity environs for extended periods. On the other hand, anti-reflection coatings may slow adsorption and exposure to the dry atmosphere at the South Pole may reduce the equilibrium water absorption within the filters. It is therefore extremely difficult to predict how much water vapor may be in any of our filters without a dedicated study.

will only vary slightly with inaccurate optical temperatures. For instance, changing the temperature of a nylon filter from 5 to 6 Kelvin (a 15% increase) could increase the total optical power by 0.004 pW (only a 0.2% increase). The hotter optics contribute a more significant fraction of the load; changing temperature of the window from 280 to 260 Kelvin (a 7% reduction) would reduce the total optical power by about 0.1 pW (about a 4% reduction). But in the aggregate, if all the cold optics temperatures are truly higher than those assumed in Table 5.1 and the window temperature is truly lower, then the combined misestimated temperatures may have disguised the true loading fraction from the window.

In the next section we will discuss the optics temperatures in the receiver over the seasons.

5.3.2 Previous Model with New Temperatures

We have a few sets of thermometers on our optics, which are read out and recorded in our raw arcfiles[15]. I went through and identified each and recorded the median temperature for nine thermometers during the first and last minute of each tag (approximately an hour of time ordered data) over the entire observation season. These thermometers are generally at the edges of the optics, so do not necessarily give us the actual equilibrium temperature at the center of the optic. However, they do give a decent estimate of

[15]Arcfiles are the raw data output of our receivers, and include everything read out through the general control program (GCP), such as raw thermometry voltages, calibrated thermometry, raw detector timestreams, etc. These data are generally just arrays of time-ordered values, so interpreting which thermometer is which requires consulting the documentation.

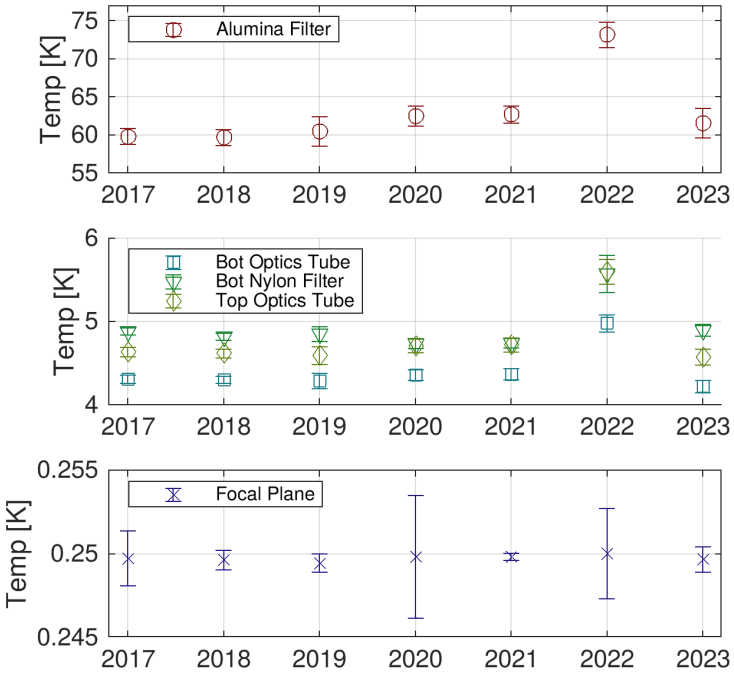


Figure 5.3: Median per year cold optics temperatures in BICEP3. Top panel is the alumina filter thermometer, middle panel are thermometers in the optics tube (at the bottom of the tube, the bottom nylon filter, and the top of the optics tube), and the lower panel is an NTD on the focal plane. Note that for every year except 2022 the optics seem to have consistent median temperatures.

the boundary condition of the optic, and it is unlikely that the center equilibrium temperature is lower than the temperature at the edge. Therefore, it's not unreasonable to estimate that the center of the optics are slightly warmer than the edge, but not significantly warmer.

In the old model the alumina filter was definitely assumed to have too low of a temperature, then, as the median yearly edge temperature of the filter in Figure 5.3 is generally 60 K, not the previously assumed 55 K. I think it's also reasonable to bump up the other cold optics temperatures slightly, particularly the nylon filters as they

Table 5.3: Measured NETs and insertion optical loads compared to the predicted NETs and optical loads from the old model outlined in Table 5.1, with temperatures that change between years.

	2017/2018 NET $[\mu K\sqrt{s}]$	2019/2020/2021 NET $[\mu K\sqrt{s}]$	2022 NET $[\mu K\sqrt{s}]$	2023 NET $[\mu K\sqrt{s}]$	31.8 mm UHMWPE [pW]	1.4 mm HMPE [pW]
Measured	261, 263	268, 269, 269	276	252	0.39	0.06
Meas. Uncer.	± 1	± 1	± 1	± 1	± 0.05	± 0.01
Old Model New Temps	262	253	268	222	0.659	0.057

have high conductance and the center equilibrium temperature is probably higher than the edge temperatures. I will admit, however, most of the regular operation bumps in temperatures (aside from the alumina filter) did not have a stronger basis than a guess. Since temperature and loss are highly degenerate, slight variations in the optics temperatures and loss tangents reported in this chapter may reproduce the same final results.

In Figure 5.3 it is evident that all the cold optics temperatures increased in 2022; all optics in the 4 K optics tube seem to have increased in temperature by about a kelvin, while the 50 K alumina filter was hotter by about 10 kelvin. Crucially, the window edge temperature did not change significantly in 2022[16]. The new baseline temperatures

[16]This does not rule out the equilibrium center temperature of the window potentially changing by a few degrees, particularly if the window center temperature is dominated by radiative thermal coupling

are outlined in Table 5.4. If we update the previous model with the new temperatures (essentially using the temperatures from Table 5.4 plus the increased temperatures in 2022, but the optical parameters in Table 5.1), we get the predictions in Table 5.3. The increased temperatures in 2022 do make the predicted NET for 2022 closer to the measured increase, but it does not quite get high enough: this likely indicates that the alumina loss should be higher than previously assumed to better match the increase in noise with a significantly warmer cold optic. Additionally, the other issues with the old model predictions remain: overestimates of the UHMWPE slab optical power and the decreased NET in 2023, and underestimates of the 2019/2020/2021 NETs (likely indicating that we are overestimating PE losses).

The very large increase in temperature in most cold optics in 2022 provides an interesting opportunity, as the fractional increases in temperature were not all created equal: while all the cold optics temperatures increased significantly (by about 20%), those that increased by 20% in temperature at 60 K will produce more optical power than the increases at 5 K. If we assume that the alumina loss stays constant at different temperatures, then we get a wonderful natural experiment: the high NET in 2022 must be explainable by a set of optics with elevated temperature that also produce the NETs and inserted material optical load at the nominal temperatures.

Table 5.4: Table of parameters in the modified BoloCalc model with the updated optical parameters, including calculated loading. Note that a given model in Table 5.5 will only have one of the italicized windows below.

Optic	Temperature [K]	$\tan\delta$ [e-4]	Thickness [mm]	Estimated Optical Load [pW]
Forebaffle	240			0.337
HDPE Window	*255*	*0.75*	*31.75*	*0.203*
UHMWPE Window	*255*	*1.38*	*25*	*0.297*
HMPE Window	*260*	*2.50*	*1.4*	*0.057*
Foam Filter	280			0.075
Alumina Filter	65	2.6	10.0	0.115
Obj Lens	6.0	2.6	20.9	0.016
Apt Stop	5.5			0.097
Upper Nylon Filter	6.5	9.5	5.0	0.011
Field Lens	5.5	2.6	26.9	0.023
Lower Nylon Filter	5.5	9.5	9.5	0.018

Table 5.5: Measured NETs and insertion optical loads compared to the predicted NETs and optical loads from the new model outlined in Table 5.4.

	2017/2018 NET $[\mu K\sqrt{s}]$	2019/2020/2021 NET $[\mu K\sqrt{s}]$	2022 NET $[\mu K\sqrt{s}]$	2023 NET $[\mu K\sqrt{s}]$	31.8 mm UHMWPE [pW]	1.4 mm HMPE [pW]
Measured	261, 263	268, 269, 269	276	252	0.39	0.06
Meas. Uncer.	± 1	± 1	± 1	± 1	± 0.05	± 0.01
New Model	263	270	274	252	0.378	0.056

5.3.3 New Model with New Temperatures

The updated optical parameters for the new model[17] are primarily composed of decreasing the loss from the windows and increasing the loss in the cold optics, which helps to better match the smaller reduction in NET than predicted with the thin window. I have increased both the alumina loss tangents (from 1.2×10^{-4} to 2.6×10^{-4}) and the nylon loss tangents (from 1.13×10^{-4} to 9.5×10^{-4}): the relatively small increase in alumina loss best reproduces both the increased noise in 2022 and keeps the optical throughput high enough that the material insertion radiometric loads are correct (if the cold optics are too lossy the measured material load goes down). Solely increasing the alumina loss does not increase the cold optics load enough to explain the smaller-than-expected

to the cold optics.

[17]In actuality, this optical model is a set of three models: one the nominal configuration, one recreating the temperature increases in 2022 and one with an additional inserted optic to recreate the radiometric material insertion load measurements. Each model shares the listed optical parameters.

Table 5.6: Difference in loading compared to 2017 of both per season median per light detector P_{meas} and modeled P_{opt} differences compared to 2017. For every year except 2023 the offset is negative, indicating that most years have a higher median optical load than what was applied to the detectors in 2017.

	2018	2019/2020/2021	2022	2023
Offset Median Light P_{meas} (pW)	-0.033	-0.073, -0.108, -0.101	-0.223	0.096
Modeled P_{opt} Diff (pW)	0.00	-0.08	-0.15	0.14

reduction NET in 2023, so the nylon loss must increase as well. Increasing the nylon loss alone does not adequately explain the increase in NET in 2022, as the nylon filters are much colder than the alumina filter. The balance between the warmer alumina filter and cold nylon filter losses must together explain the 2022 data, while also keeping the other predictions consistent in the nominal configuration.

The new model was used to predict how varying the band center of the instrument impacts optical load changes with different windows and the predicted NET varies with window thickness in Figure 2.1. Other in-lab room temperature loss measurements of HDPE and UHMWPE have found that they have similar loss tangents at around $\sim 1.4 \pm 0.2 \times 10^{-4}$, potentially constraining both PE losses at less than the $2.22{\times}10^{-4}$ assumed in the previous model (Elwood et al. 2024). Anti-reflection coat losses are not modeled on any of the optics in this chapter: it is possible that individual optics' losses are biased upwards due to this, particularly for optics with an electrical length similar

Table 5.7: Summary full instrument parameters produced by the new optical model for each year.

	2018	2019/2020/2021	2022	2023
Optical Thruput	0.315	0.314	0.314	0.318
Total $P_{\rm opt}$ (pW)	2.10	2.18	2.25	1.96
Total $P_{\rm inst}$ (pW)	0.98	0.99	1.06	0.75
Tel Temp (K_{RJ})	8.68	9.62	10.24	7.24

to the electrical lengths of the AR coats[18]. However, adding additional materials with additional parameter uncertainties is an added layer of complexity that is not necessarily required: the true loss tangent of thin windows may in actuality be lower than those reported in Table 5.4, but the optical power coming from the window system on the detectors functionally remains the same if we model the window system as a lossier window or a mildly lossy window with mildly lossy AR coats.

In Figure 5.4 we show how varied window loss tangents around 1.4×10^{-4} affect the final calculated NET and how this compares to the measured median yearly NETs. Different loss or scattering properties between HDPE and UHMWPE may explain the increase in noise seen from 2018 to 2019, when the window material was changed. I have presented how different PE loss tangents may explain the change in NET from 2018 to

[18]For instance, a 1.4 mm window with $\tan\delta = 2.5\times10^{-4}$ would produce about 0.065 pW at 150 GHz and 260 K, while a thin window ($\tan\delta = 1.4\times10^{-4}$, thickness 1.4 mm) and AR layer system (two layers, $\tan\delta = 1.4\times10^{-4}$, thicknesses 0.6 mm) would also produce combined 0.061 pW. Thus the actual loss of the thin windows may be obscured by the combined load with the anti-reflection layers.

2019 in Table 5.5, but other material or physical coupling changes could produce the same jump.

The predicted change in total P_{inst} of 0.23 pW between 2017 and 2023 from the new model is within the uncertainty of the measured difference in P_{const} of 0.06 $\pm$ 0.50 pW discussed in §4.2, while the old model predicts a larger change in P_{const} of 0.65 pW with old temperatures and 0.59 pW with new temperatures. Additionally, the predicted P_{opt} differences for each year are relatively close to the median light P_{meas} offsets, as shown in Table 5.6. These consistencies act as reasonable confidence builders in the predictive capabilities of the model, though there remain some unexplained inconsistencies, such as the under prediction of the change in P_{opt} in 2022 and over prediction of the change in P_{opt} in 2023 in Table 5.6.

The new modeled optical throughput (shown in Table 5.7) is predicted to be 0.315, which is close to the measured optical efficiency of 0.317 reported in BICEP/Keck Collaboration et al. (2022). The new model is not directly contradicted by any other related measurement, to my knowledge, which gives confidence that the model in its current form is a decent predictive tool. Direct and independent measures of the optical parameters of the optics would be welcome, even if they contradict the parameters used here, because this solution is unlikely to be truly unique. It is likely, for instance, that the $\tan\delta$'s used for our cold optics vary with temperature, or that the optics temperatures assumed are slightly inaccurate. However, the BICEP3 optical model as it currently stands seems to be a relatively robust prediction of the current suite of measured NETs and radiometric loads in BICEP3, which is a satisfactory achievement.

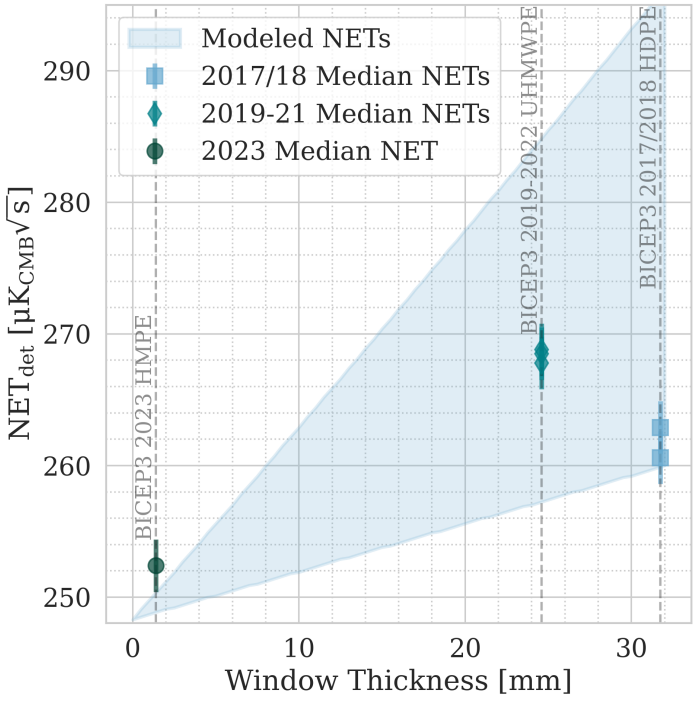


Figure 5.4: Median measured NETs (white noise) from Figure 5.1 as a function of the window thickness for that respective year. The NET from 2022 is excluded due to the elevated optics temperatures. The blue shaded band of NET models show predicted BICEP3 NET values, varying the window loss tangent between 0.6×10^{-4} to 2.4×10^{-4}. We have not yet concluded on the cause of the increase in the white noise measured between the thicker HDPE window (31.8 mm) in 2018 to the UHMWPE window (25.4 mm) in 2019: both a difference in material loss and a difference in scattering properties could explain this change. The HMPE window did definitively and significantly decrease the measured white noise compared to both thick slab windows.

Chapter 6

Receiver Bandpass Characterization

We create multiple receivers to observe at different frequencies such that we can separate out components along the line-of-sight that produce similar polarized intensity. Recall that we expect the polarized frequency spectral behavior of different sources—such as synchrotron, thermal galactic dust, or the CMB—to generally follow simple power laws at degree angular scales in our low foreground patch. As discussed in 1.2.3, we use an eight parameter model to predict the angular auto- or cross-power spectra for different frequency maps[1]. However, these models depend on our understanding of the spectral behavior of the instruments themselves; if we expect that BICEP3 has a frequency bandpass centered at 95 GHz, and it is truly centered at 93 GHz, the best fit parameters in our model will correspondingly not reflect the 'true' parameters in our patch. We therefore need to precisely determine the spectral response of our receivers.

[1]Throughout this chapter, I will attempt to be consistent in distinguishing between spatial or optical frequency spectra by referring to the angular (spatial) power spectra as power spectra, and optical frequency spectra as frequency spectra explicitly.

The end-to-end spectral response of our detectors (the bandpass) is characterized with two parameters, the bandcenter $\langle\nu\rangle$ and the bandwidth $\Delta\nu$. The bandcenter is defined as:

$$\langle\nu\rangle = \int \nu \; S(\nu) \; d\nu \tag{6.1}$$

and the bandwidth is:

$$\Delta\nu = \frac{\left(\int S(\nu) \; d\nu\right)^2}{\int S(\nu)^2 d\nu} \tag{6.2}$$

where $S(\nu)$ is the frequency dependent response of the detector, which we will call the per-detector spectrum.

We then determine a combined spectra from individual detectors by coadding the resultant per-detector spectra in a variety of combinations, such as by detector module (aka tile), polarization (recall that we have co-located orthogonal antenna which are coupled to separate bandpass filters and TES), or over the entire receiver. The full receiver coadded spectra are generally the reported spectra in publications.

The bandpass is carefully designed to avoid the atmospheric absorption lines in-band (as shown previously in Figure 1.7), but even a small amount of coupling to the absorption lines would significantly increase the load on the detectors[2], which can be seen clearly in Figure 6.1. The increased total load at high frequencies is also apparent.

[2]This is a potential area of improvement in the noise modeling from the previous chapter; I currently model the optical loading with a 'top-hat' (binary) bandpass defined by the bandcenter and fractional bandwidth; the actual frequency spectral response of the detectors is not that ideal, and therefore the integrated optical power measured by the detectors will likely be slightly different than the power predicted using the naive top-hat bandpass.

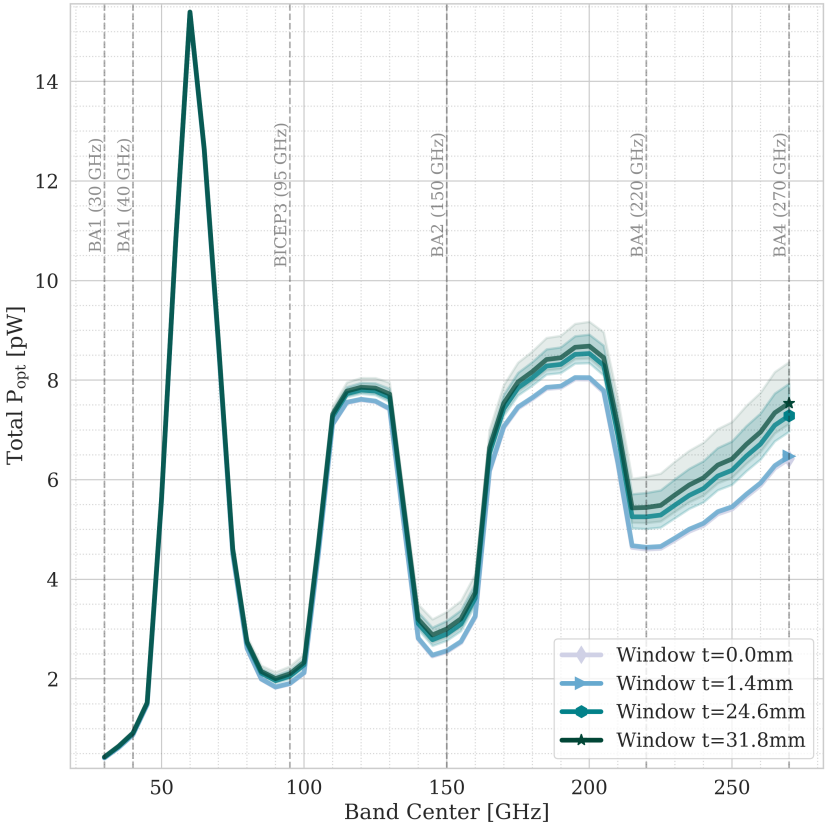


Figure 6.1: The modeled optical power on the detectors of a BICEP3-like receiver with varied bandcenter from 30 to 270 GHz, fractional bandwidth of 25% and different thicknesses of windows. The colored bands around the lines denote uncertainty in the loss of polyethylene, with the solid lines having a tanδ of 1.4×10^{-4}. The same model (which is described in §5.3) was used to predict the optical load from the windows themselves in Figure 2.1. Dashed lines denote the assumed bandcenter of our deployed receivers.

We measure the bandpasses of our detectors in-situ with a Fourier transform spectrometer (FTS) placed above our receivers. This chapter discusses the interferometer, the bandpass measurements (including changes to procedure in the last few years), and how the statistical uncertainty derived from those measurements affects our cosmological parameter estimates.

6.1 Fourier Transform Spectroscopy

6.1.1 Interferometer

The design of our FTS is a Martin-Puplett interferometer (Martin & Puplett 1970), which is conveniently a polarized interferometer. The time reversed receiver beams enter the square input port, are reflected (and potentially polarized) by an input mirror/grid into the interferometer. The beam is then split into two polarized beams by a grid rotated to lie along a cube diagonal (35.2° to the baseplate), and oriented at 45° to the input beam vector, therefore reflecting half of the power down two arms, one with a stationary rooftop mirror and one on a moving Velmex BiSlide[3] stage. The rooftop mirrors rotate the polarization of their respective beams by 90°, such that the recombined beams exit the beamsplitting grid towards the source, where they are redirected by a parabolic mirror to be split by the source wire grid to the two thermal sources, which are HR-10 microwave absorber at two different temperatures. One source is the wall of the FTS at ambient temperature, and one is submerged in liquid nitrogen in a styrofoam box at approximately 77 K. The difference in temperature between the sources sets the nominal strength of the signal. The design and initial characterization of our FTS is described further in Karkare et al. (2014). A full accounting of the electromagnetic fields in a Martin-Puplett interferometer can be found in Lesurf (1990).

The spectral resolution $\Delta\nu$ is set by the largest path length difference between the arms of the interferometer ($\Delta\nu = c/\Delta L$). For our FTS, ΔL is about 51 cm, so the resolution is about 0.5 GHz. We have specifically chosen to place the stationary mirror

[3]Velmex: `velmex.com/Products/BiSlide`

at roughly half the distance of the full range of the moving mirror from the beamsplitter; this creates symmetric two-sided interferograms for each half scan of the moving mirror, which provides a cross-check on the mirror alignment.

That alignment is crucial to optimize, as our FTS has a smaller aperture (∼127 mm) than the receiver aperture (∼550 mm), and is also situated in the near-field beam (so it is also not a point source). We thus expect that there is a frequency dependent aperture coupling effect we must account for in our spectral measurements.

The beam exiting our single-moded receiver goes like:

$$\Omega_{\rm beam}(\nu) = \iint d\theta d\phi B(\nu, \theta, \phi) \tag{6.3}$$

while we expect that the detector effective area ($A_{\rm eff}$) and beam ($\Omega_{\rm beam}(\nu)$) should be related to wavelength as:

$$A_{\rm eff}\Omega_{\rm beam}(\nu) = \lambda^2. \tag{6.4}$$

We expect that the power coupled into the receiver will depend on the beam

$$P = \int d\nu\ \eta(\nu) A_{\rm eff}(\nu) I(\nu) \iint d\theta d\phi B(\nu, \theta, \phi) \tag{6.5}$$

where $\eta(\nu)$ is the optical efficiency at different frequencies (essentially, the bandpass) and $I(\nu)$ is the spectrum of the source. We have a thermal source, and our frequencies should be firmly in the Rayleigh-Jeans tail of the blackbody, so the source should scale with ν^2.

The power coupled into the receiver from a beam filling thermal source should be:

$$P = \int d\nu\ \eta(\nu) A_{\rm eff}(\nu) I(\nu) \iint d\theta d\phi B(\nu, \theta, \phi) \tag{6.6}$$

$$= \int d\nu\ \eta(\nu) A_{\rm eff}(\nu) I(\nu) \Omega_{\rm beam}(\nu). \tag{6.7}$$

Since $A_{\text{eff}}\Omega_{\text{beam}}(\nu) = \lambda^2$, this cancels the thermal source spectrum ν^2 dependence, and the measured spectrum should just be $\eta(\nu)$.

However, we do not have a beam filling source. For a point source:

$$P = \Omega_{src} \int d\nu \; \eta(\nu) A_{\text{eff}}(\nu) I(\nu) B(\nu, \theta_0, \phi_0) \tag{6.8}$$

which is dependent on how the detector effective area behaves with frequency. If the effective area is constant with frequency, then it will drop out of the integral and the coupled power scales with $\nu^2\eta(\nu)$. But the FTS is not a point source either. It has a small aperture in the near field, so how the measured bandpass scales with frequency likely lies somewhere between the beam filling and point source cases (between ν^0 to ν^2).

The frequency dependence of the fraction of power coupled through both apertures can be computed to estimate the appropriate power law. The beams should scale with frequency in a predictable manner, and therefore it is easy to compute the integral of coupled power for beams of different scales within the band (see the left of Figure 6.2). We can take the beam-width found from near-field beam mapping, and—assuming that the measured near-field beam is band averaged, and therefore representative of the beam at the nominal center frequency—compute the frequency dependent receiver to FTS aperture throughput fraction (see the right of Figure 6.2). The slope of that coupled fraction will be the exponent (a) of the power law correction we scale the measured spectra by, which goes like ν^a.

Previously we used a pair of translation stages to physically move the FTS atop the receiver to pass the center of the Gaussian beam through the FTS. As the center of the near field beam is dependent on the pointing of a detector, the required translation of the FTS for a given tile is simple to find using the expected pointing angle and

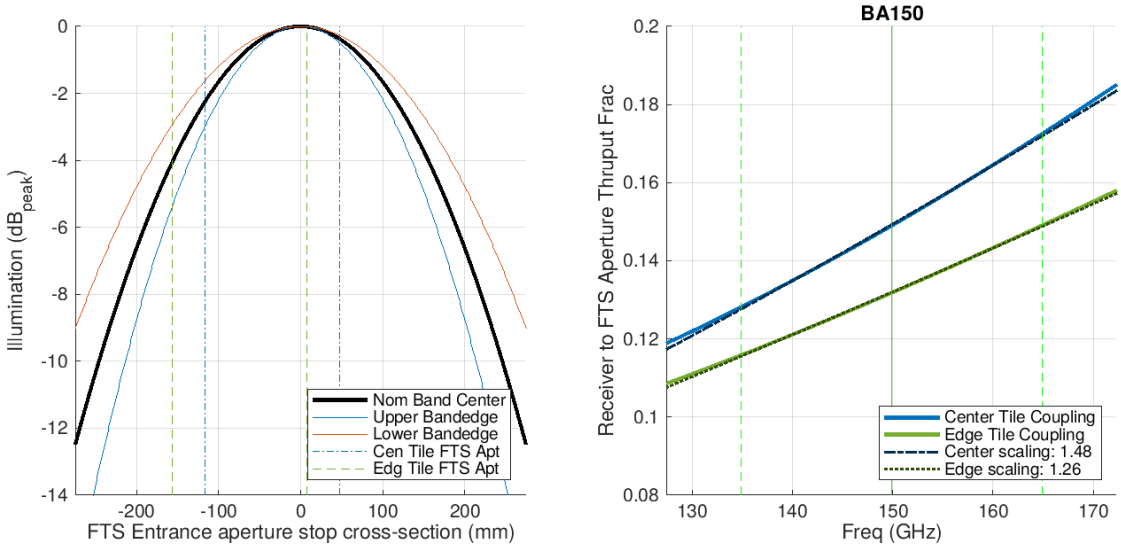


Figure 6.2: [Left] The center slice of the frequency dependent near field beam, with the FTS entrance aperture truncation overlaid at the rough positions of a center (Cen, blue dash-dotted) and edge tile (Edg, green dashed). [Right] Frequency dependent receiver to FTS aperture throughput fraction for the center of a center (blue, upper lines) and edge tile (green, lower lines). Fit aperture correction factors are shown in the legend (for a center tile, $a = 1.48$, for an edge tile, $a = 1.26$).

trigonometry. However, the translation stages were found to cause systematics in the spectral measurement and were removed prior to the deployment of BICEP Array. A carefully drawn grid on the plywood base of the FTS was meant to allow manual translation of the FTS for different tiles; this translation has not occurred for any of the BICEP Array FTS measurements. This change in procedure is illustrated in Figure 6.3, assuming that the light-path of the FTS can be essentially 'unfolded' and modeled as if the input grid/mirror moves the source aperture position to couple to the beam.

If the FTS is not translated to couple to the center of the Gaussian beams of a given

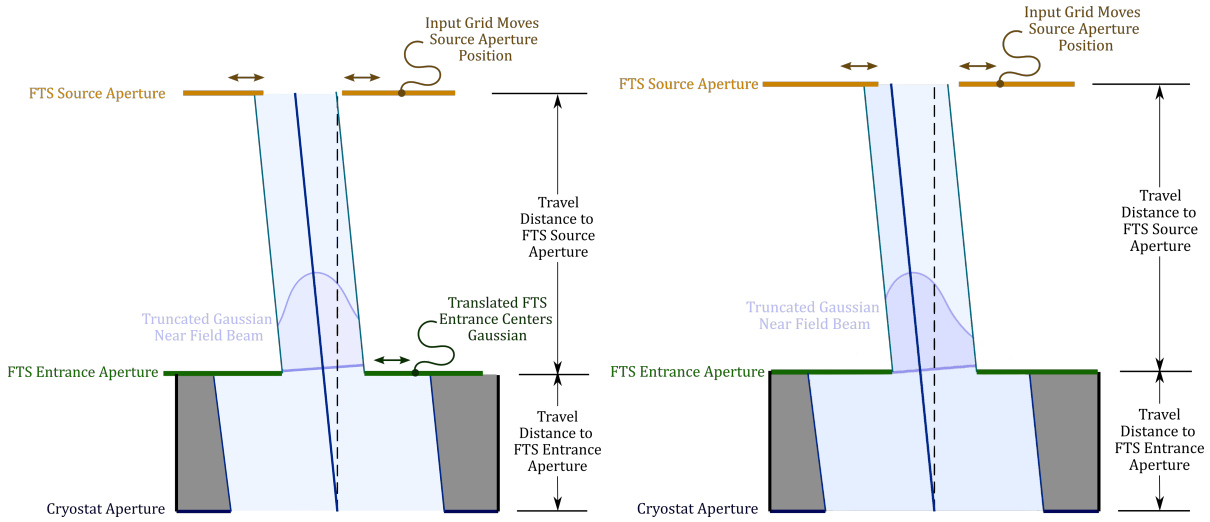


Figure 6.3: Cartoon diagrams of the instrument near field beam coupling with the FTS apertures. [Left] Previous procedure, which translated the FTS entrance aperture to center the Gaussian near field beam. [Right] Recent procedure, which centers the FTS center aperture over the receiver, relying on the input grid to couple any given detector pointing through the FTS.

tile, then the width of the beam that enters the FTS input aperture will be different across the focal plane. In other words, the FTS aperture will truncate the beam of an edge tile detector in a different place than the beam of a center tile (see the placement of the FTS aperture on the near field beam in the left of Figure 6.2, or the modeled near field truncation of four corners of an edge tile at the square input aperture in Figure 6.4). Therefore, we have to compute a different frequency scaled correction factor for center and edge tiles.

Additionally, we generally stopped taking multiple measurements of the same tile with different input grid positions (called 'pointings') due to increased time pressure. These were useful to both better couple to detectors across a tile, which may have missed

the FTS source aperture in a single center pointing, and test the model of the aperture correction[4] and measurement systematics. We do have yearly calibration set of data taken on BA150 with two pointings, one pointed at the row 3/col 3 detector (an edge pointing), one pointed at the row 9/col 9 detector (a center pointing); less than the three to nine pointings per tile that were taken for *Keck* Array FTS calibration data, but potentially useful to test the accuracy of our aperture correction calculation.

Therefore, we have to not only determine the aperture correction for an immobile FTS, but how the beams across a given detector tile will be truncated by the two apertures (the square entrance aperture and the circular source aperture) within the FTS for a single pointing[5].

The actual coupling fraction will vary across tiles dependent on the pointing. As such (and since we very rarely have more than one pointing per tile for the BA FTS campaigns) we can calculate the per detector aperture correction factor, assuming that we have one good center pointing per tile. This leads to a per detector aperture

[4]Tyler St. Germaine found better agreement between four different pointings on the same detector with slightly modified aperture correction factors applied based on the source aperture truncation, though it did not result in perfect agreement. He attributed the slight disagreement between corrected spectra to internal misalignment within the FTS, which is not currently included in the model (St. Germaine 2021).

[5]The single pointing FTS campaigns typically attempted to optimize the pointing by maximizing the amplitude of the interferograms of the detectors at the center and four corners of the tile. This appears to have resulted in the center detector beam passing through the FTS relatively unobstructed (the beam being truncated only by the source aperture), whereas detectors towards the edge suffered greater truncation by the square entrance aperture.

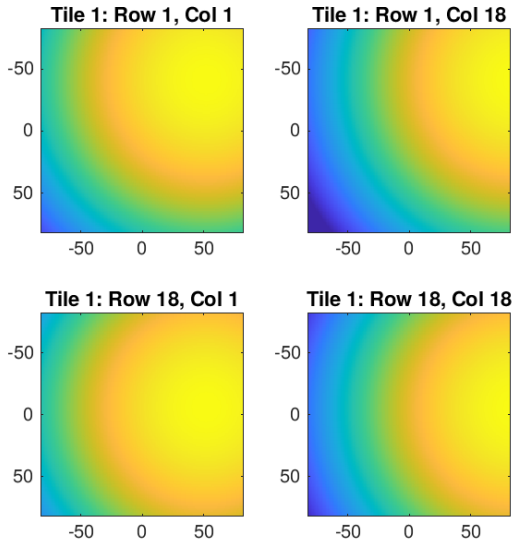


Figure 6.4: FTS entrance aperture truncation of the near field beam, centered on the entrance aperture, for all the corners of an edge tile (Tile 1) in BA150.

correction factor (see Figure 6.5 for a model of the expected bandcenter shift from the aperture truncation), instead of the per receiver aperture correction factor applied in the past; therefore the correction factor needs to be applied to the per detector spectra, not, as has been done in the past, as a single correction for any given spectra in a receiver.

6.2 BA150 Spectra

We have added tiles to BA150 during each summer season; the receiver initially deployed in austral summer 2022/2023 with five of twelve tiles. We deployed an additional five tiles in the austral summer 2023/2024, and the final two tiles in austral summer 2024/2025. FTS measurements were taken inside the building, before the receiver was lifted into the

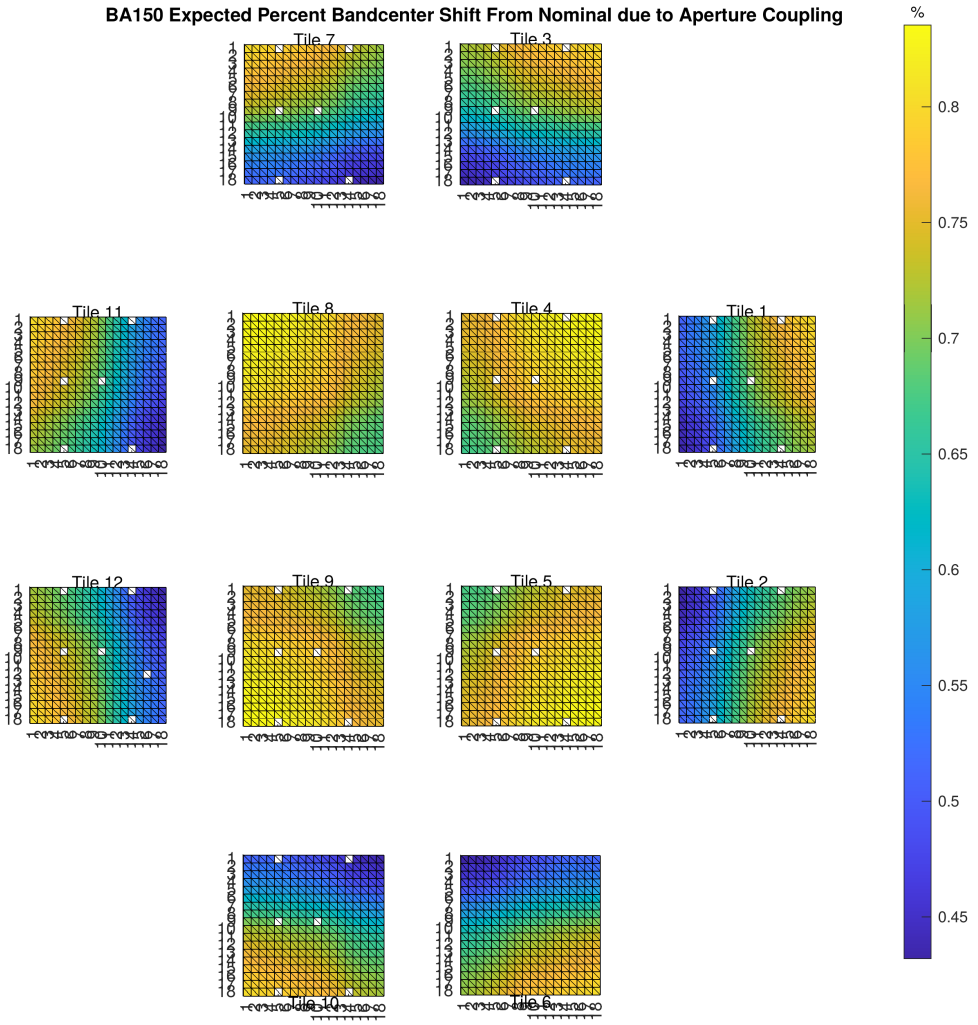


Figure 6.5: The expected percent bandcenter shift caused by the aperture truncation across the BA150 focal plane. An ideal top-hat (binary) bandpass is multiplied by the calculated ν^a for a detector, then the bandpass parameters are determined using Eq. 6.1 and 6.2.

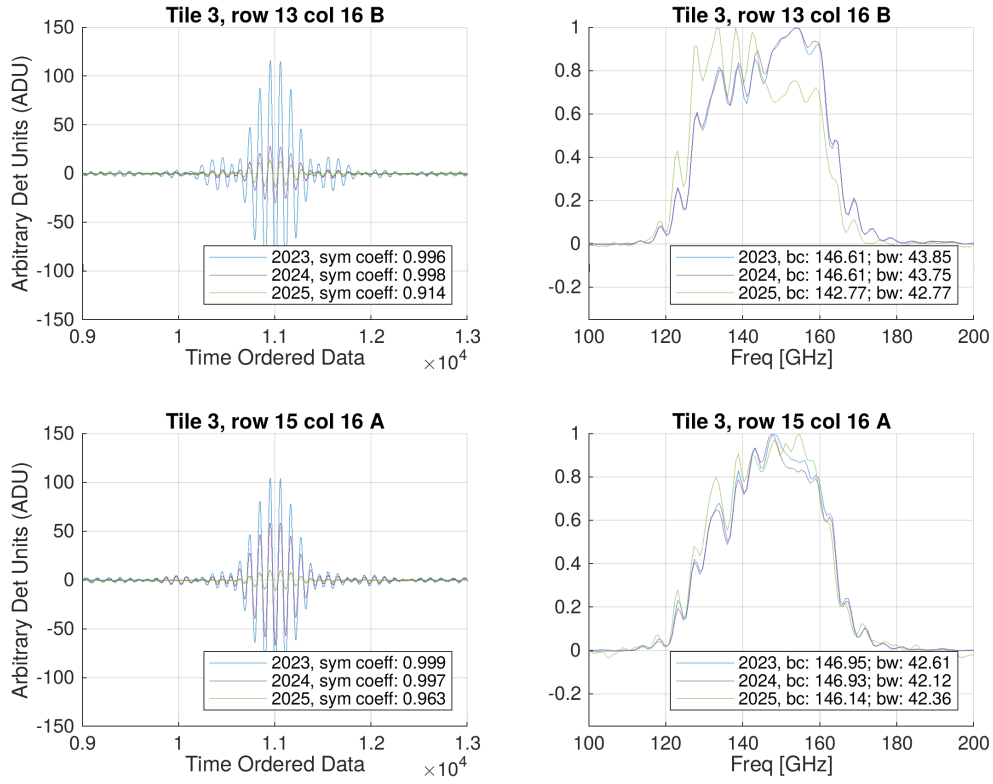


Figure 6.6: Example interferograms and spectra from two detectors (A and B polarization) for the three years of data taken on the BA150 instrument on the same tile. Color corresponds to calibration campaign year (blue is 2023, purple is 2024 and green is 2025). On the left the symmetry coefficient (Eq. 6.9, an estimate of interferogram quality) for each interferogram is reported in the legend, and the recovered band parameters are reported in the legend on the right.

mount[6]. Recent summer calibration seasons have had abbreviated calibration campaigns due to increased operational constraints. The truncated timeline has lead to significantly less FTS data being taken during recent campaigns compared to pre-COVID campaigns; previous FTS campaigns included multiple pointings per tile at the same FTS rotation. Recent years only had one or two pointings per tile. This has limited our ability to test systematics within the FTS in the same configuration[7]. We can test year to year variations by comparing the measured spectra on the same detectors.

6.2.1 Analysis

The raw FTS data are detector timestreams concurrent with the moving mirror jogging through its range, which produces interferograms. We have previously used an encoder on the moving mirror to determine the temporal location of interferograms in the detector timestreams; many of the campaigns in recent years have not had the encoder actively recording, and as such the white light fridges (WLF) of the interferograms are determined through a rough guess to narrow the range for a peak finding algorithm to search a known good detector timestream. Interferograms per detector are then subsequently identified, polynomial filtered and windowed; the third order polynomial filter is meant to correct thermal drifts in the ambient half of the thermal load. After

[6]Many of the *Keck* spectra were taken outside on the roof (St. Germaine 2021). BICEP Array receivers typically take calibration measurements before being loaded in the mount, either in the compressor room below the mount or in lab spaces adjacent to the mount.

[7]The FTS is fully disassembled between campaigns, and therefore alignment is likely to vary campaign to campaign.

the WLF for a given interferogram is identified (more on that below), the timestream is windowed with either a triangle (for 30/40 GHz detectors) or Blackman-Harris window (all higher frequency detectors) to reduce "ringing" in the spectral domain, at the cost of spectral resolution. The interferogram is then passed through a unitary Fourier transform, with real and imaginary parts recorded separately. All resultant spectra from a set of interferograms—generally four to seven in a single measurement—are then averaged together to average down noise. Examples of single raw interferograms and subsequent frequency spectra for the same A and B detectors are shown in Figure 6.6.

Two FTS data quality metrics are useful to define here. The first is the symmetry coefficient. We take an interferogram, divide it into a left (L) and flipped right (R) portion around the suspected area of the WLF and then take a covariance matrix of the two resultant vectors. If the halves of the sides of the interferogram are more symmetric, the covariance between them will be close to the variance of each side. The metric is therefore just the covariance scaled by the geometric mean of the variances, or mathematically:

$$c_{sym} = \frac{\mathrm{cov}(L, R)}{\sqrt{\mathrm{var}(L) + \mathrm{var}(R)}}. \tag{6.9}$$

This symmetry can be reduced by either alignment mismatches within the FTS[8] or a low signal to noise ratio, as noise is inherently non-symmetric. Therefore, this metric is best utilized as a qualitative estimate of data quality for a single interferogram or—as it was designed to do—find the precise location of the WLF, but not necessarily a quantitative

[8]If, say, the moving mirror is misaligned such that more power couples through the source aperture during the first half of the interferogram compared to the second half, then we would see a less symmetric interferogram.

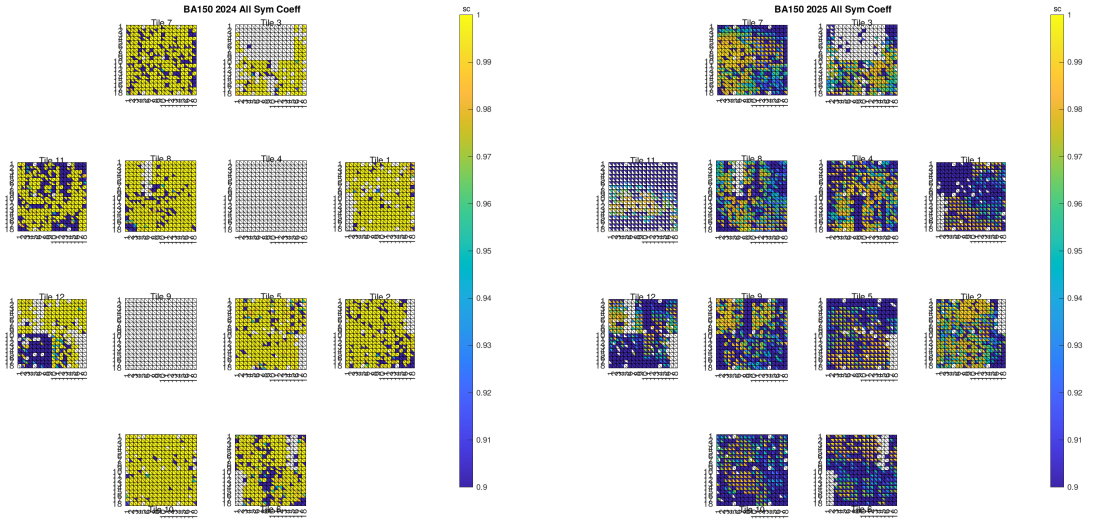


Figure 6.7: Per detector symmetry coefficients (defined in Eq. 6.9) for all measured BA150 detectors from the 2024 (left) and 2025 (right) FTS campaigns. Brighter colors (yellow) indicate a high symmetry coefficient close to 1, darker colors (dark blue) indicate a low symmetry coefficient below 0.9. The 2025 campaign exhibits significantly less consistently high symmetry coefficients.

measure for data cuts, or assessing which of a set measurements on a given detector are best.

The second metric is the mean of the raw spectra within a predetermined band (for BA150, between 100 and 200 GHz). This is a measure of the signal to noise, as a higher signal interferogram will also result in a higher signal spectra. The mean spectrum is used when evaluating the 'best' spectra for a given detector, if we have multiple measurements of the same detector from different pointings or rotations[9] of the FTS. We

[9]Rotations of the FTS are from changing the orientation of the FTS relative to the boresight. The FTS can either be aligned such that the moving mirror is aligned with one linear polarized detectors (A

normalize the spectra when comparing them, either with peak normalization (dividing the spectra by the maximum value within a band) or integral normalization (dividing the spectra by the area under the curve).

Both these metrics having been defined, it is worth stating that, for both metrics, the 2025 BA150 campaign shows consistently worse performance for both symmetry coefficients and mean spectra (see Figure 6.7 for a comparison between the symmetry coefficients across the focal plane in 2024 and 2025): this is likely due to the very low amplitude interferograms, particularly compared to previous years. I am uncertain at this time what exactly is the cause of this severe reduction in data quality in the 2025 dataset: the FTS could have been misaligned (which would cause the low amplitude interferograms and the low symmetry coefficients across the focal plane), but the interferograms are not visibly less uniform in 2025. The low symmetry coefficients could be biased down because of the low signal to noise ratio. That being said, there are cases in previous datasets where we have had low power interferograms with excellent symmetry that produce very consistent spectra year to year; the found spectra for like detectors in 2025 tend to significantly diverge from their previous measurements.

The only optical change in BA150 for the 2025 campaign beyond the addition of the final two detector modules was another low-pass edge (LPE) filter (Ade et al. 2006). This new filter was intended to removed a suspected high frequency blue leak, and to hopefully correct the high NETs measured in BA150 from calibrated timestream noise[10]. The new LPE filter was added in addition to the previously installed LPE

or B, called a 0/90 deg rotation), or 45 degree offset from the detectors.

[10]These NETs were determined using the same technique described in §5.2.

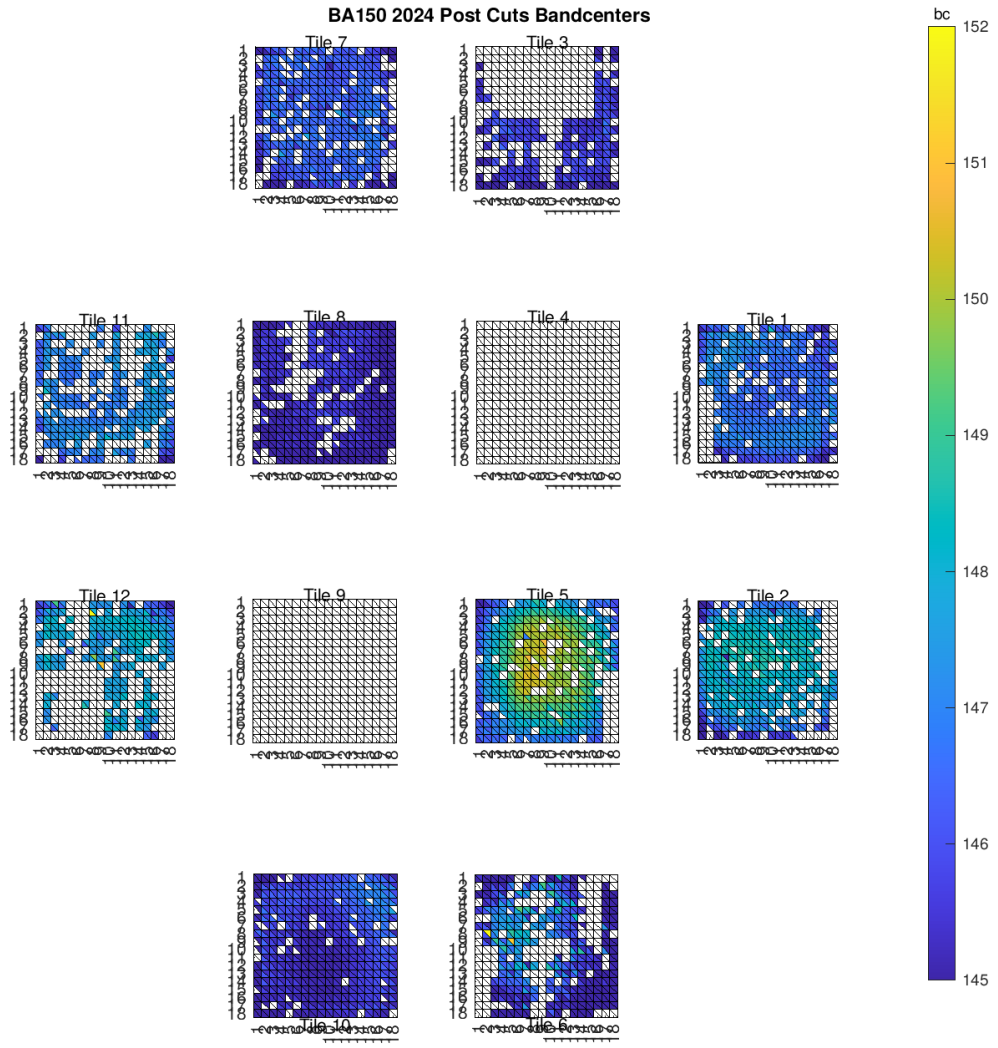


Figure 6.8: The aperture corrected bandcenters of all detectors that passed cuts in 2024. Variation in calculated bandcenter across a given tile is high (generally greater than a few percent), so it is difficult to determine if the applied aperture correction reduces the expected bandcenter variation across the focal plane as predicted in Figure 6.4.

filter. High pass filter measurements with thick grill filters before and after the LPE filter installation seem to have confirmed a reduction in high frequency coupled power, and optical efficiency measurements seem to indicate a 20% reduction in overall optical efficiency. The focal plane wide reduction in interferogram amplitude may be in part due to this reduced optical efficiency (which should also reduce per detector gain), but we seem to have observed an order of magnitude reduction in gain between the 2024 and 2025 interferograms. Additionally, the 2025 near field beam maps show a variety of odd features, which may be a sign of unexpected consequences of the new filter on the beam profile[11].

Regardless, as the 2025 observation season is ongoing the CMB data have not yet been reduced, and are unlikely to be incorporated into the current work on the next BICEP/*Keck* data release. As such, I will focus on what can be gleaned from the 2023 and 2024 FTS campaigns for a systematic study of the impact of spectral uncertainties on our cosmological parameter recovery.

6.2.2 Spectral Uncertainties

For BA150, we have three FTS campaigns from the three austral summers since the receiver was deployed. Comparing the same detectors between years, we get very similar bandcenters in 2023 and 2024, as shown in Figure 6.9, with the median difference

[11]Each of these claims (high frequency power, optical efficiency, and near field beam maps) are preliminary results derived by others and subject to reevaluation. I am merely attempting to report the current survey of measurements that may point to explanations for why the 2025 FTS dataset seems to be so poor.

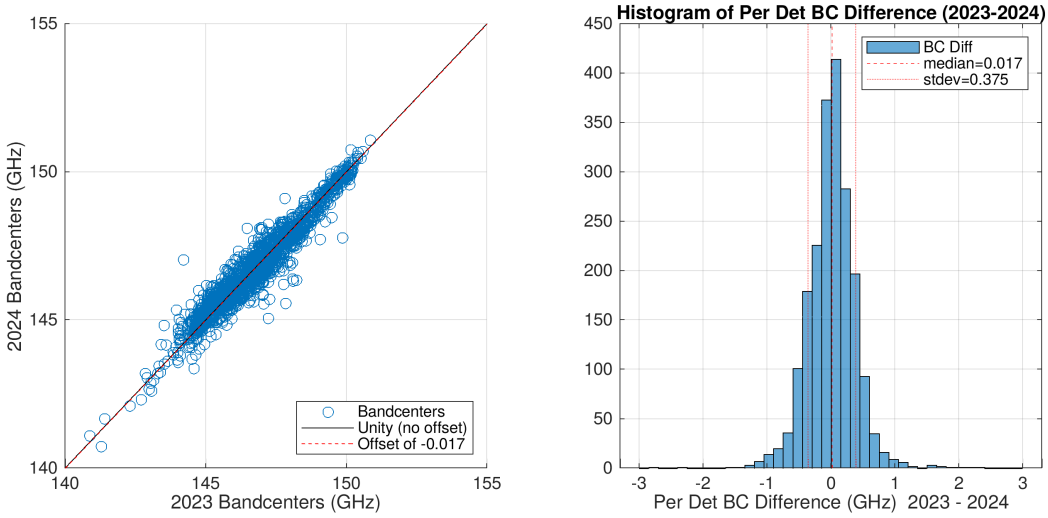


Figure 6.9: Per detector bandcenters (BC) from 2023 and 2024 FTS campaigns, with per detector aperture correction applied. [Left] Scatter plot of found bandcenters compared between 2023 (x axis) and 2024 (y axis). Unity is noted with a solid black line, the median offset with a dashed red line. [Right] Histogram of the difference of the 2023-2024 per detector bandcenters. The median of the difference between the per detector bandcenters is 0.017 GHz (dashed red line), and the standard deviation of the difference is 0.375 GHz (dotted red lines).

between bandcenters at 0.017 GHz. There is a small amount of spread in the calculated bandcenters between years, which is likely sourced from measurement uncertainty: the standard deviation of the difference in bandcenter between 2023 and 2024 is 0.375 GHz, or about 0.25%, and the differences between bandcenters has a Gaussian distribution. It is encouraging that the year to year variation is very small, particularly since the FTS is fully disassembled and realigned between campaigns. This does not rule out the possibility, however, that there exists a systematic shift in the bandcenters due to our measurement procedure: recall that without aperture correction we expect the calculated

bandcenters to shift upwards between 0.4-0.8% (Figure 6.5). If the applied per detector aperture correction incorrectly corrects the detector spectra then these data may also be systematically biased away from the true bandcenters.

The full receiver coadded spectra in Figure 6.10 have very consistent shape and bandcenters, barring some minor differences in shape at the lower frequencies. The new detectors added in 2024 appear to have shifted the coadded receiver spectra by about 0.2 GHz. This is likely representative of a real shift in the receiver spectral response between the 2023 and 2024 seasons. I expect that the final two new tiles added in 2025 will likely contribute to a significantly smaller shift in the coadded receiver bandpass, as there are was a smaller percentage increase in detector count (in 2024 the nominal detector count increased 100%, in 2025 the nominal detector count increased 20%). However, given my current understanding of the data quality metrics for the 2025 campaign, I am not confident that the data is of sufficient quality to accurately determine the coadded receiver bandpass, and we will not compare the coadded bandpasses.

There are two features of the aperture corrected coadded BA150 spectra from 2023 and 2024 in Figure 6.10:

1. The coadded spectra for the two orthogonal polarizations have different peaks (as visible on the left side of Figure 6.10). The separation of the peaks remain in both 2023 and 2024: it is possible that this differential frequency response between the detectors has resulted in some kind of polarized systematic effect, such as increased temperature to polarization leakage in BA150[12].

[12]T-to-P leakage is the main source of systematic uncertainty on our measure of r (BICEP2 and Keck Array Collaborations et al. 2019; BICEP/Keck Collaboration et al. 2021). A study of the effect of T-to-P

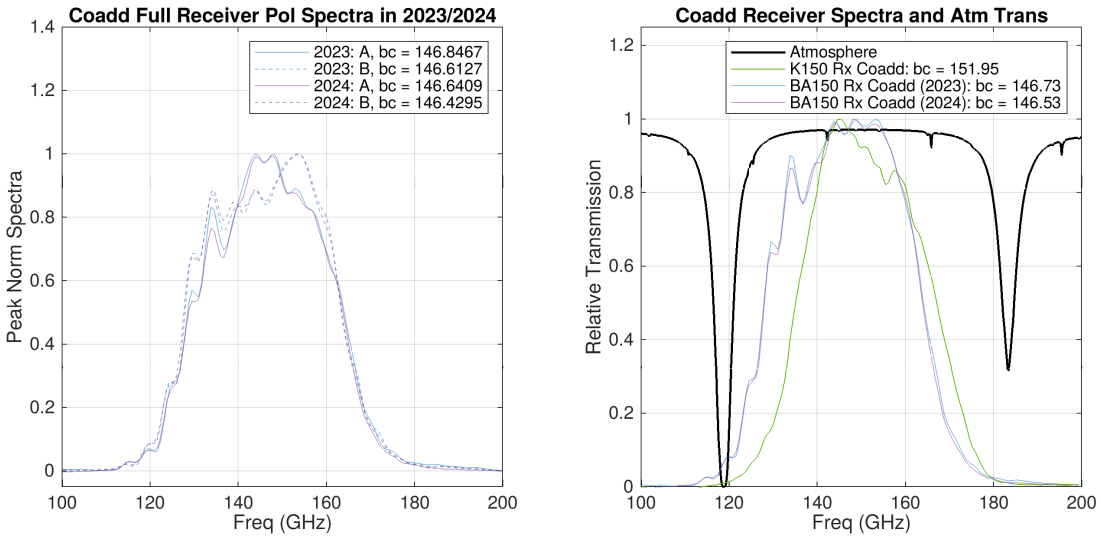


Figure 6.10: Coadded per detector aperture corrected spectra for BA150. [Left] The per polarization (A detectors denoted with solid lines, B detectors with dashed) coadded spectra and peak normalized for 2023 (blue) and 2024 (purple). [Right] Full receiver coadded spectra for K150 (green) is shown as a comparison to the BA150 full receiver coadds for 2023 (blue) and 2024 (purple). Modeled atmospheric transmission at South Pole at an elevation of 60 degrees shown in black (Paine 2024).

2. There appears to be significant pick up in the 118 GHz oxygen line (shown on the right of Figure 6.10). The relative transmission at the oxygen line is about 6% of the peak for the aperture corrected coadded spectra (and about 4% for the non-aperture corrected coadded spectra, which are not shown).

The coupling to the oxygen line may partially contribute to the high noise measured in the receiver, as the coupling to the band would add optical load, as previously discussed and shown in Figure 6.1. We have measured higher than expected noise in the

leakage on the dataset up to 2024 is currently underway.

receiver, potentially due to higher than expected optical loading which we previously attributed to a blue leak, but high optical load might as well come from this coupling to a lower frequency atmospheric line. That low frequency coupling would not have been corrected by the additional LPE filter added in the 2024/2025 austral summer season. Both of these features could significantly impact the performance of the receiver by introducing extra noise due to differential pickup on the atmosphere or on the oxygen line. More work is needed to relate these features to our current understanding of the performance of the receiver.

6.3 Component Separation

Recall that the main reason that we observe the sky with multiple frequency bands is to separate out the components along the line-of-sight. There are three main sources of polarized emission at the scales and frequencies we observe. They are galactic dust and synchrotron, and the CMB. We expect these sources to have distinct optical and spatial frequency dependence, which allows us to separate these components.

Therefore, we conduct an eight parameter maximum likelihood search on our data, which consist of the bandpowers at 9 multipoles of a variety of auto– and cross–spectra. The previous released dataset (BICEP/Keck Collaboration et al. (2021), hereafter BK18), used five external (WMAP 23 and 33 GHz; *Planck* 30, 44, 143, 217, and 353 GHz) and four internal maps (BICEP2/*Keck* Array 95, 150, and 220 GHz, and BICEP3 95 GHz)[13]. The next anticipated data release (hereafter referred to as BK24), is expected

[13]I will be referring to these maps as K95, K150, K220 and B95.

to include data up to either the 2023 or 2024 season, which will add an additional 270 GHz *Keck* band, as well as the BICEP Array 30, 40 and 150 GHz bands[14]. The eight parameter model consists of amplitudes and optical frequency/spatial spectral exponents (A_d, A_s, β_d, β_s, α_d, α_s) for the galactic components, a parameter to characterize spatial correlations between the polarized dust and synchrotron (ϵ) and the primordial tensor to scalar ratio (r).

For the previous dataset, we determined that a conservative upper limit on the bandpass uncertainties of $\pm 2\%$ (most bandpass systematics were found to be below the 1% level) would lead to a maximum bias on r of about 0.0009 (BICEP/Keck Collaboration et al. 2021). For simplicity's sake, I have kept that uncertainty for the current study on the impact of bandpass uncertainties on the next dataset.

6.3.1 Pseudosimulations

Under regular circumstances the bandpowers evaluated by the maximum likelihood (ML) search would either be real data, or simulated data with real noise properties. However, for BK24, which is expected to include data up to either the 2023 or 2024 season, we do not yet have the noise maps required to generate true simulations of our data. Therefore, we attempt to approximate the data with the expected bandpowers estimated from the fiducial signal model[15]. We attempt to approximate the noise with Gaussian random

[14] I will be referring to the BA maps as BA30, BA40, and BA150.

[15] The fiducial model used to generate both the BK18 and BK24 pseudosims has the parameters $r = 0.0$, $A_d = 3.6$, $\beta_d = 1.59$, $\alpha_d = -0.42$, $A_s = 0.0$, $\beta_s = -3.0$, $\alpha_s = -0.6$ and $\epsilon = 0.0$. The delensed BK24 pseudosims use $A_L = 0.337$.

noise, with correlations between bandpowers set by a predicted bandpower covariance matrix (BPCM). The predicted BPCM is projected from achieved performance based on our usual forecasting architecture (Buza 2019).

We generate these "pseudosims" by estimating the amplitude of the bandpowers with our model and estimating the noise with a Cholesky decomposition of the BPCM. The Cholesky decomposition gives us an estimated variance on the bandpowers in the BPCM. By multiplying that set of variances (M) by a set of Gaussian distributed random number (y, with μ=0 and σ=1) we can approximate Gaussian random noise on a set of model expectation value bandpowers we derive from the fiducial model. Therefore, a single pseudosim signal (s) is a linear combination of the expectation values and the Gaussian random noise like:

$$s = x + yM. \tag{6.10}$$

A full set of 499 pseudosims should adequately sample the uncertainty on the bandpowers; parameter constraints are then found by determining the parameter values that produce the maximum likelihood for a given set of bandpowers (one 'realization' of the pseudosims). The ML search on the pseudosims should return a set of parameters distributed around the initial parameters used to generate the initial set of model expectation values (x) if nothing is changed, which we term the baseline. We can then perturb the model in some way, such as shifting the bandcenters of the bandpasses used to produce the bandpowers by some percent, to evaluate how that changes the recovered parameters on a realization to realization basis.

The likelihood used for real bandpowers is an approximated Hamimeche-Lewis function, which has particular properties set by the physical universe (Hamimeche

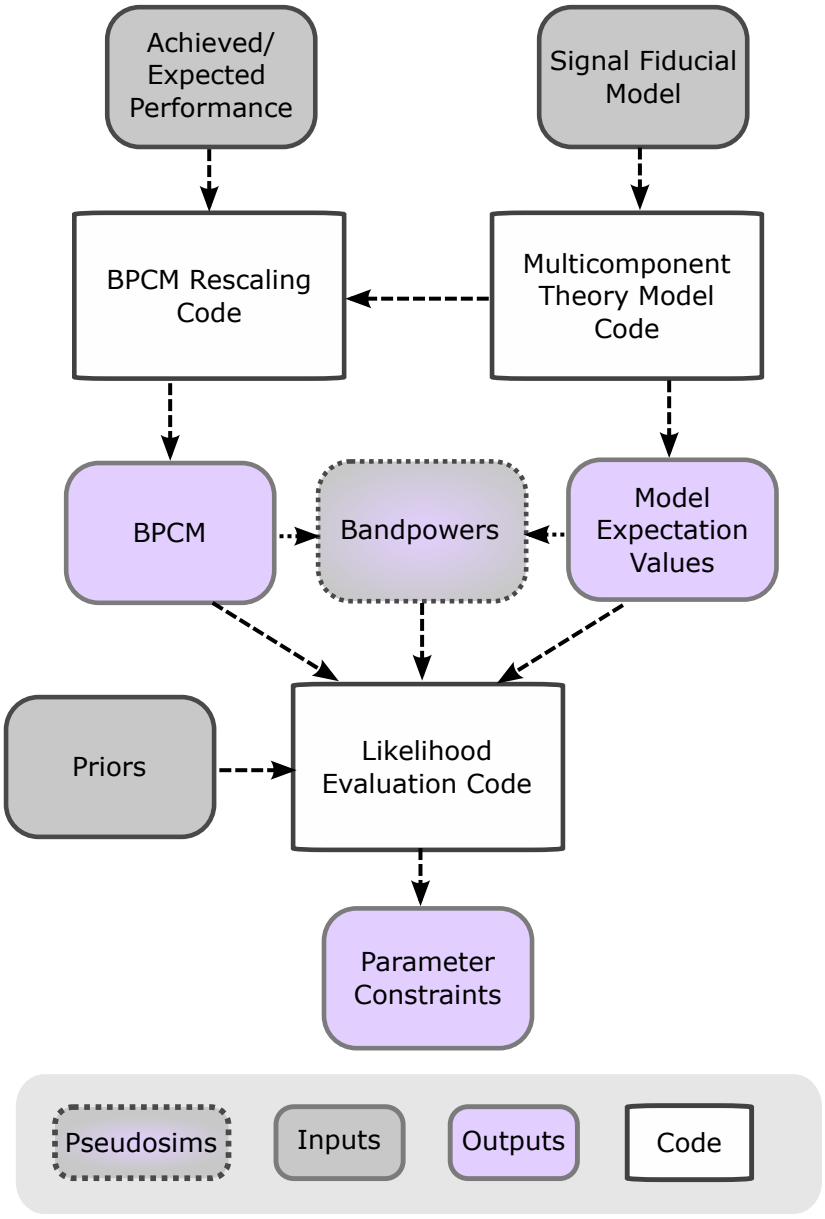


Figure 6.11: The multi-frequency BICEP/*Keck* maximum likelihood framework for multi-component parameter constraints. The pseudosims replace the inputs with the dotted outlines, which are the bandpowers used in the likelihood evaluation code. Schematic based on those in Buza (2019).

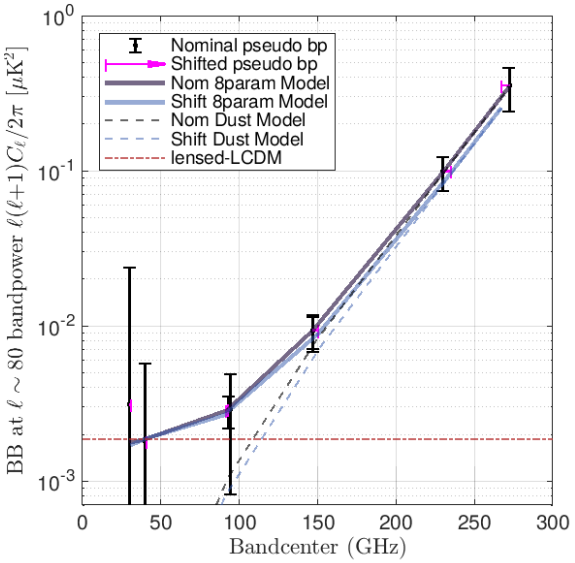


Figure 6.12: The mean BK24 pseudosim BB auto-spectra BICEP/*Keck* bandpowers plotted at the nominal (black points) and shifted bandcenter frequencies (magenta). The lines denote the expectation values derived from the model with the mean ML search parameters: solid lines are the full 8 parameter model expectation values, dashed lines are the portion of the model values associated with the dust, dash-dotted line are the portion of the model values associated with lensed-ΛCDM. The model fit to the nominal bandcenters is dark gray, and the model fit to the shifted bandcenters is blue. The realization to realization difference between the nominal and shifted bandcenter parameters are reported in the bottom of Figure 6.13.

& Lewis 2008). However, the pseudosims do not have the same distribution, as Gaussian random noise is applied to the expectation values; negative (i.e. non-physical) bandpowers are possible, particularly for low-ℓ or low signal frequencies. Therefore, the likelihood evaluated for the maximum likelihood search on the pseudosims is a χ^2, not the usual Hamimeche-Lewis.

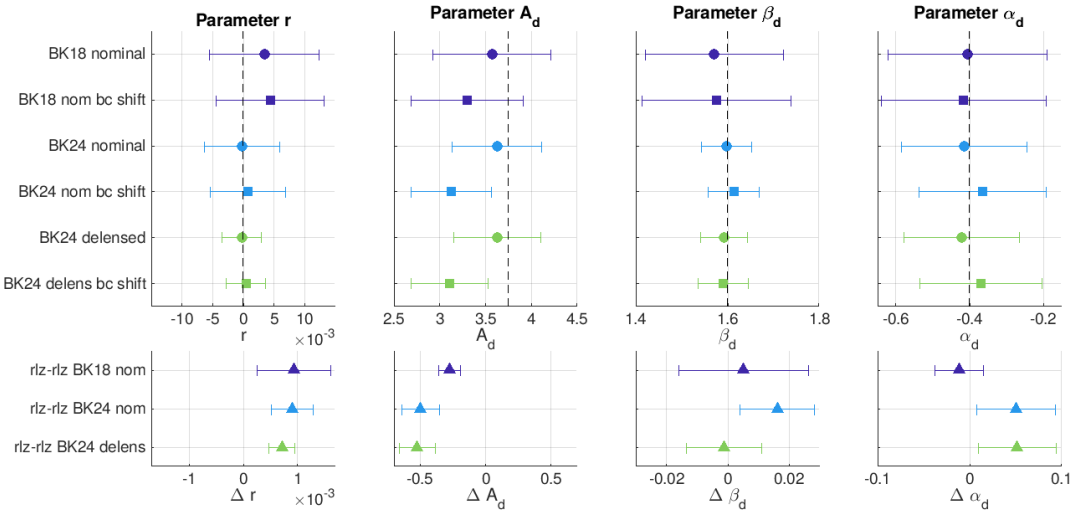


Figure 6.13: [Top] Parameter constraints for each of the pseudosim bandpowers—BK18 in purple, BK24 in blue, and delensed-BK24 in green—with and without bandcenter shifts (square and circles, respectively). Black dashed line denotes the expectation value for the parameter (and thus the basis from the fiducial model for the pseudosim). [Bottom] Realization to realization parameter biases from bandcenter shifts.

When evaluating a systematic such as bandpass uncertainties, we can either add the bandcenter as another free parameter to evaluate, or shift the bandcenter and reevaluate the likelihood with that shift. For this analysis, we have evaluated the likelihood with shifted bandcenters, and attempted to find the combination of bandcenter shifts that result in the largest bias on r compared to the baseline with no bandcenter shifts. The BK24 shifted bandcenters and the expectation values modeled from the ML search best fit parameters are shown in Figure 6.12 at one multipole. The full set of recovered parameters for all parameters for the delensed BK24 case (discussed further below) with and without bandcenter shifts, and the realization to realization bias on each parameter is shown in Figure 6.14.

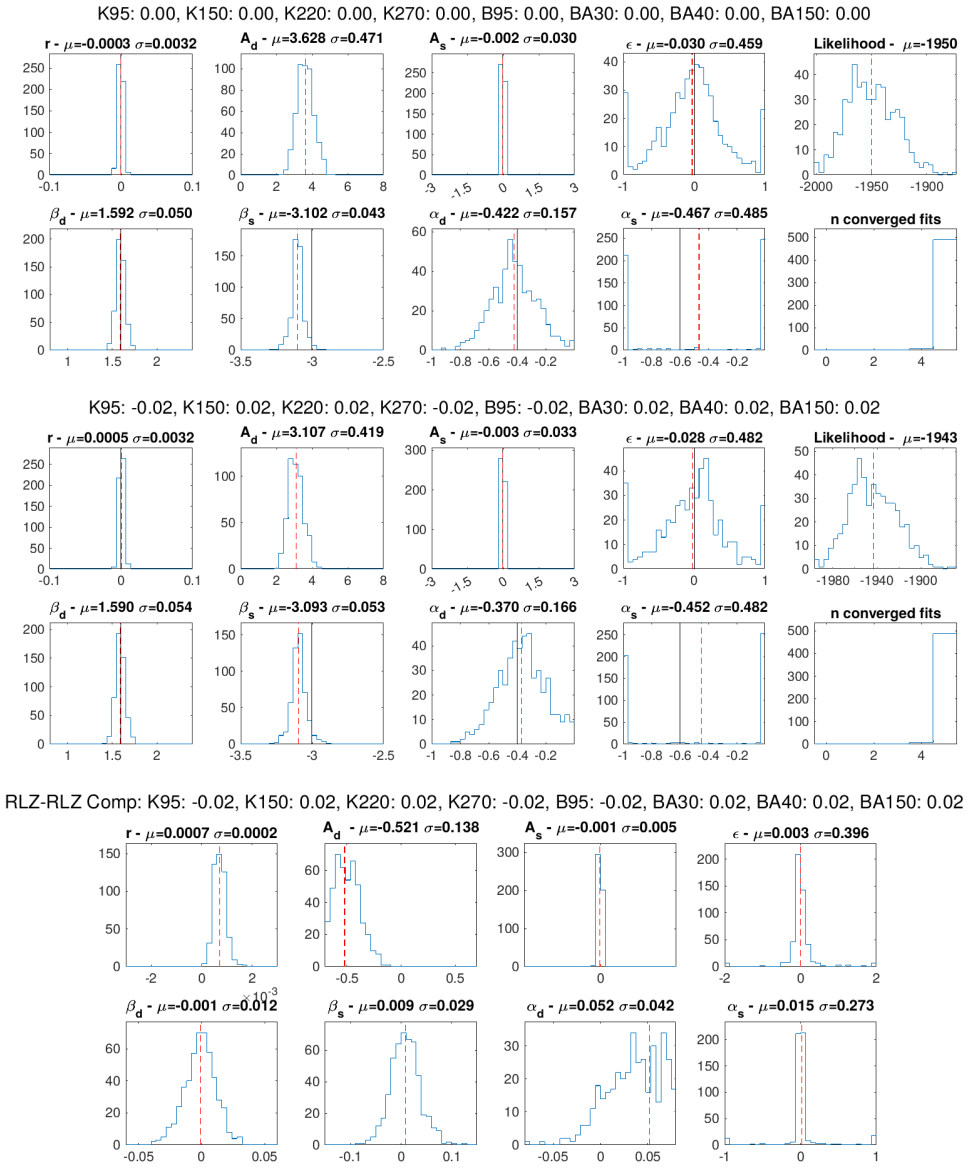


Figure 6.14: [Top] Maximum likelihood parameter constraints on the nominal delensed pseudosim bandpowers (no bandcenter shifts). [Middle] Maximum likelihood parameter constraints on the delensed pseudosim bandpowers with shifted bandcenters. Negative 2% shifts to K95, B95, and K270, and positive 2% shifts to BA30, BA40, K150, BA150, and K220. [Bottom] Realization to realization comparison of the above maximum likelihood parameters.

Table 6.1: Bias on r from 2% shifts on individual nominal (not delensed) BK24 bands.

	K95	K150	K220	K270	B95	BA30	BA40	BA150
Δr	-0.3e-4	0.3e-4	2.7e-4	-0.6e-4	-3.7e-4	0.04e-4	0.02e-4	3.5e-4

We first evaluated a set of BK18 pseudosims to recreate the previous worst bias on r and validate the method: the parameter constraints of two BK18 ML searches (one with and one without bandcenter shifts) is shown on the top of Figure 6.13 and the realization to realization comparison of those searches are shown on the bottom of the figure. The BK18 pseudosim parameter constraint on r showed a small consistent bias on r (not consistent with the parameter constraints from the ML search on regular simulation), but the realization to realization bias from bandcenter shifts was consistent with ML searches on regular simulations, reported in BICEP/Keck Collaboration et al. (2021). The bandcenter shifts that produced the largest bias on r were negative 2% to K95 and B95, and positive 2% shifts to K150 and K220.

In the BK24 dataset, we add additional low frequency bands (BA30 and BA40), the high frequency K270, and a 150 GHz band in BA150 to the previous set of four bands. The individual influence of BK24 bands on the bias on r is reported in Table 6.1. There are only four BICEP/*Keck* bands in BK24 that have a strong influence on the model parameters with bandcenter shifts applied: K220, K270, B95, and BA150. To produce a large Δr, B95 must be shifted in the opposite direction from the higher frequency bands. Since the B95 band is highly weighted, shifting the B95 band down while the high frequency bands (particularly K220 and BA150) are shifted up moves the bandpower associated with dust down enough that a small bias on r forms to compensate.

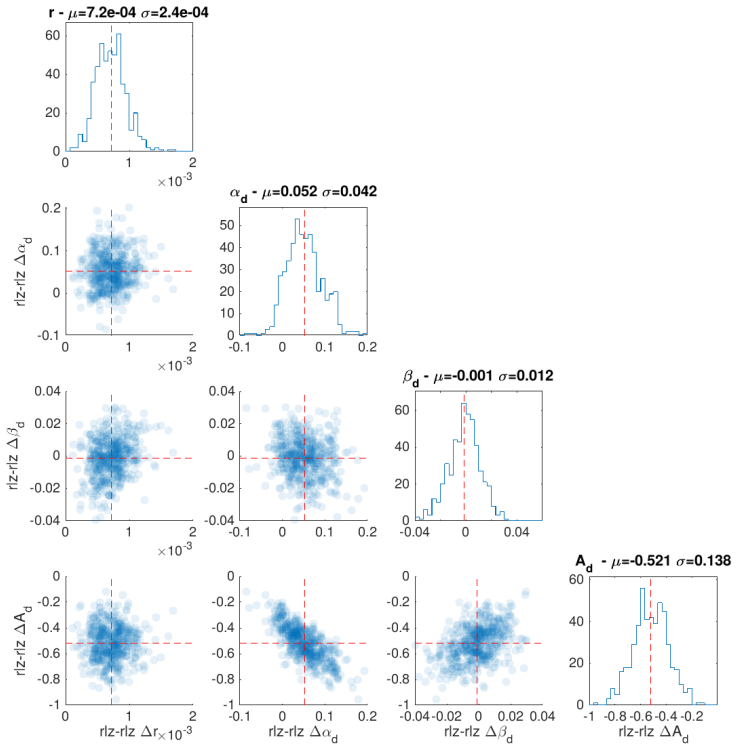


Figure 6.15: Corner plot of realization to realization shifts in maximum likelihood parameters r, A_d, β_d, and α_d with bandcenter shifts on delensed BK24 pseudosims. Maximum likelihood fit to delensed BK24 pseudosims with the same bandcenter shifts as those reported in the bottom of Figure 6.14, namely negative 2% shifts to K95, B95, and K270, and positive 2% shifts to BA30, BA40, K150, BA150, and K220. Red dashed lines denote the mean shift of the parameter.

Previously we were only concerned with the bias on r as a result of uncertain bandpasses as a systematic separate from other effects. Since BK18, however, we need to consider that we are running into the lensing foreground: the maps we use in the future are going to be delensed. We can generate a new set of pseudosims with the delensed BPCM and test the effect of bandpass shifts on that delensed set of bandpowers. I

assume that the worse case set of bandpass shifts from the nominal not-delensed case will produce the same largest shift in r in the delensed case: the bias on r reduces slightly in the delensed case from 9×10^{-4} to 7×10^{-4} (see Figure 6.13).

The parameters for synchrotron remain unbiased no matter which bandcenter shifts are applied: this is likely because the pseudosims were generated with the assumption that there is no synchrotron in our field (A_s was set to 0). It is possible if synchrotron is detected in the real data, that uncertainties at the low frequency will change the resultant parameters. The variance on the pseudosim bandpowers at the low frequency auto-spectra is very high. Even if synchrotron was detected, it is expected to drop off to be very low at 95 GHz, and likely will not bias r.

The dust parameters, by contrast, shift significantly with the shifted bandcenters, particularly A_d. Interestingly, a large shift in any of the dust parameters does not necessarily result in a large shift in r; dust parameters are only correlated with each other, not with the cosmological tensor to scalar ratio, as seen in Figure 6.15.

We can project when the bandcenter uncertainty induced bias on r (Δr) will become unacceptably close to the statistical uncertainty on r ($\sigma(r)$). The projected $\sigma(r)$ for the next decade of South Pole Observatory (combined BICEP/*Keck* and South Pole Telescope) observations is shown in Figure 6.16. The maximum 2% shifts in bandcenters resulted in a bias on r of 7×10^{-4}. We would prefer to keep a systematic effect to less than 25% of the statistical uncertainty; the conservative 2% bandcenter shifts therefore are a problem at $\sigma(r) \approx 2.8\times10^{-3}$, a statistical uncertainty we expect to reach in the first half of 2026. The bias on r is roughly linear with the magnitude of the shifted bandcenters. Therefore, 1% bandcenter shifts should result in a Δr of 3.5×10^{-4}, and

become problematic at $\sigma(r) \approx 1.4\times10^{-3}$, which we expect to reach in late 2029. To keep the bandcenter uncertainty induced bias on r less than 25% the statistical uncertainty until after 2035, the bandcenter uncertainty needs to be at or below 0.5%.

Bandpass uncertainty will become an issue for our constraint on r in short order—within the next year[16]. The constraint on band parameters is set by the measurement uncertainty of recent datasets: a better constraint on the FTS systematics uncertainty on BA receivers will hopefully allow for a tighter bandcenter uncertainty, and a lower resultant bias on r.

[16]This is assuming that we want to keep the bias to less than 25% of the statistical uncertainty. If we want to keep the bias to less than 20% of the statistical uncertainty, the uncertainty on bandcenter is already a problem for the data we are taking *now*, in 2025.

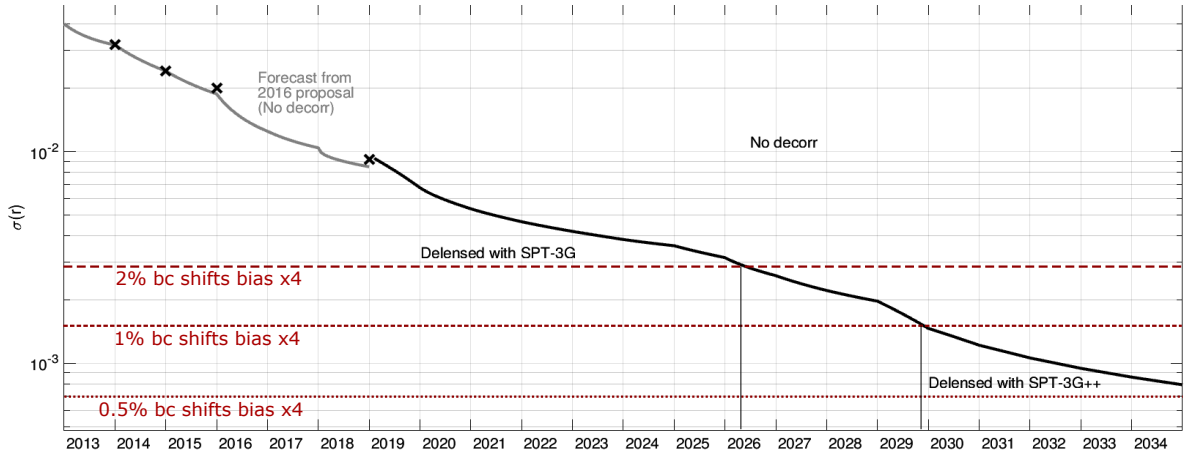


Figure 6.16: The projected statistical uncertainty on r ($\sigma(r)$) for the BICEP/*Keck* experiment—with delensing at high resolution provided by South Pole Telescope—until 2035 (solid black line) to be compared with the expected bias on r (Δr) from bandcenter (bc) uncertainties at the 2% (dashed line), 1% (short dashed line) and 0.5% (dotted line) level. The bias from bandcenter uncertainty is multiplied by a factor of four, as we would prefer to keep a systematic bias to less than 25% of the statistical uncertainty. The current conservative 2% bandcenter shifts resultant Δr reaches 25% of $\sigma(r)$ in the first half of 2026, the expected 1% bc shifts Δr reaches 25% of $\sigma(r)$ in late 2029, and the expected 0.5% Δr reaches 25% of $\sigma(r)$ after 2035. Forecasting of $\sigma(r)$ done by Kenny Lau on behalf of the BICEP/*Keck* Collaboration.

Chapter 7

Conclusions, Next Steps

7.1 Laminate Windows

The laminate window project accounted for the majority of the work in this thesis. I have retired many of the risks that we initially identified when starting the project in earnest, plus a few more. These initially identified problems were as follows:

1. Mechanical issues due to poor lamination between layers seemingly resulting in a slow leaks during creep tests (covered in §2.2 and §2.5);

2. The safety factor of windows operating at the South Pole (covered in §2.2, §2.3, §2.4, and §2.6.);

3. Uncertain optical properties, particularly the refractive indices reported in Barkats et al. (2018) (covered in §4.1.1), loss tangent/transmission (covered in §5.3), scattered power (covered in §4.2.1), and polarization (covered in §4.1.2);

4. The design of anti-reflection coatings, particularly for higher frequency receivers (covered in §3.3 and §3.4).

There still remains a handful of open questions about the windows. Let us discuss these remaining questions, and the potential next steps to answer them.

Thicknesses

While we have done some studies to predict the final thickness of the windows, it is currently a fully empirical estimate. We have fit linear relationships to the combined 'starting' and final thicknesses, which fit well (see Figure 7.1), and we find that the linear slope is close to unity. There is some amount of variation in the thicknesses of windows with nominally the same starting thicknesses: I suspect that this is due to minor variations in LDPE film thicknesses, and potentially processing differences (if pressure is not held through the lamination, for instance, the LDPE may not fully penetrate the HMPE).

From the fit linear relations there is a bulk negative offset of the final thicknesses; I suspect that this offset is related to the amount of LDPE that gets wicked into the HMPE. Therefore, a certain amount of the LDPE thickness will 'disappear' into an HMPE layer. Different number HMPE layer laminates will have a different-but-predictable offset in the thickness relationship, which is backed up anecdotally by one-off laminates I have made. This offset seems to be around 3.5 mil per HMPE layer, which makes some sense if about half of the volume of the HMPE needs to be filled.

The fully melted LDPE will also viscously flow outwards during lamination: this results in a 'skirt' of LDPE around the laminate. The volume lost to the skirt can reduce

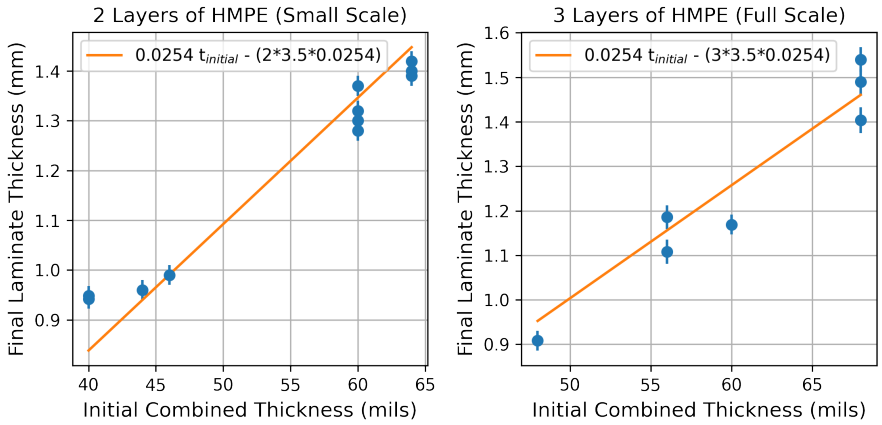


Figure 7.1: Initial combined thicknesses in mils versus the final measured laminate thicknesses of two layer samples (left) and three layer full scale windows (right). We typically assume that the HMPE layers start at 8 mil, so the total combined thicknesses is the summed initial LDPE layers plus the HMPE layers.

the final thickness of the laminate. This is a worse effect for smaller laminates, because a larger fraction of the volume flows out over the edge[1]. This becomes a worse issue if you fully melt the LDPE during AR lamination, but it's possible to mitigate the flow by not fully melting the LDPE.

The volume of LDPE wicked into the HMPE sets the likely lower limit of laminate window thickness. If we attempt to make a window without enough LDPE, it is likely that the HMPE fibers will not be fully laminated and therefore the window will not be vacuum tight. However, there is interest in the community for determining the lower limit of the window thickness, particularly for small aperture wide-bandwidth windows

[1]Recall that the circumference of a circle scales with the radius but the volume of a cylinder scales with radius squared. Therefore, if the LDPE can flow the same distance from the edge when melted, the skirt will be a larger fraction of the volume at smaller radii.

on millimeter interferometers.

If I had to guess, given the linear empirical thickness relationships, the absolute minimum amount of LDPE per layer of HMPE is somewhere around 3.5 mil for each HMPE layer[2], but this has not been tested. The lower limit I have been willing to test is two 6 mil layers sandwiching a single layer laminate: that resulted in the high pressure single layer sample used for the birefringence study (§4.1.2), which had a final thickness of about 0.4 mm. It may be worth attempting to find the actual lower limit of laminate window thickness with a dedicated study.

Failure Strength

The hydrostatic test described in §2.3 adequately determined a lower limit safety factor of a full scale window, but did not find the actual failure point. I highly doubt we will ever push a deployed window above the 85 psi limit we found, but it would also be useful to determine what, exactly, the limit is.

To push the window above 85 psi will require a new apparatus, as the previous set up has obviously reached its operational and safety limit. One quick way to get the apparatus to higher pressure would to either use a thicker aluminum backing plate or better support the existing plate. I suspect that the leaks at the edge are due to the

[2]I would also expect that the laminating LDPE layers have to sandwich the HMPE to surround it. I do not think putting 3.5 mil of LDPE on one side of an HMPE layer would fully laminate the HMPE because the LDPE would not flow through the HMPE to the other side. It would be like putting a pile of spaghetti on chicken-wire mesh; some strands of spaghetti may make it through the mesh, but a lot would get caught above the mesh and stay there.

frame and plate flexing away from each other because the extra mechanical support from the C-clamps helped push the final achieved pressure higher (the C-clamps helped push the frame and plate together). Plus, the safety factor for the half inch aluminum plate was only 3 at 4 atm; too much more force will bring the safety factor below safe limits.

Any new test will likely need additional safety precautions; 85 psi over the large surface area of the window is a lot of force (the equivalent weight of something like 18 Honda Civics), and when the window fails it is likely to do so spectacularly. It may be best to fully submerge the entire apparatus in water to adsorb the energy of the rupture. At the very least, some kind of blast shield would not go amiss.

Absorption Loss

While I have determined roughly what the absorption loss of the full thin window system is from optical tests (likely somewhere around a tanδ of 2.4×10^{-4}, as discussed in §5.3 and shown in Figures 2.1 and 5.4), we do not yet have a good direct measure of the loss of the laminate window alone. There has been some progress within the Kovac group on loss measurements of polymers with a quasi-optical cavity, but laminate windows have not yet had good constraints placed on their loss properties in that device (Elwood et al. 2024). I suspect—given that HDPE and UHMWPE have been shown to have very similar loss—that the HMPE laminates themselves likely have similar loss tangents to other polyethylene grades, and that the 'high' loss estimates sourced from the material insertion load/NET measurements is biased upwards by the relatively thick ePTFE AR coats. However, I do not currently have a well founded estimate of the ePTFE loss properties; Lamb (1996) puts the loss of bulk PTFE somewhere around 3×10^{-4} at 150

GHz, but how loss scales with density is an open question.

High Frequency Instrument Noise

Measuring and understanding the instrument noise at high frequency is a project in and of itself; BA150 in particular has had significantly higher measured noise than modeled noise since its deployment (with a thin window) due to a suspected blue leak (though, as discussed in §6.2.2, some of the excess optical load may be from coupling with the 118 GHz oxygen absorption line). Another low pass filter was added in the austral summer of 2024/2025, but the effect of the filter on CMB timestream derived NET is yet to be determined. Further breaking down the individual component contributions, including the thin window, will hopefully be conducted this season.

BA220/270 has also deployed with a thin window in the austral summer of 2024/2025, but as of the writing of this thesis the winter observing campaign has only just begun. Hopefully the optical load from the window is proven to be as low as expected, but only time will tell.

Updating our BICEP Array NET model with better estimates of the low temperature loss properties of the optics is the natural next step in modeling our instrumental noise performance; there is active work in the group to take the previously mentioned quasi-optical cavity to the temperatures our optics are held, and to conduct numerous other studies on dielectric losses. This information will be invaluable in predicting instrument performance of current and next generation CMB receivers.

Other improvements to the NET modeling may include using a more accurate estimate of the detector bandpass for optical load estimates or a fully differential study

of the loss tangents' uncertainty impact on the noise.

7.2 Anti-reflection Coats

The anti-reflection coat design process was a significant side project, and the focus of the longest chapter in this thesis (Chapter 3). The solutions I have determined for the current aperture sizes are relatively robust and close to ideal at their fractional bandwidth (particularly for the BICEP3 thin window and the BA150 or BA220/270 HDPE AR coats). I think that we may be able to further improve our AR coat design and production techniques to further reduce the band average reflections.

There is a balancing act for thin windows between additional anti-reflection layers and increased transmission loss from the thicker optical system. Our typical design process is to reduce the band average reflected power as much as possible, with a requirement to keep the reflected power below 1% (-20 dB) across the band. The absorbed power lost through a thin AR layer should be very small compared to the power lost through reflection off even a decent AR coat. But the added photon noise is difficult to predict on a general basis, as the optical throughput to a detector is instrument specific. I believe that thin windows are entering a regime where we must consider carefully what the dominant sources of transmission loss are.

Dual Band 95/150 GHz AR Coats

The largest fractional bandwidth covered in Figure 3.4 is the proposed CMB-S4 dual band mid-frequency (95/150 GHz) receiver ($\frac{\Delta\nu}{\nu} = 0.7$). This very wide bandwidth poses

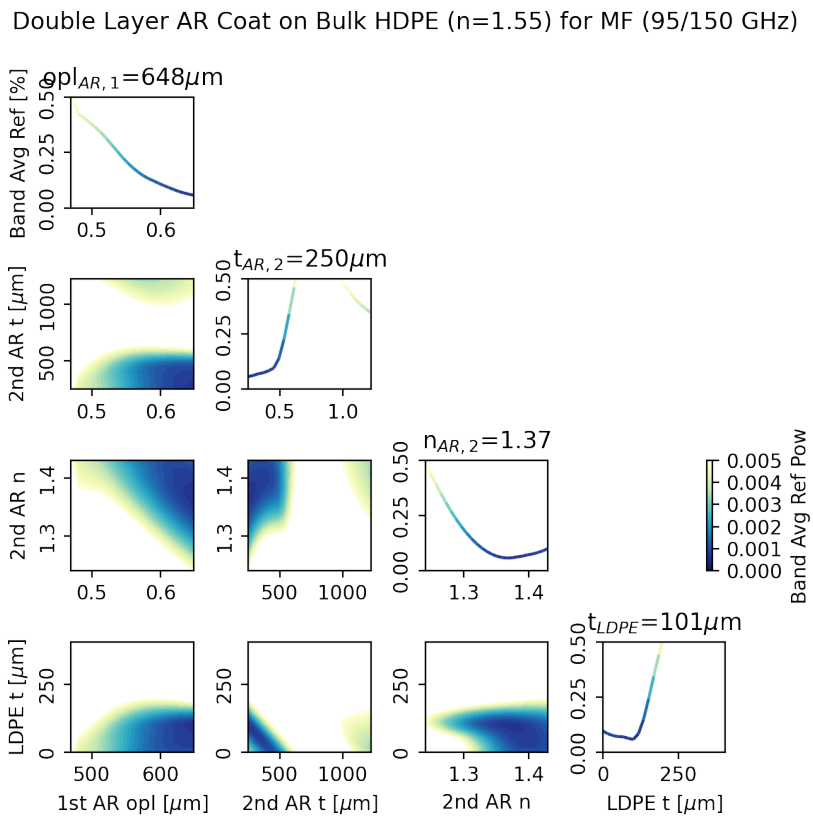


Figure 7.2: Design corner plot for a hypothetical dual band 95/150 GHz receiver like those proposed for CMB-S4 or PreSAT (Abazajian et al. 2019; Petroff et al. 2024). The design assumes SH Teadit for the outer layer (AR,1) and some sintered PTFE membrane for the inner layer (AR,2).

a significant design challenge to AR coat; however the bands are separated to avoid the oxygen absorption line between them, so a double layer AR coat may be designed to center a null in each band. This results in the proposed AR coat for HDPE shown in Figure 7.2. The same SH Teadit used for the BA150 AR coats acts as the outer layer, and some form of sintered, relatively high density (potentially even bulk PTFE, using appropriate caution) PTFE membrane can be used for the inner layer. The nylon AR coat would likely be similar to the coatings used in BA220 (outer layer low density

SH Teadit, inner LDPE acting as both adhesion and an AR layer), but instead with a different thickness of LDPE.

Alumina AR Coats

Layered alumina AR coats are particularly difficult due to both the material's high index of refraction (necessitating high index AR layers) and low coefficient of thermal expansion (making it difficult to adhere layers). We currently use sets of epoxy layers with different amounts of fill material to tune the index of refraction, cured between two rigid plates. These epoxy sheets must be laser cut in a grid once adhered to the alumina surface such that the squares can contract separately (Dierickx et al. 2023). These epoxy sheets are very thin, particularly for high frequency, and we regularly find squares off epoxy have delaminated during cool-downs.

Purely speculatively, I suspect that we might be able to make an alumina AR coat with a pseudo-double layer design like those used for the BA220 window, with relatively high index sintered PTFE membrane and adhesion layers. The ideal outer layer index for a double layer alumina AR coat has an index around 1.33 to 1.39 (see Table D.5), which is within the achievable range of a sintered billet[3]. The intermediate layer is significantly more difficult to generate with low index materials and is made even harder by the addition of the required adhesion layers, which add to the electrical length. Potentially a relatively high index outer layer (1.34), a very low index inner layer (1.15), and thick LDPE adhesion layers would produce a decent anti-reflection coat; the thick LDPE layers require a sufficiently different impedance layer between to provide anti-reflection.

[3]Higher density sPTFE is easier to skive, and may be possible to be cut thinner.

The coefficient of thermal expansion mismatch would be the biggest issue with such an AR coat. I would hope that a sufficient amount of LDPE would be able to deform under the large amount of shear stress, but LDPE will become brittle below the glass transition temperature, which should be around 200 K. I also would not recommend laser cutting stress relieving features into a PTFE based AR coat, as the volatiles off melted PTFE are known to be severely damaging to human health (National Academies of Sciences, Engineering, and Medicine 2022; Zahm et al. 2024). It may be possible to do an epoxy layer, laser cut that layer, then adhere another PTFE outer layer, but I am uncertain if that would make the thermal contraction issues better or worse[4].

Other Uses of DeWAL

DeWAL, as a low density low thickness membrane, could be used for other optics. Physically separated sheets of uncompressed DeWAL could potentially be used as a Radio Transparent Multi-layer Insulation (RT-MLI) (Choi et al. 2013). DeWAL's very small pore size (likely similar to other ePTFEs at around 12 μm) could eliminate the scattering concerns that are raised by the approximately 400 μm cell sizes that the current BA RT-MLI material (Goldfinger et al. 2022). The number of layers of physically separated DeWAL that would be required to be an appropriately efficient IR filter is an open question.

A many layered (on order tens of layers) stack of DeWAL and LDPE could be used to generate a low pass edge filter. Developing this filter would likely require an

[4]Pro: the epoxy layer can contract separately with the stress relieving laser cut grid. Con: the LDPE adhesion layer on top might enter the laser cut gaps and cause issues.

easier procedure to lay DeWAL flat for these many layers such that it is not egregiously time-consuming to generate.

7.3 Spectra Analysis

The FTS project is the one that I have sat with the least, but I have some thoughts on which steps should be the next ones taken. The FTS is a important piece of calibration equipment, which takes a vital measurement of our receivers' response to the sky, and we can leverage it for high precision measurements.

Translation of the FTS

The fact that the FTS has not been translated above the BICEP Array receivers to couple with the center of the Gaussian beams is a major problem: I hope that I have corrected the resultant complex frequency dependent aperture truncation that occurs, but it would be far simpler to just translate the FTS during a campaign to prevent the issue in the first place. I suspect that the reason this has not occurred recently is two-fold; the FTS is currently difficult to move accurately above the BICEP Array receiver and there has been a miscommunication between the instrumentalists who operate the FTS and those who analyze the FTS data.

As such, I believe that the next major step that should be taken with the FTS is to overhaul its documentation and measurement procedure. If we can make translating the FTS physically easier, it will incentivize the proper data taking, and significantly reduce the difficulty in determining the aperture correction factor on BA data. Additionally,

simpler and more reliable alignment procedure will hopefully ease the complexity in taking FTS measurements, which primarily boils down to it being difficult to translate and point to detectors in the receiver.

However, as constraints on summer deployments do not look like they will ease in the coming years, I think it wise to explicitly state a decision tree for future FTS measurements. These decisions assume that we will be translating the FTS for future measurements. The higher on the list, the more valuable the measurement is:

1. Both polarizations
    - If time is limited, prioritize 45 deg rotated measurements over 0/90 aligned measurements, as that seems to generate fine data quality and gets *both polarizations at the same time*.
2. Multiple pointings
    - Getting multiple measurements pointed at different positions on a tile will hopefully better couple with those different detector pointings, and help constrain systematics for a given measurement campaign.
3. Multiple interferograms
    - Getting many halfscans to average together does help reduce the noise of a single measurement, which is useful if we can only get one, but also increases the time to take data considerably. If time is limited, consider, for instance, only doing 4 halfscans instead of 7.
4. Multiple rotations

- It is nice to get both a 45 deg and a 0/90 deg rotated measurement on a detector to probe rotation systematics.

5. Other systematics (mirror speed, etc)

- These are nice to get, but not required.

Other Spectra

We have done FTS campaigns on the other BA receivers, and some work has gone into pushing those data through the first pass reduction. I prioritized the band that would be in the next dataset and has high survey weight (BA150), but the other BA bands should be scrutinized more closely.

BA30/40, being the oldest receiver, has the most work put into its spectra: the issue with analyzing and comparing between those FTS campaigns is the many iterations of detector modules that have moved through that receiver in the last five years. Year to year comparison will mostly be a bureaucratic exercise for the BA30/40 modules, though the low detector counts will hopefully not make the issue too difficult. I did not prioritize the low frequency bands as work had already been done on them by Zhang (2023) and they were expected to have minimal impact on cosmological parameter recovery, which was confirmed with the pseudosim ML searches.

BA220/270 was only deployed this last austral summer and therefore it will likely not be incorporated into the next data release. I will be working on some systematics measurements taken recently with BA220/270 to hopefully evaluate the potential systematic impact from the stationary FTS at high frequency. An FTS systematics

campaign was conducted at Stanford in Fall 2024 and at South Pole in the austral 2024/2025 season on BA220/270 and may help illuminate the issues with the BA150 2025 campaign: if the FTS was misaligned, for instance, it likely would be misaligned for both receivers.

Individual Band Contribution to Bias on r

Some bands, having more survey weight, have a larger impact on the ML search parameter recovery. In general, it seems, that B95 and BA150 have the largest single band influence on the bias on r. In Figure 7.3, the realization to realization compared fit parameters are shown for shifting BA150 and K95 by 2%: the BA150 bandcenter shift produces a significant bias on r (and the dust parameters) while K95 does not shift any of the parameters significantly.

We currently are including quite a few maps in BK24, which results in a large amount of stored data in the form of data products like observation matrixes, simulation sets, and maps. It is possible that we may remove maps from the dataset in the future if they are not contributing significantly to the survey, and the lower frequency *Keck* bands (K95 and K150) seem to have dropped in weight such that they may be removable.

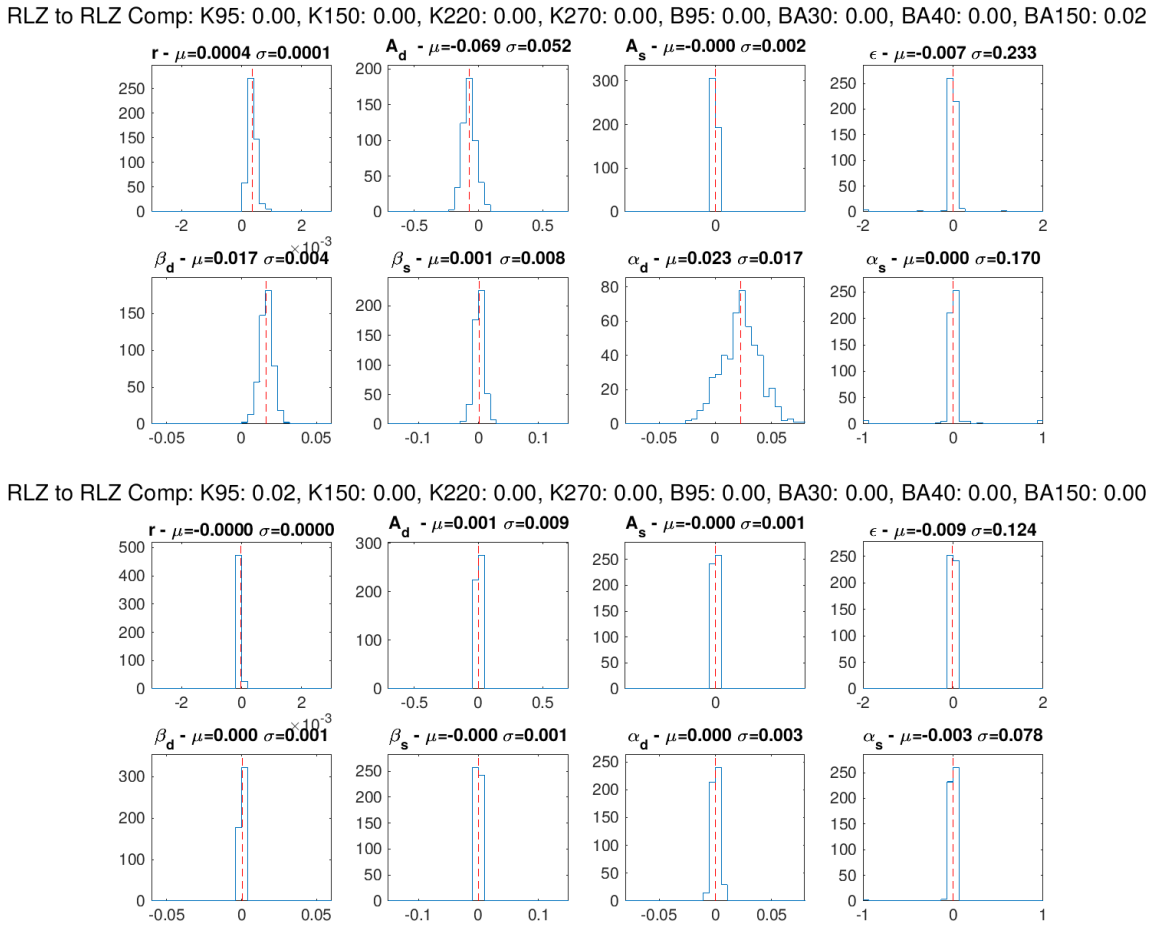


Figure 7.3: Realization to realization comparison of maximum likelihood parameters fit to pseudosims with and without shifted bandcenters. [Top] BK24 pseudosims maximum likelihood search with single band (BA150) shifted by 2%. [Bottom] BK24 pseudosims maximum likelihood search with single band (K95) shifted by 2%.

Appendix A

How to Build a Laminate Window

Making the windows is a many step process, and requires some specialty equipment.

Another important note: we currently have two rolls of HMPE weave which is 90" wide and 70 meters long. As we slowly approach the end of the first roll, I have noticed artifacts into the weave; they appear to be fibers of some other kind that have gotten caught in the weave (see Figure A.1). These artifacts are difficult to spot, being white on white, but I suspect that they were the source of the spot in the prototype window used for the hydrostatic test in §2.3 (visible as a red circled white spot in Figure 2.5), which was measured to have a different index of refraction compared to the rest of the laminate, hence why it was not made the science grade window.

I have had some success removing some artifacts in the weave with tweezers and careful fiddling, but some are so knotted that they are impossible to disentangle. In those cases, I have had to cut around the artifacts, which can be very difficult (and wasteful) depending on the placement of the artifacts.

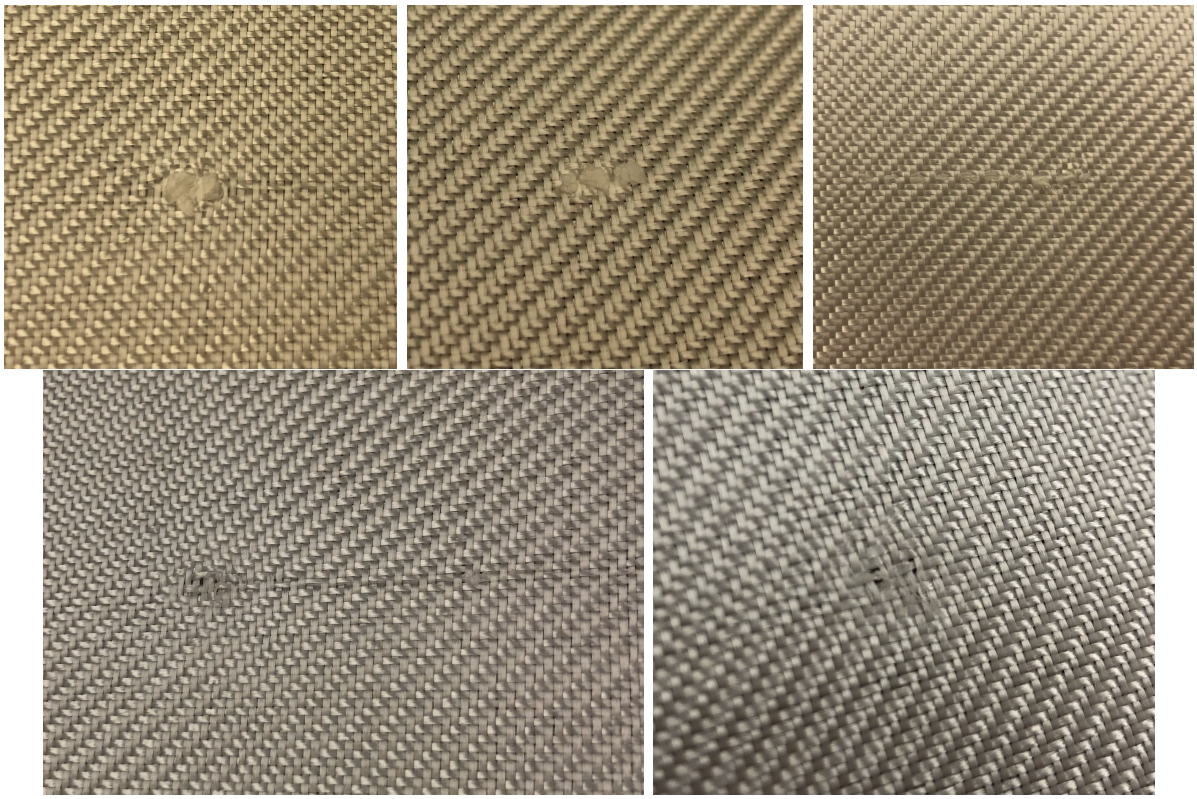


Figure A.1: Examples of 'artifacts' woven into the HMPE weave. The lower two pictures are the weave after I attempted to remove the inclusions. In general, the 'spongier' artifacts (top left and top center) are easier to remove, while the more fibrous (top right) artifacts are more difficult. I suspect that the 'spongy' artifacts are a different material, while the 'fibrous' artifacts are tangled HMPE.

A.1 Designing the Window

The designed window thickness is typically dependent on the designed anti-reflection; the thin windows, as discussed in Chapter 3, can act like another anti-reflection layer, and the placement of the null is important! The measured index of refraction (discussed in §4.1.1) of the laminate windows are typically constrained between 1.537 to 1.557, while the final thickness is dependent on the number of LDPE and HMPE layers in the laminate (see the discussion in 7.1, and specifically Figure 7.1). The final thickness varies with the number of HMPE layers, as LDPE is wicked into the fibers during lamination, which is generally predictable by a linear relationship with a slope of one and an offset (of about 3.5 mil) multiplied by the number of HMPE layers [1].

The number of HMPE layers is set by the desired safety factor and, to some extent, the creep of the window. I have some evidence that the creep will change depending on the number of HMPE layers (see Figure A.2), but this is from a quick one off study comparing a two layer, three layer and four layer HMPE laminate with the same designed thickness.

[1] So, for example, a three layer HMPE laminate with desired thickness of 1.2 mm (47.2 mils) needs an initial thickness of $47.2 + 3(3.5) = 57.7 \approx 58$ mils, 24 mils of which (3×8) is HMPE. Therefore, we need about 34 mils of LDPE, which can be accomplished with 5 layers of 6 mil and 1 layer of 4 mil LDPE. The laminate stack for such a window might be 66D6D6D64 (the 'D' denotes where in the stack the HMPE or 'Dyneema' layers are).

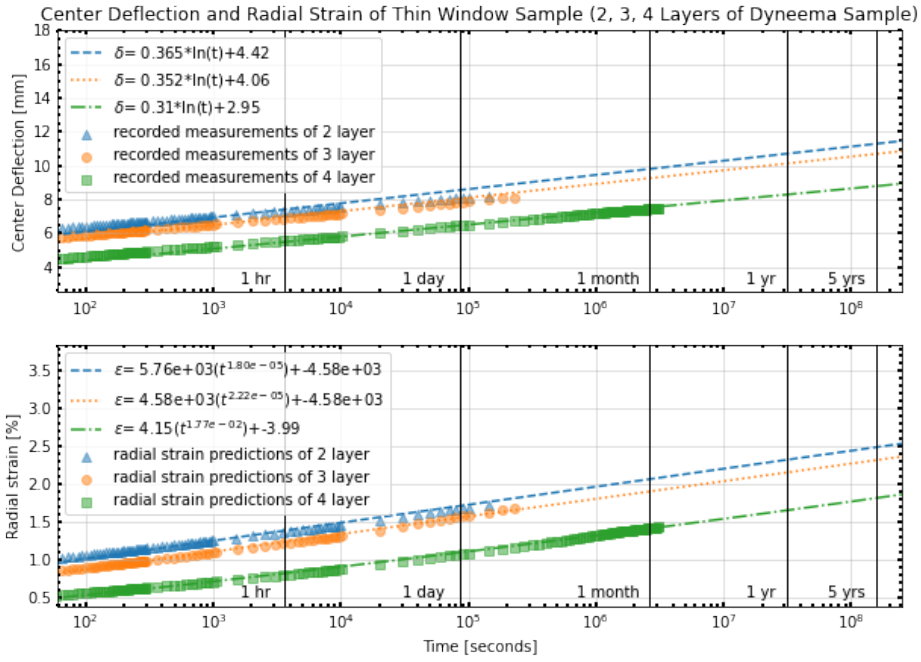


Figure A.2: Small scale creep of laminates with three different layers of HMPE; blue triangles are two layer (0/45 orientation) laminate, orange circles are three layer (0/22.5/45 orientation) laminate, and green squares are four layer (0/45/0/45 orientation).

A.2 Cutting the HMPE (or anything else) on the Rotating Template

We use the same template to cut full scale window materials, such as LDPE sheets, the HMPE layers, and the release layers. The template is also used to lay-up DeWAL (see Appendix C), and to cut AR layers to shape (the inner diameter of the annulus is the same diameter of the o-ring in the window frame and clamp, and therefore is useful as a guide). The template jig is made up of four primary components:

1. The template; a circular sheet of aluminum that has a diameter about 0.5" wider than the final window diameter in the window frame and clamp.

2. The turntable; a lazy susan that goes below the template to rotate it.

3. The annulus; an aluminum annulus with handles that acts both as the clamping mechanism between a material and the reference for AR coat diameters for the windows.

4. The magnets; a set of a few hundred small (0.5" diameter) magnets used to clamp materials between the template and the annulus.

I will note when some step is specific to a material with *For* {*material*}. Most often, there are specific steps related to the HMPE cutting, which requires specific hardware and considerations.

A.2.1 Using the Template for Window Materials

1. Set up the aluminum circular template with the good side (no scratches) up on the turntable, such that it rotates about the center (marked with a small black 'x'). *For HMPE*, make sure the template is close to one side or the other of the HMPE roll (as the roll is much wider than the template you should alternate left and right to save as much material as possible).

2. Place magnets on the table below the template a few centimeters in from the edge and about ten centimeters apart (see Figure A.3). Note which pole you are facing upwards.

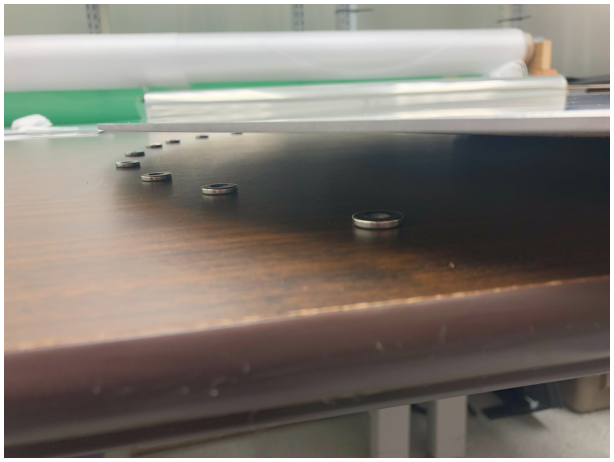


Figure A.3: Photo of the magnets placed below the edge of the template on the table.

3. Unroll the material to be cut (in Figure A.4, HMPE) over the template. *For HMPE*, give a few centimeters of margin around the edge of the template.

4. *For HMPE*, roughly cut off any trailing excess that hangs off the edge of the table to prevent the draped fabric from distorting the fabric over the template.

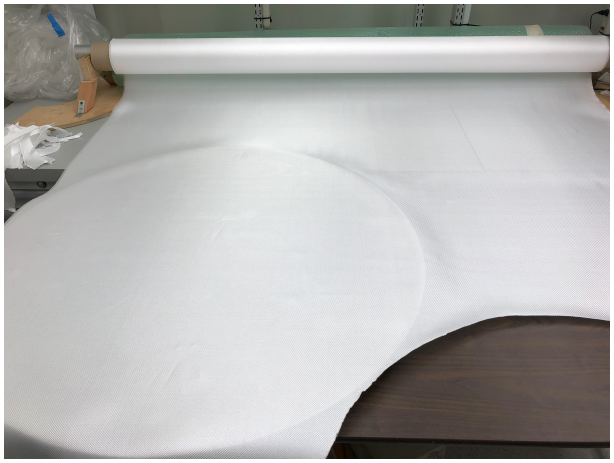


Figure A.4: Photo of HMPE sheet over the template.

5. Inspect the material over the template for defects, particularly *for HMPE* for the artifacts included in the weave (examples shown in Figure A.1). Attempt to remove

any artifacts atop the template surface, and maneuver the template to a different portion of the fabric if the artifacts are not removable.

6. Place the annulus over the template, taking care to align the outer circumferences. *For HMPE*, inspect the weave for distortions; the fabric has a tendency to stretch and warp when a long portion has been draped over the edge of the table for an extended period of time. You can nudge it straighter by hand.

7. Place magnets above the annulus, with the same pole facing upwards as the ones placed below the template. For the first magnet, you may have to move it around a bit and press firmly to get the two magnets to get close enough to attract (see Figure A.5). Work in a circle around the annulus until there is an even ring of magnets holding the annulus, material, and template together. If the template is not fully over the table, add two magnets to hold the annulus down over the portion that you could not put magnets below.

8. Inspect the annulus; it is not perfectly flat, and there may be portions of the ring that are not held against the material. You can slide a greater density of magnets to the elevated portions to push the ring down.

9. *For HMPE*, rough cut with the hot wire cutter a few inches away from the template. The best attachment for this is the stiff wire, with a flattened, roughly straight portion that can be evenly placed over the HMPE, cutting a few inches at a time. The HMPE sheet after a rough cut is shown in Figure A.6

10. Close cut the material.

    (a) *For HMPE*, use the wooden wire cutter, with the power source at 5 amps.



Figure A.5: Photo of placing magnet on the annulus over a sheet of polyester (release material).

Move the wire cutter up and down against the template in a sawing motion and slowly spin the template to cut off all the excess material.

(b) *For LDPE*, use scissors and cut to the template, taking care to not create jagged edges. You should be able to glide the scissors along the edge smoothly once the cut is started.

(c) *For release layers (or any material that needs a diameter larger than the window)*, cut with scissors with a set of small magnets the length of the desired increase in radius off the window outer radius. The magnets act as a guide to the extra radius off the template outer diameter. See Figure A.7. Take care to carefully cut slowly, so as to not create a jagged edge.

11. Carefully remove the magnets atop the annulus (if the magnets are not evenly

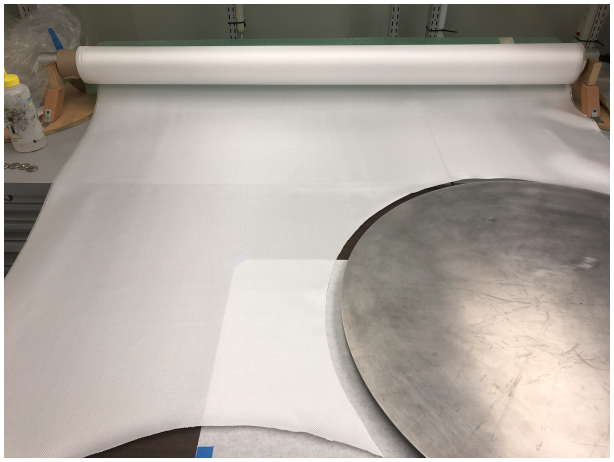


Figure A.6: Photo of the HMPE roll around the template after a rough cut.

spread on the annulus, evenly space them before removing the top), allowing the magnets below to fall to the table. Catch any that would fall to the floor from the portion of the template not above the table.

12. Remove the annulus and the material from atop the template. Store the material appropriately.

A.2.2 Using the Template for Cutting AR Layer

Cutting AR coats is a slightly different process, but very similar.

1. Set up the aluminum circular template with the bad side (circular scratches) up on the turntable, such that it rotates about the center (marked with a small black 'x'). Make sure the surface is clean.

2. Place the AR material on the template, centering it as much as possible on the template.

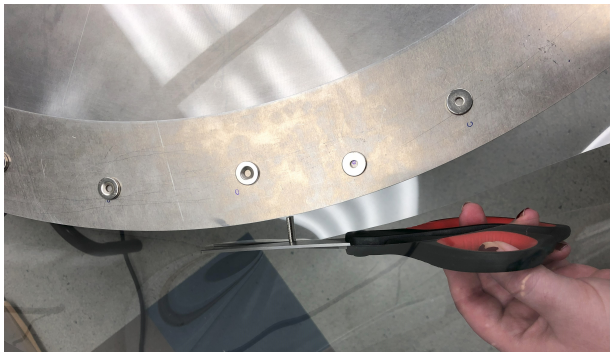


Figure A.7: Photo of cutting polyester (release material) with scissors and a small magnet stack to a desired increase in radius off of the window outer radius. Particularly useful for release layer cutting.

3. Place the annulus over the AR layer (portions of the AR layer may not be under the annulus, and that's okay), and align the outer diameters.

4. Again, check the centering of the AR layer; there should be the same distance between the edge to inner annulus on all sides.

5. Place magnets above/below the annulus/template to hold them together. You don't need as many as for the other cutting procedures, as you just need to hold the annulus down.

6. Mark and/or cut the desired diameter of the AR layer.

   (a) If marking (which I've done for AR layers that have radii at the inner radius of the annulus), take a marker to the AR coat at the inner radius of the annulus. Then remove the annulus and cut at the mark carefully with scissors.

   (b) If cutting (which I've done for AR layers that have a radii much further in than the inner radius of the annulus), take a retractable razorblade and add a length of small magnets with the desired reduction of radius to the AR layer



Figure A.8: Photo of a close cut layer of HMPE.

(see Figure A.9). Then carefully and evenly run the razor blade through the AR layer by spinning the template below the razor blade[2], using the length of magnets to keep the radius consistent.

7. Remove the annulus if not done so far, and store the cut AR layer appropriately.

A.3 Making the Laminate Stack

After determining the window recipe, and cutting all required materials for the window, the final step is putting it all together.

1. Place the initial set of LDPE layers (up to the first HMPE layer) onto the template, aligning all layers to the edge up to the template.

2. Place the first layer of HMPE onto the LDPE, and note the angle of the layer.

[2]This is why the 'bad' side of the template should be facing up, as we don't want to scratch the good side with the razorblade.

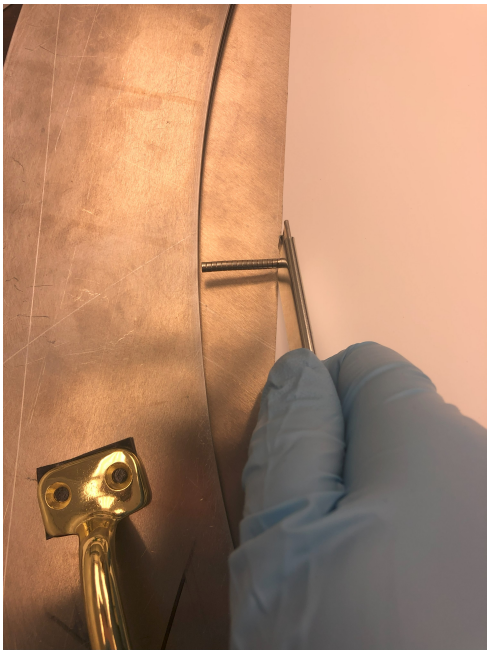


Figure A.9: Photo of the razor blade with a length of magnets to make a circular cut of an AR layer.

3. Place the next set of LDPE layers (up to the next HMPE layer) onto the template.

4. Place the next layer of HMPE onto the LDPE, now rotated by the desired amount. I have done this by aligning layers using an angle ruler. By eye align one layer fiber orientation to one axis of the angle ruler, and the other layer fiber orientation to the other axis of the angle ruler, set to the desired offset (typically 30 degrees).

5. Repeat steps 4 and 5 until the full recipe is laid out.

6. Temporarily tack the full stack (window laminate layers only, not the aluminum template) together with a set of paperclips, evenly spaced around the circumference of the laminate.

7. Move the stack to the large autoclave by lifting the template up with the stack on

top (best done with two people, with an additional person opening doors).

8. 'Pizza slide' the laminate stack onto the autoclave baseplate (Step 2 of §B.2 Large Autoclave procedure).



Figure A.10: Photo of full laminate stack in large autoclave, with paper clips around the edge.

9. Follow the rest of the steps in §B.2 to bake the laminate.

10. Record the recipe and final measured thicknesses in the PE window inventory in the BA Optics google drive.

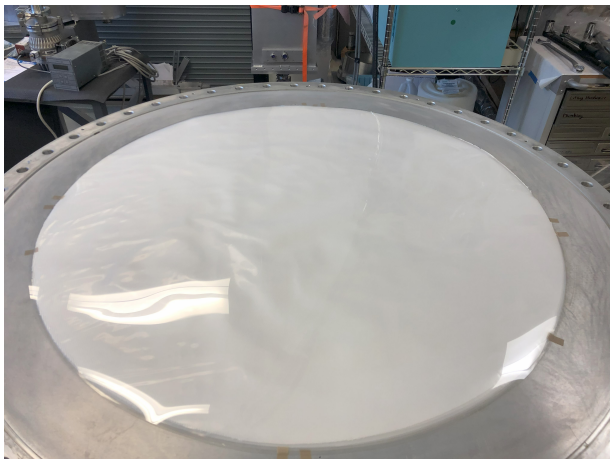


Figure A.11: Photo of full laminate stack in large autoclave, with top release layer.

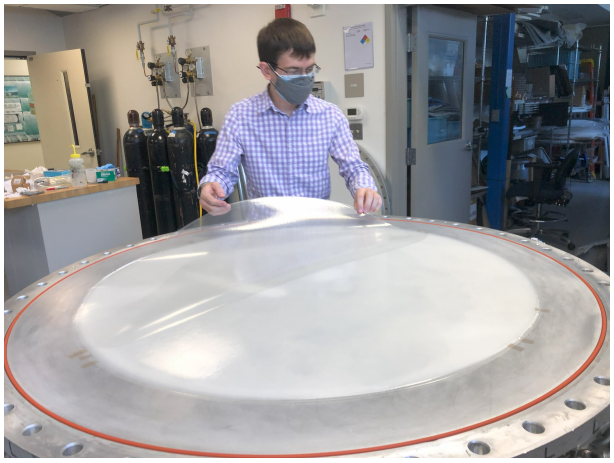


Figure A.12: Photo of full laminate stack in large autoclave after lamination bake. Matthew Petroff is removing the top release layer.

Appendix B

How to Run the Optics Autoclaves

The optics autoclaves are two very useful devices we use to bake materials that require high heat and pressure: these include laminate windows (discussed in Chapter 2) and compressed ePTFE AR coats (discussed in Chapter 3). They can also be used for laminating many-layer stacks (like a thin window, or a multi-layer AR coat), as the double-volume design allows for pulling vacuum above and below a vacuum barrier simultaneously (zero-pressure evacuation[1]). Each autoclave has a slightly different operating procedure, which are covered in separate sections.

[1]Zero-pressure evacuation pulls a vacuum above and below a vacuum barrier simultaneously, which prevents bubbles from becoming trapped in many-layer stacks such as thin windows or multi-layer AR coats. After a sufficient vacuum is pulled in the entire chamber, the volumes above and below the barrier are separated, and the top volume can be pressurized.

B.1 Small Autoclave

The small autoclave is a square chamber defined by a 1.5" thick square aluminum "ring" with two 1" thick plates. The ring has two o-ring grooves on the top and the bottom, which are silicon'ed in to ease assembly (this has led to some amount of o-ring shedding, as the o-ring is overly compressed by the volume taken up by the silicon and is squeezed in the groove). The bottom plate has threaded holes for the clamping 1/2" bolts, and two vacuum ports to pull and monitor the pressure below the vacuum barrier, respectively. The top chamber pressurized through a 'yor-lok' copper pipe fitting on one side of the square ring.

The procedure for using the small autoclave is as follows.

1. Clean the bottom plate of the autoclave thoroughly with isopropyl out to the o-ring surface (should be visible by a faint impression).

2. Place the to-be-baked optics stack on the baseplate, between the vacuum port holes (including release layers).

3. Triple check that stack up is adequately far from vacuum holes in base plate. LDPE will typically flow ∼few millimeters during baking.

4. Tack down the stack with teflon tape to prevent layer slippage (3-4 pieces of tape, only on the very edge).

5. Place a small piece of breather cloth over the two vacuum port holes to prevent plastic from flowing into the vacuum port.

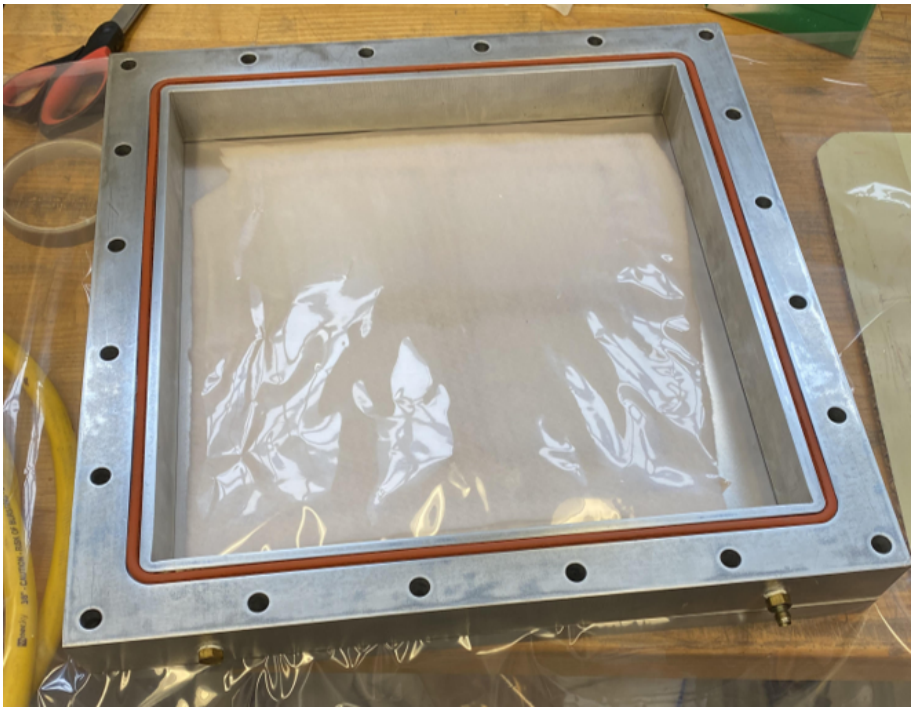


Figure B.1: Picture of the interior of the small autoclave.

6. Cut and place breather cloth such that it covers the stack completely, does not interfere with autoclave ring, and also adequately covers vacuum ports.

7. Cover entire stack with Stretchlon (vacuum barrier material), making sure that the vacuum barrier fully covers the base plate.

8. Wipe down the bottom o-ring on the autoclave ring.

9. Place autoclave ring over vacuum barrier, such that the yor-lok fitting is in front and to the right hand side (that o-ring should have just been cleaned).

10. Pull the Stretchlon taut. Ensure that no wrinkles are over the optic, nor travel below the autoclave ring.

11. Clean the top plate of the autoclave thoroughly with isopropyl out to the o-ring

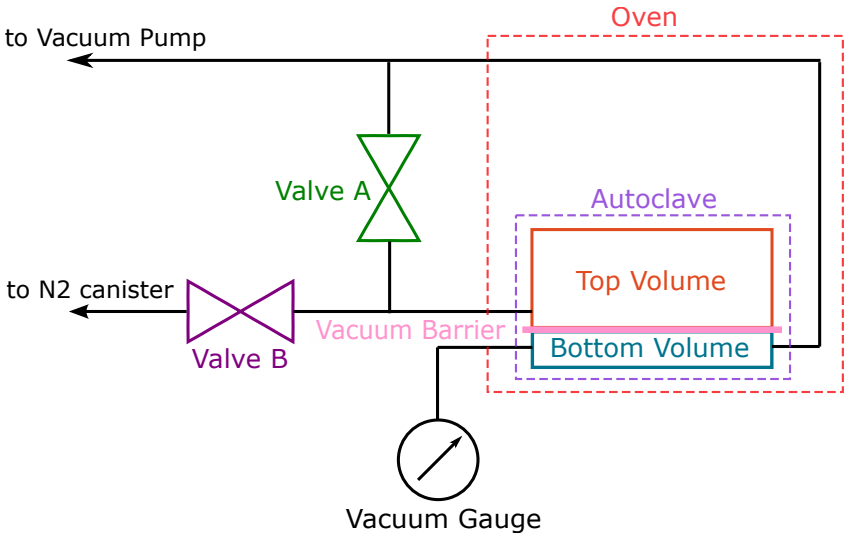


Figure B.2: Pipe circuit of small autoclave.

surface (should be visible by a faint impression).

12. Place the top plate on the autoclave ring, clean side down.

13. Punch through the vacuum barrier in the 1/2" clamping holes.

14. Thread the 1/2" bolts into the bottom plates from the top by hand, then tighten the bolts with hex wrenches in a four pointed star pattern.

15. Put the autoclave in the oven yor-lok fitting and baseplate ports facing out, taking care to not bang the baseplate vacuum ports on the sides of the oven.

16. Attach both vacuum tubes to the sawtooth vacuum ports on the bottom plate (one goes to the vacuum gauge on top of the oven, the other is the vacuum line to the pump).

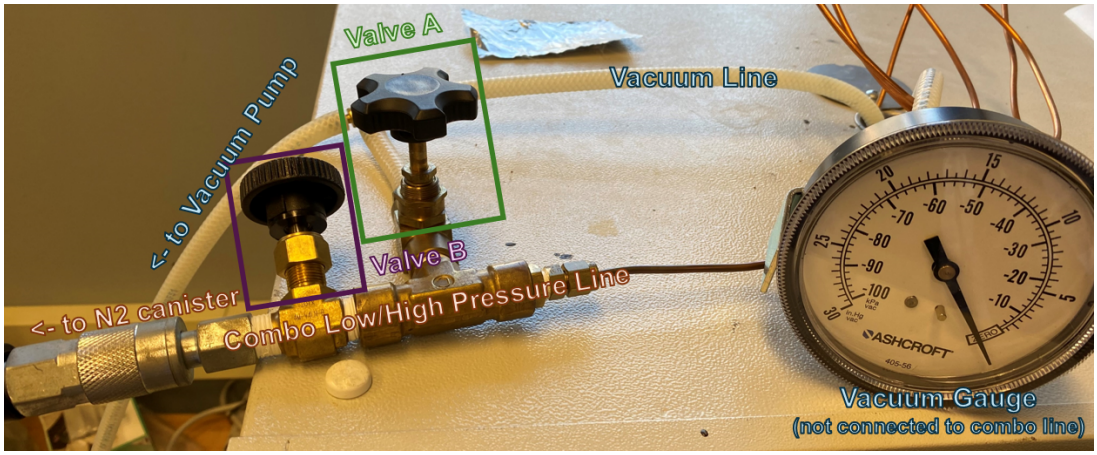


Figure B.3: Photo of the pipe circuit for the small autoclave on top of the oven. When vacuum pump is activated, vacuum is always pulled in the vacuum line, but the combo low/high pressure line can be switched from vacuum to the atmosphere or high pressure nitrogen. The vacuum gauge to the right is connected to the bottom plate, and the only path to the gauge from the pump is through the vacuum barrier.

17. Attach the copper pipe yor-lok fitting to the front fitting. Tighten with both open ended 7/16 wrenches, taking care not to twist the fitting on the autoclave side.

18. Do a zero pressure evacuation. Open Valve A and close Valve B (see Figure B.3) to pull vacuum on both sides of the vacuum barrier simultaneously, and turn on the pump with the extension cord on the bench.

19. Give the pump at least fifteen minutes to pull vacuum. All vacuum gauges (the one above the pump[2] and the separated one above the oven) should get below at least $\sim$ 25 inHg. They are unlikely to get all the way to 29 inHg.

20. Check leak tightness with a valved off pressure rise test. When the pump is valved

[2]This one is miscalibrated: good vacuum is down between the J. and L in Kurt J. Lesker.

off, the pressure should not increase visibly in a minute. If the pressure does increase significantly while the valve is closed, give it another few minutes on the pump[3] and attempt a pressure rise test again. If it fails again, may need to perform leak checks if the autoclave is going to be pressurized.

21. Pressurize the top volume by first closing Valve A (isolating the top volume from the lower volume) and then opening Valve B (again see Figure B.3) to either the atmosphere or the nitrogen cylinder. If pressurizing, increase pressure in nitrogen cylinder regulator slowly (ie, if goal is 150 psi set nitrogen regulator to ∼130 psi and wait, then when both the top pressure gauge and regulator agree increase to 140 psi, then 150 psi).

22. Select the desired heat cycle either pre-programmed in the oven.

23. When the cycle is finished (usually the next day), un-pressurize the chamber (if pressurized) by removing the quick connect high pressure fitting. Valve off the cylinder and un-pressurize the line as well.

24. After pressure is released, turn off the pump and open the autoclave in a rough reversal of steps 9 thru 17.

---

[3]As the inner diameter of the copper pipes are very small, and the pipe is very long, flow through it is very slow. A slight pressure rise can indicate there is still gas in the top chamber, which has taken time to reach the pump.

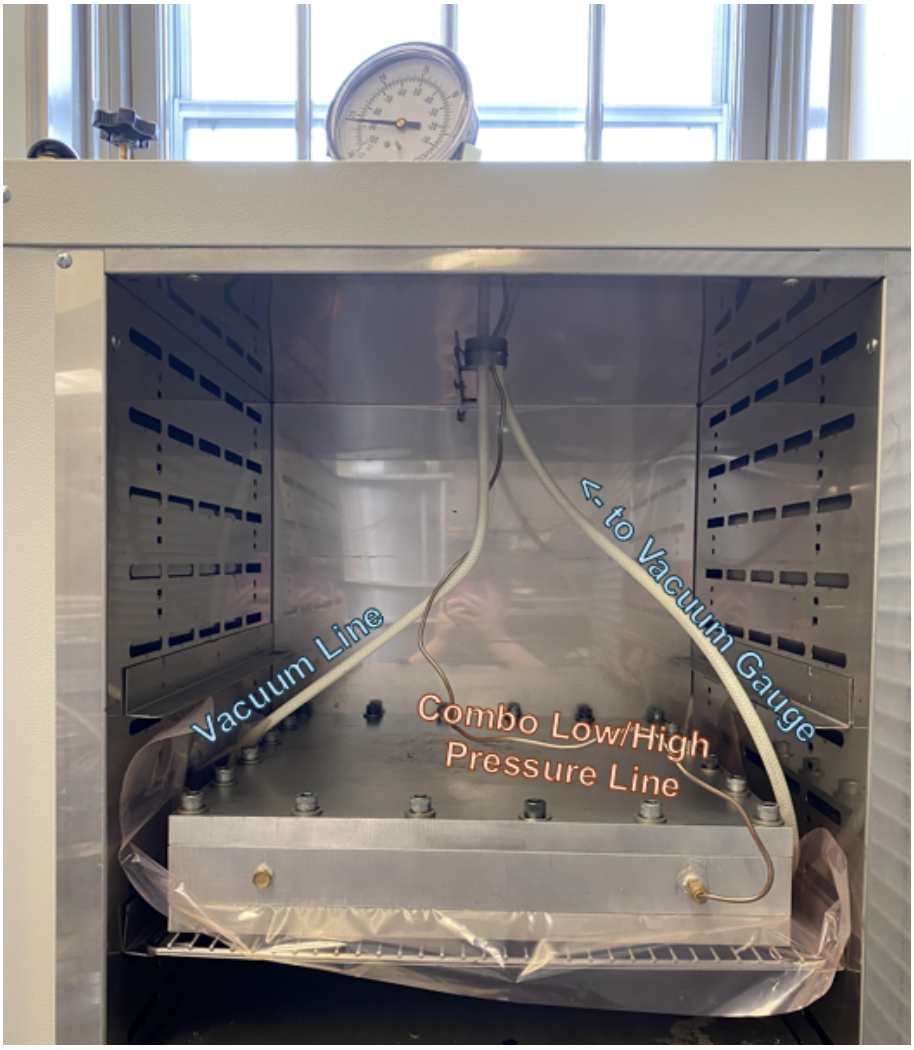


Figure B.4: Picture of the oven with the small autoclave set up inside.

B.2 Large Autoclave

The large autoclave is specifically designed to manufacture laminate windows, and as such is somewhat limited to be used for relatively thin optics. The central ring (which defines the interior volume) is only 1.5" high. Additional spacers could be added in the future to increase the inner volume, if needed, but currently the large optics oven—which is used for vacuum conformed AR lamination or annealing thick plastic

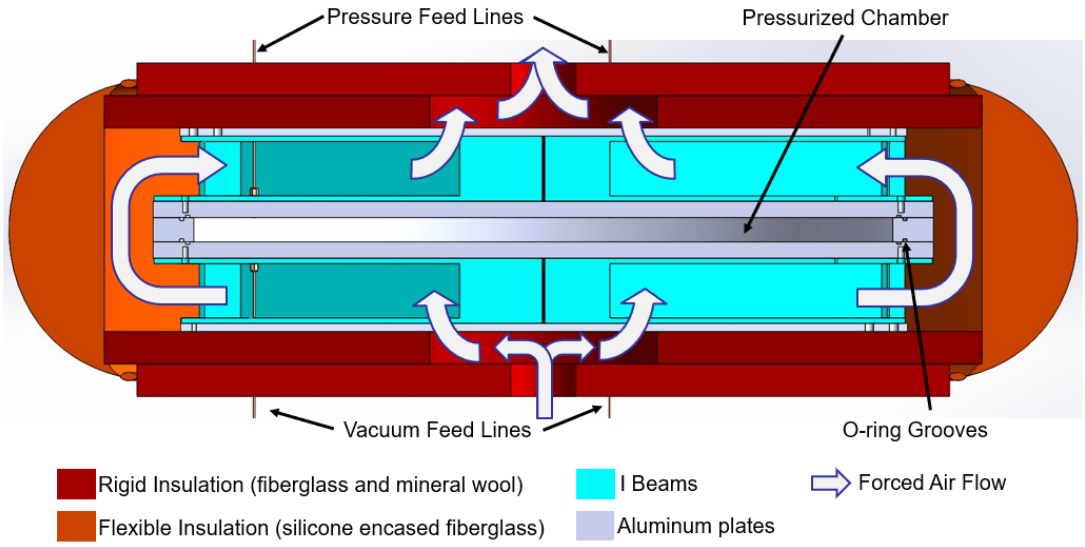


Figure B.5: Cutaway view of the large optics autoclave, with components color-coded and arrows pointed at small features. Large block arrows show direction of forced air flow.

stock—in the high bay and the autoclave—used for thin window lamination and thin optics zero-pressure evacuation AR lamination—sufficiently cover the current use cases needed for BA optics generation.

The large optics autoclave is itself both an oven and a pressure vessel (hence, 'autoclave' which is a heated pressure vessel). It is designed to be heated up to 140°C (with safety measures such as thermal breakers incorporated into the thermal circuit to prevent heating above 150°C) and pressurized to 10 atm (though it has been pressurized to 12 atm before safely). The majority of the exterior of the autoclave is either thermal insulation or mechanical support (see Figure B.5). Air is circulated in a closed loop within the thermal insulation to help equilibrate the temperature throughout the mechanical supports. Additional information about the design and testing of the large optics autoclave can be found in Eiben et al. (2022).

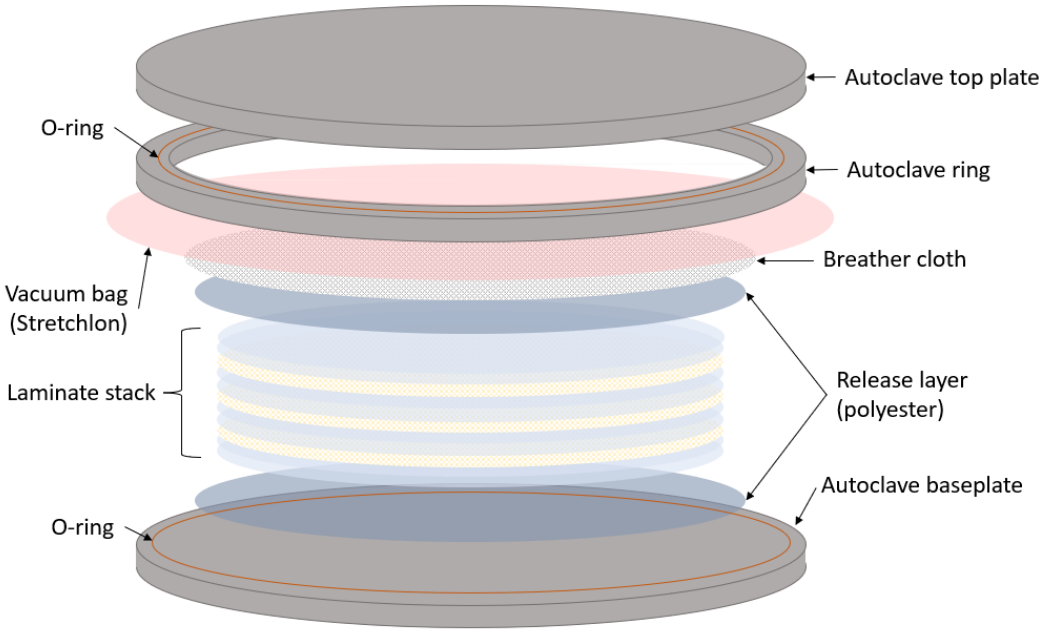


Figure B.6: Cartoon of a full stack in a thin window autoclave run. The laminate stack will change depending on the optic being baked, but the release layers, breather cloth, vacuum barrier and autoclave plates are universal.

This entire procedure works best with two people working in parallel. Some activities (*italicized*) require at least two people, and some are simply easier with two people [noted with square brackets].

1. Clean bottom release layer of autoclave thoroughly with isopropyl. It should already be tacked down to baseplate with teflon tape; replace tape if needed.

2. [Place the to-be-baked optics stack on the bottom release layer.]

3. Triple check that stack up is adequately far from vacuum holes in base plate. LDPE will typically flow ∼few millimeters during baking.

4. Tack down the stack with teflon tape to prevent layer slippage (3-4 pieces of tape,

only on the very edge).

5. Clean top release layer thoroughly with isopropyl. Place on top of stack. Make sure that the release layer does not cover vacuum ports: cut the release layer around the ports, if necessary.

6. Tack down release layer with teflon tape to prevent layer slippage (3-4 pieces of tape, only on the very edge)

7. Place a small piece of breather cloth over the four vacuum port holes to prevent plastic from flowing into the vacuum port.

8. [Cut and place breather cloth such that it covers the stack completely, does not interfere with autoclave ring, and also adequately covers vacuum ports. This can be done by unrolling the breather cloth over the stack, marking a rough circle with sharpie (the bottom o-ring and stack should be visible through the breather cloth) and cutting with scissors.]

9. Clean and inspect bottom plate o-ring. **Do not grease it**, it's not worth it.

10. [Cover entire stack with Stretchlon (vacuum barrier material), making sure that the vacuum barrier fully covers the base plate.]

11. *Working on opposite sides, pull the Stretchlon taut and tack the edges down with blue painters tape below the baseplate of the autoclave. Take care to not tape over the thru bolt holes on the baseplate.*

12. *Place autoclave ring over vacuum barrier, o-ring side up. Make sure that the clocking is correct (0/90/180/270 marks on ring should align with marks on base plate).*



Figure B.7: Picture of the interior of the large autoclave at Step 14. Keara Carter for scale.

13. Check that vacuum barrier has remained taut with the ring placed over it, and no wrinkles have formed. Ensure that no wrinkles are over the optic, nor travel below the autoclave ring.

14. **Take a picture, at this step or before it.**

15. [Punch through the ring holder (1/4") threaded holes (the small ones in the ring). Then clamp down the ring with the ring holder bolts from underneath the baseplate, with the small ratcheting wrenches. Work roughly in a star pattern.]

16. Clean the autoclave ring o-ring. **Grease the o-ring groove over the patch (darker/rougher portion of the o-ring groove in the ring) with Apiezon High Temperature (E) Vacuum Grease**, at the place marked 'patch' in red on

the exterior of the ring (between the 0/270 marks).

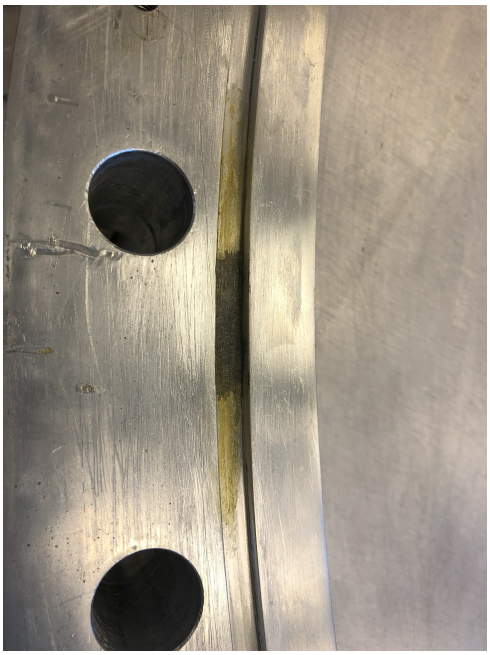


Figure B.8: The patch in the top o-ring groove on the autoclave ring greased with high temperature E grease.

17. Punch through the 3/4" main autoclave clamping holes (the big ones in the ring) with the sharpened pipe and a hammer.

18. Do a quick inspection to make sure you've punched through all the holes, and to double check that the stack up looks good. This is the last time you will see the stack up pre-bake. Wish it well on its journey!

19. *Pull over the top of the autoclave with the hoist. Nudge the base to be under the hoist as evenly as possible if it's not centered. Make sure clocking of top matches ring/baseplate.* One person should be the operator of the hoist, one person the spotter.

20. ***Clearly communicate motion of the top plate. DO NOT put any appendages (fingers, etc) below the top plate.***

21. *Lower topplate (make sure that bolts do not touch the exposed o-ring), watching for tilting of the top plate as the 3/4" bolts enter the ring. Minor tilts may be corrected by lightly pulling down on the high side of the top plate; major tilts may require backing out and readjusting the alignment of the autoclave halves. Tap bolts through if they get stuck.*

22. [Thread nuts on 3/4" bolts by hand. Make sure to inspect nuts for stripping of edges, and bolts for difficult threads. Bolts can be re-threaded at this step.]

23. [Tighten nuts/bolts with large wrenches.]

24. **If only pressurizing to 1 atm (opening top chamber to room), this step can be skipped; if pressurizing above 1 atm, this step is REQUIRED.** *Tighten 3/4" bolts with torque wrench to 280 ft-lbs. Torque wrench should be taped to appropriate torque (to prevent it from slipping off spec).* One person should be using the torque wrench on the 3/4" bolt heads, and one person holding nuts steady (an aluminum tube may be used on the nut wrench to aid leverage). Work steadily, without jerking the torque wrench and listening carefully for the click indicating torque reached, in a circle around the autoclave.

25. *Put on flexible insulation. Use a taut-line hitch*[4] *knot to keep tension in tie lines. Cover open hole with flexible insulation pillow.*

[4] `https://en.wikipedia.org/wiki/Taut-line_hitch`

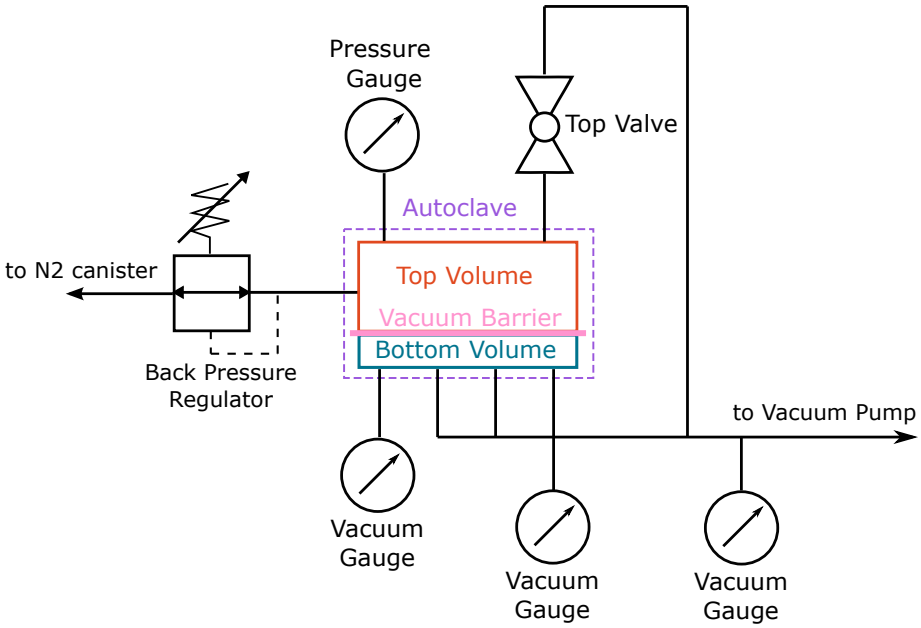


Figure B.9: Pipe circuit of large autoclave.

26. Attach the vacuum tube open to the top plate port (an open sawtooth pipe port with a ball valve around the 270 mark above the insulation on the top plate), and the pressure line to the nitrogen cylinder to the open quick connect fitting on the back pressure valve (this closes off that port).

27. Plug in the vacuum pump (it should already be plugged into an extension cord with the closed loop exhaust fan).

28. Open top vacuum valve (in-line with pipe is open, perpendicular is closed) to pressure chamber. Turn on the vacuum pump by turning on the pump extension cord. This pulls vacuum on both top and bottom of vacuum barrier (zero-pressure evacuation).

29. Give the pump at least fifteen minutes to pull vacuum. All vacuum gauges[5] should get below at least ∼20 inHg, if not 25 inHg. They are unlikely to get all the way to 29 inHg.

30. Check leak tightness with a valved off pressure rise test. When the pump is valved off, the pressure should not increase visibly in a minute. If the pressure does increase significantly while the valve is closed, give it another few minutes on the pump and attempt a pressure rise test again. If it fails again, may need to perform leak checks if the autoclave is going to be pressurized.

31. Close top vacuum valve (isolates the top volume). Monitor if vacuum pressure changes. It shouldn't.

32. Attach circulating vent hose.

33. **If doing a 1 atm bake**, remove the pressure line to the nitrogen cylinder and allow the top volume to reach room pressure. **If doing a pressurized run,** pressurize top volume from nitrogen cylinder slowly (ie, if goal is 150 psi set nitrogen regulator to ∼130 psi and wait, then when both the top pressure gauge and regulator agree increase to 140, then 150). Back pressure valve may start

[5]There are four pressure gauges on the autoclave (see the large autoclave pipe circuit in Figure B.9): one a pressure gauge (with a rough vacuum range) connected to the pressure chamber (top volume) and three in the vacuum circuit: one directly next to the pump, one connected to the manifold between three of the vacuum holes in the bottom plate, and one connected only to the remaining vacuum port strapped to a strut next to the pump (as the only path from the gauge to the pump is through the bottom volume, this one ensures that a good vacuum is being pulled under the vacuum barrier). The disconnected gauge and the pump gauge should agree if good vacuum is pulled.

releasing pressure prematurely, so monitor for releasing gas and increase set pressure if not at desired pressure yet. This process may take tens of minutes.

34. Take pictures/record the pressure gauges above the autoclave/on the nitrogen regulator to note the starting pressures.

35. Plug in the heater/thermometer cables to the top plate, and plug in the high voltage (240 V) line to the outlet. Flip the right breaker (with the same black outlet cord plugged in below it) in the compressor room.

36. Plug in the heater PID control ethernet cable to autoclave CalControl laptop, and start the heat control cycle python script. Monitor the script for a few cycles: none of the heater set values should be 1 (maximum), and most thermometers should read out room temperature. If that is not the case, abort and troubleshoot.

37. **Record the run in the pressure cooker logbook.**

38. Return after heat cycle is finished (this is generally between 5.5 to 7 hours after the script is started) to remove circulating hose from top, to allow for faster cooling.

39. Compare the pressures to those recorded earlier. Note any differences (**be advised**: if the back-pressure valve was too close to the set pressure, it may be emptying the nitrogen cylinder).

40. Once cool (typically the next day), terminate control script, flip the high voltage breaker, carefully depressurize the top volume (remove the high pressure line to the nitrogen cylinder, valve off the cylinder and un-pressurize the line), and turn off and unplug any power cables. Then open the autoclave in a rough reversal of steps 12-26.

Appendix C

How to Layer DeWAL

DeWAL is a very strange material; under heat compression, layers of DeWAL stick together in a process we've taken to calling 'bonding' and it, like other ePTFEs, is fine-tunable with heat compression. However, it is also the thinnest form of ePTFE I have found that is manufacturer on the required BA scales, which makes it uniquely difficult to work with. I've written up the procedure we have developed to lay up DeWAL effectively, after a decent amount of trial and error, but this process is subject to improvement.

There are three main pieces of equipment required for this procedure:

- The frame; the DeWAL 'parchment stretching' (as it's designed like a traditional tanner parchment stretcher) frame, which is a wooden frame with thin steel rods on two sides (the short sides), and magnetic clips looped around elastic held in place with thumb tacks along the other two sides (the long sides).
- The roll; the DeWAL roll, placed on a roll holder to ease unrolling.

- The template; the same template used for cutting window materials. The turntable should be below the template to ease rotation between DeWAL layers. The current template may be able to be replaced with a smaller diameter template.
- Additional materials required are the blue lab jacks (used to hold the frame above the template while working), good sharp scissors, and the appropriate thick green Nitrile texture-less gloves.

This procedure is best accomplished in the clean room to limit dust between layers. The procedure currently takes about 20 minutes per layer and two people, but both may be reduced with practice and procedural changes (a smaller frame and template would reduce the scale and make everything easier to handle with a single person).

1. Pull DeWAL over the frame in the direction of the rods. Measure the distance roughly edge to edge within the frame to keep the sheet roughly centered in the frame.
2. Attach magnet strips to first the further free end, then to the closer roll end.
3. Cut the length loose from the roll.
4. Put the frame up on the 4 lab jacks above the template.
5. By hand, pull the sheet tight (this currently requires two people). One holds the corners of the sheet and very carefully stretches the edge to be mostly straight and mostly flat, the other picks up and puts down the magnet strip to lock the tension. Try to keep the sheet as centered as possible.
6. On the other end of the sheet, repeat the by-hand-sheet-tightening.

7. Working from one end of the sheet, in tandem, attach the 'parchment stretching' magnets.

   (a) Lift one of the thumb tacks to release the tension in the elastic string BEFORE attaching the magnets to the edge.

   (b) Try to align the magnet to be roughly centered between the two securing thumb tacks.

   (c) Then, after the edge (within the cm folded area) is secured, tighten the elastic by pulling on one end gently, until the DeWAL edge between the last secured point and this one is roughly parallel with the frame. Having the elastic looped once around the thumb tack helps keep the tension even as you pull.

   (d) Secure the tension on that magnet by pushing down the thumb tack with the elastic looped around it.

   (e) Repeat until all magnets on the frame are attached to the edge of the DeWAL.

8. Repeat the by-hand-sheet-tightening, now trying to pull as many ripples out of the center as possible. Lock the center tension first by lowering just the center of the magnet strip down, and work the edges mostly flat individually. The ripples at the edge may remain: the primary concern is getting the center as flat as possible.

9. Repeat the above step on the other straight side.

10. Tighten any loose parchment magnets until there is even tension on the edge.

11. By now the center should be mostly flat, with some ripples to the edge. Repeat the previous 3 steps if the center still has many ripples.

12. Lower the frame onto the template. Be careful that the frame does not impact with the template (there is low clearance on the edges). There should be about a half inch of space between the top of the frame and the bottom of the circle plate. This should very lightly tension the sheet, smoothing out any further wrinkles in the center.

13. GENTLY glide gloved hands over the center, to push out any air. Gloved hands should be the thick green smooth-tipped nitrile gloves: other gloves tend to catch on the DeWAL or impart a texture. Ensure that your gloved hands don't have any wrinkles where you touch the DeWAL surface.

14. Working by hand from the center out, gently smooth out the edge ripples. If you accidentally press in a wrinkle, STOP! Then gently push up the edge of the sheet with your fingers until that area is no longer in contact with the surface below it. Stretch the wrinkle out carefully. You should be able to smooth out the wrinkle with your finger tips.

15. Once all ripples are gone from the surface of the DeWAL in contact with the circle plate, carefully pull off the parchment-magnets that are above the plate. DO NOT release the magnets that are not above the plate (there should be 3 remaining on each corner, 12 total). Releasing the magnets works best by gently pushing down on the edge of the sheet and pulling the top magnet up, then pulling the parchment-magnet away.

16. Cut off the edge of the DeWAL (roughly a cm in from the edge) with scissors, but only the edge that is above the circle plate.

17. Gently smooth out any ripples that may have formed from the cutting process.

18. Release all the parchment-magnets.

19. Gently pull up the strip magnet on one edge, and gently roll the end of the sheet within the frame.

20. No part of the frame should be in contact with the DeWAL sheet. Lift the frame, and begin again.

Appendix D

Ideal (Chebyshev) AR Indexes for Millimeter Optics

I thought it may be useful[1] to compile some tables of the ideal Chebyshev anti-reflection layer indices of refraction for common optical materials used in millimeter receivers. I am only exploring a handful of bandwidths and core materials; the bandwidths are roughly the same bandwidths as the latest generation of BICEP/*Keck* receivers, and core materials are all materials that I know have been used for millimeter optics in the past, or close equivalents. If the reader is interested in another material or fractional bandwidth outside the bounds of this appendix they are encouraged to run the Chebyshev solutions outlined in Chapter 3 of this thesis.

I also only report the ideal index of a given AR layer, not the corresponding physical

[1] Also, I thought it would be funny to have an appendix in my thesis modeled after an engineering textbook appendix (many of which are just reference tables).

thickness, which is wavelength dependent. To find the required thickness is a simple operation:

$$h = \frac{\lambda}{4n} = \frac{c}{\nu}\frac{1}{4n} \tag{D.1}$$

where λ is the center wavelength/ν is the center frequency of the band, n is the index of refraction of the layer, and h is the required thickness to generate a quarter wavelength layer.

Core materials are listed in order of index of refraction, from lowest to highest.

D.1 PTFE (Teflon): $n \sim 1.43$

Used previously in millimeter receivers for filters (BICEP2 and Keck Array Collaborations et al. 2015; Grimes et al. 2024)

Table D.1: Chebyshev single, double, and triple anti-reflection coat indexes on PTFE (n=1.43) for fractional bandwidths ranging from 0.25 to 0.8.

Layer	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.25**	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.4**	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.6**	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.8**
Single	$n_1 = 1.196$	$n_1 = 1.196$	$n_1 = 1.196$	$n_1 = 1.196$
Double	$n_1 = 1.095$	$n_1 = 1.098$	$n_1 = 1.105$	$n_1 = 1.114$
	$n_2 = 1.305$	$n_2 = 1.302$	$n_2 = 1.294$	$n_2 = 1.284$
Triple	$n_1 = 1.047$	$n_1 = 1.049$	$n_1 = 1.054$	$n_1 = 1.062$
	$n_2 = 1.196$	$n_2 = 1.196$	$n_2 = 1.196$	$n_2 = 1.196$
	$n_3 = 1.365$	$n_3 = 1.363$	$n_3 = 1.356$	$n_3 = 1.346$

D.2 HDPE: $n \sim 1.55$

Used previously in millimeter receivers for windows and lenses (BICEP/Keck Collaboration et al. 2022; Eiben et al. 2024; Hui et al. 2018; Moncelsi et al. 2020). Similar in index to other forms of polyethylene (LDPE, UHMWPE, HMPE) and polypropylene (Lamb 1996).

Table D.2: Chebyshev single, double, and triple anti-reflection coat indexes on HDPE (n=1.55) for fractional bandwidths ranging from 0.25 to 0.8.

Layer	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.25**	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.4**	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.6**	$\frac{\Delta\nu}{\nu_{cen}}$ **= 0.8**
Single	$n_1 = 1.245$	$n_1 = 1.245$	$n_1 = 1.245$	$n_1 = 1.245$
Double	$n_1 = 1.118$	$n_1 = 1.122$	$n_1 = 1.130$	$n_1 = 1.142$
	$n_2 = 1.386$	$n_2 = 1.382$	$n_2 = 1.372$	$n_2 = 1.358$
Triple	$n_1 = 1.058$	$n_1 = 1.061$	$n_1 = 1.067$	$n_1 = 1.077$
	$n_2 = 1.245$	$n_2 = 1.245$	$n_2 = 1.245$	$n_2 = 1.245$
	$n_3 = 1.465$	$n_3 = 1.461$	$n_3 = 1.453$	$n_3 = 1.44$

D.3 Nylon: $n \sim 1.73$

Used previously in millimeter receivers for filters (BICEP/Keck Collaboration et al. 2022; Abazajian et al. 2019)

Table D.3: Chebyshev single, double, and triple anti-reflection coat indexes on nylon (n=1.73) for fractional bandwidths ranging from 0.25 to 0.8.

Layer	$\frac{\Delta\nu}{\nu_{cen}} = 0.25$	$\frac{\Delta\nu}{\nu_{cen}} = 0.4$	$\frac{\Delta\nu}{\nu_{cen}} = 0.6$	$\frac{\Delta\nu}{\nu_{cen}} = 0.8$
Single	$n_1 = 1.315$	$n_1 = 1.315$	$n_1 = 1.315$	$n_1 = 1.315$
Double	$n_1 = 1.150$	$n_1 = 1.155$	$n_1 = 1.165$	$n_1 = 1.180$
	$n_2 = 1.504$	$n_2 = 1.498$	$n_2 = 1.485$	$n_2 = 1.466$
Triple	$n_1 = 1.073$	$n_1 = 1.077$	$n_1 = 1.084$	$n_1 = 1.097$
	$n_2 = 1.315$	$n_2 = 1.315$	$n_2 = 1.315$	$n_2 = 1.315$
	$n_3 = 1.612$	$n_3 = 1.607$	$n_3 = 1.595$	$n_3 = 1.577$

D.4 Quartz: $n \sim 2.10$

Used previously in millimeter receivers for windows (Barkats et al. 2018).

Table D.4: Chebyshev single, double, and triple anti-reflection coat indexes on quartz (n=2.10) for fractional bandwidths ranging from 0.25 to 0.8.

Layer	$\frac{\Delta\nu}{\nu_{cen}} = \mathbf{0.25}$	$\frac{\Delta\nu}{\nu_{cen}} = \mathbf{0.4}$	$\frac{\Delta\nu}{\nu_{cen}} = \mathbf{0.6}$	$\frac{\Delta\nu}{\nu_{cen}} = \mathbf{0.8}$
Single	$n_1 = 1.449$	$n_1 = 1.449$	$n_1 = 1.449$	$n_1 = 1.449$
Double	$n_1 = 1.208$	$n_1 = 1.215$	$n_1 = 1.230$	$n_1 = 1.251$
	$n_2 = 1.738$	$n_2 = 1.728$	$n_2 = 1.708$	$n_2 = 1.678$
Triple	$n_1 = 1.100$	$n_1 = 1.105$	$n_1 = 1.116$	$n_1 = 1.133$
	$n_2 = 1.449$	$n_2 = 1.449$	$n_2 = 1.449$	$n_2 = 1.449$
	$n_3 = 1.908$	$n_3 = 1.900$	$n_3 = 1.882$	$n_3 = 1.853$

D.5 Alumina: $n \sim 3.10$

Used previously in millimeter receivers for filters and lenses (BICEP/Keck Collaboration et al. 2022; Hui et al. 2018; Moncelsi et al. 2020). Similar in index to sapphire (used for half-wave plates).

Table D.5: Chebyshev single, double, and triple anti-reflection coat indexes on alumina (n=3.10) for fractional bandwidths ranging from 0.25 to 0.8.

Layer	$\frac{\Delta\nu}{\nu_{cen}} = 0.25$	$\frac{\Delta\nu}{\nu_{cen}} = 0.4$	$\frac{\Delta\nu}{\nu_{cen}} = 0.6$	$\frac{\Delta\nu}{\nu_{cen}} = 0.8$
Single	$n_1 = 1.761$	$n_1 = 1.761$	$n_1 = 1.761$	$n_1 = 1.761$
Double	$n_1 = 1.334$	$n_1 = 1.346$	$n_1 = 1.371$	$n_1 = 1.395$
	$n_2 = 2.323$	$n_2 = 2.303$	$n_2 = 2.262$	$n_2 = 2.222$
Triple	$n_1 = 1.156$	$n_1 = 1.164$	$n_1 = 1.182$	$n_1 = 1.21$
	$n_2 = 1.761$	$n_2 = 1.761$	$n_2 = 1.761$	$n_2 = 1.761$
	$n_3 = 2.681$	$n_3 = 2.662$	$n_3 = 2.623$	$n_3 = 2.562$

D.6 Silicon: $n \sim 3.44$

Used previously in millimeter receivers for lenses (Iuliano et al. 2018).

Table D.6: Chebyshev single, double, and triple anti-reflection coat indexes on silicon (n=3.44) for fractional bandwidths ranging from 0.25 to 0.8.

Layer	$\frac{\Delta\nu}{\nu_{cen}} = 0.25$	$\frac{\Delta\nu}{\nu_{cen}} = 0.4$	$\frac{\Delta\nu}{\nu_{cen}} = 0.6$	$\frac{\Delta\nu}{\nu_{cen}} = 0.8$
Single	$n_1 = 1.855$	$n_1 = 1.855$	$n_1 = 1.855$	$n_1 = 1.855$
Double	$n_1 = 1.370$	$n_1 = 1.383$	$n_1 = 1.411$	$n_1 = 1.432$
	$n_2 = 2.511$	$n_2 = 2.487$	$n_2 = 2.438$	$n_2 = 2.403$
Triple	$n_1 = 1.172$	$n_1 = 1.181$	$n_1 = 1.200$	$n_1 = 1.232$
	$n_2 = 1.855$	$n_2 = 1.855$	$n_2 = 1.855$	$n_2 = 1.855$
	$n_3 = 2.934$	$n_3 = 2.913$	$n_3 = 2.866$	$n_3 = 2.793$

References

Abazajian, K., Addison, G., Adshead, P., et al. 2019, arXiv e-prints, arXiv:1907.04473

Ade, P. A. R., Pisano, G., Tucker, C., & Weaver, S. 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 6275, Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III, ed. J. Zmuidzinas, W. S. Holland, S. Withington, & W. D. Duncan, 62750U

Afsar, M. N. 1987, IEEE Transactions on Instrumentation and Measurement, IM-36, 530

Badr, Y., Ali, Z., & Khafagy, R. M. 2000, Radiation physics and chemistry (Oxford, England : 1993), 58, 87

Barkats, D., Dierickx, M. I., Kovac, J. M., et al. 2018, in Proc. SPIE, Vol. 10708, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX, ed. J. Zmuidzinas & J.-R. Gao, 107082K

BICEP2 and Keck Array Collaborations, Ade, P. A. R., Aikin, R. W., et al. 2015, The Astrophysical Journal, 806, 206

BICEP2 and Keck Array Collaborations, Ade, P. A. R., Ahmed, Z., et al. 2019, The Astrophysical Journal, 884, 114

BICEP2 Collaboration, Ade, P. A. R., Aikin, R. W., et al. 2014, Phys. Rev. Lett., 112, 241101

BICEP2 Collaboration, Keck Array Collaboration, SPIDER Collaboration, et al. 2015a, ApJ, 812, 176

BICEP2 Collaboration, Ade, P. A. R., Aikin, R. W., et al. 2015b, ApJ, 814, 110

BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., et al. 2016, ApJ, 825, 66

BICEP/Keck Collaboration, Ade, P. A. R., Ahmed, Z., et al. 2021, Phys. Rev. Lett., 127, 151301

—. 2022, Astrophys. J., 927, 77

—. 2024a, arXiv e-prints, arXiv:2411.10428

—. 2024b, arXiv e-prints, arXiv:2410.12089

BICEP/Keck Collaboration, Nakato, Y., Ade, P. A. R., et al. 2024c, arXiv e-prints, arXiv:2409.02296

Birch, J., Dromey, J., & Lesurf, J. 1981, Infrared Physics, 21, 225

Bryan, S., Austermann, J., Ferrusca, D., et al. 2018, in Proc. SPIE, Vol. 10708, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX, ed. J. Zmuidzinas & J.-R. Gao, 107080J

Buza, V. 2019, PhD thesis, Harvard University

Cheshire, J. 2024, PhD thesis, University of Minnesota

Choi, J., Ishitsuka, H., Mima, S., et al. 2013, Rev. Sci. Instrum., 84, 114502

Cornelison, J., Vergès, C., Ade, P. A. R., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, ed. J. Zmuidzinas & J.-R. Gao, 121901X

Cornelison, J. A. 2023, PhD thesis, Harvard University

Corrigan, L. 2019, Development of a microwave scatterometer and investigation of optical properties of BICEP3 materials

D'Alessandro, G., Paiella, A., Coppolecchia, A., et al. 2018, Infrared Physics and Technology, 90, 59

Datta, R., Chuss, D. T., Eimer, J., et al. 2021, Review of Scientific Instruments, 92, 035111

DeWAL. 2023, DeWAL ePTFE Membrane

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, ApJ, 142, 414

Dierickx, M., Ade, P. A. R., Ahmed, Z., et al. 2023, J. Low Temp. Phys., 211, 366

Dyneema. 2008, Dyneema ® high-strength, high-modulus polyethylene fiber, Tech. rep., Dyneema

Ebnesajjad, S. 2016, in Expanded PTFE Applications Handbook, Plastics Design Library (Oxford: William Andrew)

Eiben, M., Barkats, D., Balkanski, A., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, ed. J. Zmuidzinas & J.-R. Gao, 121902L

Eiben, M., Carter, K., Dierickx, M., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13102, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. J. Zmuidzinas & J.-R. Gao, 131022M

Einstein, A. 1916, Annalen der Physik, 354, 769

Ekin, J. W., & Zimmerman, G. O. 2007, Physics Today, 60, 67

Elwood, B. D., Grimes, P. K., Kovac, J., Eiben, M., & Meiners, G. 2024, in Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XII, ed. J. Zmuidzinas & J.-R. Gao, Vol. PC13102, International Society for Optics and Photonics (SPIE), PC1310220

Findley, W. N., & Tracy, J. F. 1974, Polymer engineering and science, 14, 577

Fixsen, D. J. 2009, ApJ, 707, 916

Frank, W., Schmidt, H., & Wulff, W. 1977, Journal of polymer science. Polymer symposia, 61, 317

Galeazzi, M. 2011, IEEE Transactions on Applied Superconductivity, 21, 267

Giannakopoulos, C., Vergès, C., Ade, P. A. R., et al. 2024, arXiv e-prints, arXiv:2409.16440

Goldfinger, D. C., Ade, P. A. R., Ahmed, Z., et al. 2022, in Proc. SPIE, Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, ed. J. Zmuidzinas & J.-R. Gao, 121901V

Golec, J. E., McMahon, J. J., Ali, A., et al. 2020, in Proc. SPIE, Vol. 11451, Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation IV, ed. R. Navarro & R. Geyl, 114515T

Grimes, P., Blundell, R., Leiker, P., et al. 2020, in Proc. 31st Int. Symp. Space Terahertz Technol., 60–66

Grimes, P. K., Keating, G. K., Blundell, R., et al. 2024, arXiv e-prints, arXiv:2406.17192

Guth, A. H. 1981, Phys. Rev. D, 23, 347

Halpern, M., Gush, H. P., Wishnow, E., & de Cosmo, V. 1986, Applied Optics, 25, 565

Hamimeche, S., & Lewis, A. 2008, Phys. Rev. D, 77, 103013

Hill, C. A., Bruno, S. M. M., Simon, S. M., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10708, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX, ed. J. Zmuidzinas & J.-R. Gao, 1070842

Hu, W., & Dodelson, S. 2002, ARA&A, 40, 171

Hubble, E. 1929, Proceedings of the National Academy of Science, 15, 168

Hui, H., Ade, P. A. R., Ahmed, Z., et al. 2018, in Proc. SPIE, Vol. 10708, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX, ed. J. Zmuidzinas & J.-R. Gao, 1070807

Inoue, Y., Matsumura, T., Hazumi, M., et al. 2014, Appl. Opt., 53, 1727

Irwin, K. D., Cho, H. M., Doriese, W. B., et al. 2012, Journal of Low Temperature Physics, 167, 588

Iuliano, J., Eimer, J., Parker, L., et al. 2018, in Proc. SPIE, Vol. 10708, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX, ed. J. Zmuidzinas & J.-R. Gao, 1070828

Jeong, O., Plambeck, R., Raum, C., Suzuki, A., & Lee, A. T. 2023, Appl. Opt., 62, 1628

Kamionkowski, M., & Kovetz, E. D. 2016, Annual Review of Astronomy and Astrophysics, 54, 227

Karkare, K. S., Ade, P. A. R., Ahmed, Z., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, ed. W. S. Holland & J. Zmuidzinas, 91533B

Keating, B. G., Shimon, M., & Yadav, A. P. S. 2013, The Astrophysical Journall, 762, L23

Kernasovskiy, S., Ade, P. A. R., Aikin, R. W., et al. 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8452, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI, ed. W. S. Holland & J. Zmuidzinas, 84521B

Kovac, J. M. 2004, PhD thesis, University of Chicago

Kovac, J. M., & Barkats, D. 2007, arXiv e-prints, arXiv:0707.1075

Lamb, J. W. 1996, Int. J. Infrared Millim. Waves, 17, 1997

Langley, S. P. 1880, Proceedings of the American Academy of Arts and Sciences, 16, 342

Lesurf, J. C. G. J. C. G. 1990, Millimetre-wave optics, devices, and systems (Bristol, England ; New York: A. Hilger)

Linde, A. D. 1982, Physics Letters B, 108, 389

Lorentz, H. A. 1881, Annalen der Physik, 248, 127

Lorenz, L. 1880, Annalen der Physik, 247, 70

Martin, D., & Puplett, E. 1970, Infrared Physics, 10, 105

Mather, J. C. 1982, Applied Optics, 21, 1125

Moncelsi, L., Ade, P. A. R., Ahmed, Z., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11453, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy X, ed. J. Zmuidzinas & J.-R. Gao, 1145314

Nadolski, A., Vieira, J. D., Sobrin, J. A., et al. 2020, Appl. Opt., 59, 3285

National Academies of Sciences, Engineering, and Medicine. 2022, Guidance on PFAS Exposure, Testing, and Clinical Follow-Up (Washington, DC: The National Academies Press), doi:10.17226/26156

Nguyen-Tuong, V. 1993, Journal of Vacuum Science & Technology A, 11, 1584

Paine, S. 2024, The am atmospheric model, doi:10.5281/zenodo.13748391

Paine, S. N., & Turner, D. D. 2013, IEEE Transactions on Geoscience and Remote Sensing, 51, 5187

Penzias, A. A., & Wilson, R. W. 1965, The Astrophysical Journal, 142, 419

Petroff, M. A., Ahmed, Z., Bock, J. J., et al. 2024, arXiv e-prints, arXiv:2406.13587

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, AA, 641, A1

Planck Collaboration, Akrami, Y., Arroja, F., et al. 2020b, A&A, 641, A9

—. 2020c, A&A, 641, A10

Planck Collaboration, Akrami, Y., Andersen, K. J., et al. 2020d, Annual Review of Astronomy and Astrophysics, 643, A42

Polder, D., & van Santeen, J. H. 1946, Physica, 12, 257

Pozar, D. M. 1997, Microwave engineering, 2nd edn. (New York: Wiley)

Rees, M. J. 1968, ApJ, 153, L1

Rubin, V. C., & Ford, Jr., W. K. 1970, ApJ, 159, 379

Ryden, B. 2016, Introduction to Cosmology (Cambridge University Press)

Schillaci, A., Ade, P. A. R., Ahmed, Z., et al. 2023, Journal of Low Temperature Physics, 213, 317

Schnabel, W. 2014, in Polymers and Electromagnetic Radiation (Weinheim, Germany: Wiley), 5–54

Shitvov, A., Savini, G., Hargrave, P. C., et al. 2022, in Proc. SPIE, Vol. 12190, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI, ed. J. Zmuidzinas & J.-R. Gao, 121900D

St. Germaine, M. T. 2021, PhD thesis, Harvard University

Staniszewski, Z., Aikin, R. W., Amiri, M., et al. 2012, Journal of Low Temperature Physics, 167, 827

Starobinsky, A. A. 1980, Physics Letters B, 91, 99

Stokes, G. G. 1851, Transactions of the Cambridge Philosophical Society, 9, 399

Takahashi, Y. D., Ade, P. A. R., Barkats, D., et al. 2010, ApJ, 711, 1141

Thornton, R. J., Ade, P. A. R., Aiola, S., et al. 2016, Astrophys. J., Suppl. Ser., 227, 21

Timoshenko, S. 1959, Theory of plates and shells, 2nd edn., Engineering societies monographs (New York: McGraw-Hill)

Vasile, C. 2005, Practical Guide to Polyethylene (Shawbury: Smithers Rapra)

Wietzke, S., Jansen, C., Reuter, M., et al. 2011, Journal of Molecular Structure, 1006, 41

Zahm, S., Bonde, J. P., Chiu, W. A., et al. 2024, The lancet oncology, 25, 16

Zhang, C. 2023, PhD thesis, California Institute of Technology

Zmuidzinas, J. 2003, Appl. Opt., 42, 4989